



Intel[®] Math Kernel Library

Reference Manual

Copyright © 1994-2003 Intel Corporation
All Rights Reserved
Issued in U.S.A.
Document Number: 630813-6004

World Wide Web: <http://developer.intel.com>

Intel® Math Kernel Library Reference Manual

Revision	Revision History	Date
-001	Original Issue.	11/94
-002	Added functions crotg, zrotg. Documented versions of functions ?her2k, ?symm, ?syrk, and ?syr2k not previously described. Pagination revised.	5/95
-003	Changed the title; former title: "Intel BLAS Library for the Pentium® Processor Reference Manual." Added functions ?rotm, ?rotmg and updated Appendix C.	1/96
-004	Documents Intel Math Kernel library release 2.0 with the parallelism capability. Information on parallelism has been added in Chapter 1 and in section "BLAS Level 3 Routines" in Chapter 2.	11/96
-005	Two-dimensional FFTs have been added. C interface has been added to both one- and two-dimensional FFTs.	8/97
-006	Documents Intel Math Kernel Library release 2.1. Sparse BLAS section has been added in Chapter 2.	1/98
-007	Documents Intel Math Kernel Library release 3.0. Descriptions of LAPACK routines (Chapters 4 and 5) and CBLAS interface (Appendix C) have been added. Quick Reference has been excluded from the manual; MKL 3.0 Quick Reference is now available in HTML format.	1/99
-008	Documents Intel Math Kernel Library release 3.2. Description of FFT routines have been revised. In Chapters 4 and 5 NAG names for LAPACK routines have been excluded.	6/99
-009	New LAPACK routines for eigenvalue problems have been added in chapter 5.	11/99
-010	Documents Intel Math Kernel Library release 4.0. Chapter 6 describing the VML functions has been added.	06/00
-011	Documents Intel Math Kernel Library release 5.1. LAPACK section has been extended to include the full list of computational and driver routines .	04/01
-6001	Documents Intel Math Kernel Library release 6.0 beta. New DFT interface (chapter 8) and Vector Statistical Library functions (chapter 7) have been added.	07/02
-6002	Documents Intel Math Kernel Library 6.0 beta update. DFT functions description (chapter 8) has been updated. CBLAS interface description was extended.	12/02
-6003	Documents Intel Math Kernel Library 6.0 gold. DFT functions have been updated. Auxiliary LAPACK routines' descriptions were added to the manual.	03/03
-6004	Documents Intel Math Kernel Library release 6.1.	07/03

This manual as well as the software described in it is furnished under license and may only be used or copied in accordance with the terms of the license. The information in this manual is furnished for informational use only, is subject to change without notice, and should not be construed as a commitment by Intel Corporation.

Intel Corporation assumes no responsibility or liability for any errors or inaccuracies that may appear in this document or any software that may be provided in association with this document. Except as permitted by such license, no part of this document may be reproduced, stored in a retrieval system, or transmitted in any form or by any means without the express written consent of Intel Corporation.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, or life sustaining applications. intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

Processors may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available upon request.

Celeron, Dialogic, i386, i486, iCOMP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Inside, Intel Inside logo, Intel NetBurst, Intel NetStructure, Intel Xeon, Intel XScale, Itanium, MMX, MMX logo, Pentium, Pentium II Xeon, Pentium III Xeon, and VTune are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

* Other names and brands may be claimed as the property of others.

Copyright © 1994-2003, Intel Corporation. All Rights Reserved.

Chapters 4 and 5 include derivative work portions that have been copyrighted:
© 1991, 1992, and 1998 by The Numerical Algorithms Group, Ltd.

Contents

Chapter 1 Overview

About This Software	1-1
Technical Support	1-2
BLAS Routines.....	1-2
Sparse BLAS Routines	1-3
Fast Fourier Transforms	1-3
LAPACK Routines	1-3
VML Functions	1-3
VSL Functions.....	1-4
DFT Functions	1-4
Performance Enhancements	1-4
Parallelism.....	1-5
Platforms Supported	1-5
About This Manual.....	1-6
Audience for This Manual	1-6
Manual Organization.....	1-6
Notational Conventions	1-8
Routine Name Shorthand	1-8
Font Conventions	1-8
Related Publications	1-9

Chapter 2 BLAS and Sparse BLAS Routines

Routine Naming Conventions	2-2
Matrix Storage Schemes	2-3
BLAS Level 1 Routines and Functions	2-4

?asum	2-5
?axpy	2-6
?copy	2-7
?dot	2-8
?dotc	2-9
?dotu	2-10
?nrm2	2-11
?rot	2-12
?rotg	2-14
?rotm	2-15
?rotmg	2-17
?scal	2-18
?swap	2-20
i?amax	2-21
i?amin	2-22
BLAS Level 2 Routines	2-23
?gbmv	2-24
?gemv	2-27
?ger	2-30
?gerc	2-31
?geru	2-33
?hbmv	2-35
?hemv	2-38
?her	2-40
?her2	2-42
?hpmv	2-44
?hpr	2-47
?hpr2	2-49
?sbmv	2-51
?spmv	2-54
?spr	2-56
?spr2	2-58

?symv	2-60
?syr	2-62
?syr2	2-64
?tbmv	2-66
?tbsv	2-69
?tpmv	2-72
?tpsv	2-75
?trmv	2-77
?trsv	2-79
BLAS Level 3 Routines	2-82
Symmetric Multiprocessing Version of Intel® MKL.....	2-82
?gemm	2-83
?hemm	2-86
?herk	2-89
?her2k	2-92
?symm	2-96
?syrk	2-100
?syr2k	2-103
?trmm	2-107
?trsm	2-110
Sparse BLAS Routines and Functions	2-114
Vector Arguments in Sparse BLAS	2-114
Naming Conventions in Sparse BLAS	2-115
Routines and Data Types in Sparse BLAS	2-115
BLAS Routines That Can Work With Sparse Vectors .	2-116
?axpyi	2-116
?doti	2-118
?dotci	2-119
?dotui	2-120
?gthr	2-121
?gthrz	2-122
?roti	2-123

?sctr 2-124

Chapter 3 Fast Fourier Transforms

One-dimensional FFTs	3-1
Data Storage Types	3-2
Data Structure Requirements	3-3
Complex-to-Complex One-dimensional FFTs.....	3-3
cfft1d/zfft1d	3-4
cfft1dc/zfft1dc	3-5
Real-to-Complex One-dimensional FFTs.....	3-7
scfft1d/dzfft1d	3-8
scfft1dc/dzfft1dc	3-10
Complex-to-Real One-dimensional FFTs.....	3-12
csfft1d/zdff1d	3-13
csfft1dc/zdff1dc	3-15
Two-dimensional FFTs.....	3-17
Complex-to-Complex Two-dimensional FFTs	3-18
cfft2d/zfft2d	3-19
cfft2dc/zfft2dc	3-20
Real-to-Complex Two-dimensional FFTs	3-21
scfft2d/dzfft2d	3-22
scfft2dc/dzfft2dc	3-24
Complex-to-Real Two-dimensional FFTs	3-27
csfft2d/zdff2d	3-28
csfft2dc/zdff2dc	3-29

Chapter 4 LAPACK Routines: Linear Equations

Routine Naming Conventions	4-2
Matrix Storage Schemes	4-3
Mathematical Notation	4-3
Error Analysis	4-4
Computational Routines	4-5
Routines for Matrix Factorization	4-7
?getrf	4-7

?gbtrf	4-10
?gttrf	4-12
?potrf	4-14
?pptrf	4-16
?pbtrf	4-18
?pttrf	4-20
?sytrf	4-22
?hetrf	4-25
?sprtf	4-28
?hptrf	4-31
Routines for Solving Systems of Linear Equations	4-33
?getrs	4-34
?gbtrs	4-36
?gttrs	4-38
?potrs	4-41
?pptrs	4-43
?pbtrs	4-45
?pttrs	4-47
?sytrs	4-49
?hetrs	4-51
?sptrs	4-53
?hptrs	4-55
?trtrs	4-57
?tptrs	4-59
?tbtrs	4-61
Routines for Estimating the Condition Number	4-63
?gecon	4-63
?gbcon	4-65
?gtcon	4-67
?pocon	4-70
?ppcon	4-72
?pbcon	4-74

?ptcon	4-76
?sycon	4-78
?hecon	4-80
?spcon	4-82
?hpcon	4-84
?trcon	4-86
?tpcon	4-88
?tbcon	4-90
Refining the Solution and Estimating Its Error	4-92
?gerfs	4-92
?gbrfs	4-95
?gtrfs	4-98
?porfs	4-101
?pprfs	4-104
?pbrfs	4-107
?ptrfs	4-110
?syrf	4-113
?herf	4-116
?sprf	4-119
?hprf	4-122
?trrf	4-124
?tprf	4-127
?tbrf	4-130
Routines for Matrix Inversion	4-133
?getri	4-133
?potri	4-135
?pptri	4-137
?sytri	4-139
?hetri	4-141
?sptri	4-143
?hptri	4-145
?trtri	4-147

?tptri	4-148
Routines for Matrix Equilibration	4-150
?geequ	4-150
?gbequ	4-153
?poequ	4-155
?ppequ	4-157
?pbequ	4-159
Driver Routines	4-161
?gesv	4-162
?gesvx	4-163
?gbsv	4-170
?gbsvx	4-172
?gtsv	4-179
?gtsvx	4-181
?posv	4-186
?posvx	4-188
?ppsv	4-193
?ppsvx	4-195
?pbsv	4-200
?pbsvx	4-202
?ptsv	4-207
?ptsvx	4-209
?sysv	4-212
?sysvx	4-215
?hesvx	4-220
?hesv	4-224
?spsv	4-227
?spsvx	4-230
?hpsvx	4-235
?hpsv	4-239

Chapter 5 LAPACK Routines: Least Squares and Eigenvalue Problems

Routine Naming Conventions	5-4
Matrix Storage Schemes	5-5
Mathematical Notation	5-5
Computational Routines	5-6
Orthogonal Factorizations	5-6
?geqrf	5-8
?geqpf	5-11
?geqp3	5-14
?orgqr	5-17
?ormqr	5-19
?ungqr	5-21
?unmqr	5-23
?gelqf	5-25
?orglq	5-28
?ormlq	5-30
?unglq	5-32
?unmlq	5-34
?geqlf	5-36
?orgql	5-38
?ungql	5-40
?ormql	5-42
?unmql	5-45
?gerqf	5-48
?orgrq	5-50
?ungrq	5-52
?ormrq	5-54
?unmrq	5-57
?tzzf	5-60
?ormrz	5-62
?unmrz	5-65

?ggqrf	5-68
?ggrqf	5-71
Singular Value Decomposition	5-74
?gebrd	5-76
?gbbnd	5-79
?orgbr	5-82
?ormbr	5-85
?ungbr	5-88
?unmbr	5-91
?bdsqr	5-94
?bdsdc	5-98
Symmetric Eigenvalue Problems	5-101
?sytrd	5-105
?orgtr	5-107
?ormtr	5-109
?hetrd	5-111
?ungtr	5-113
?unmtr	5-115
?sptrd	5-117
?opgtr	5-119
?opmtr	5-120
?hptrd	5-122
?upgtr	5-124
?upmtr	5-125
?sbtrd	5-128
?hbtrd	5-130
?sterf	5-132
?steqr	5-134
?stedc	5-137
?stegr	5-141
?pteqr	5-146
?stebz	5-149
?stein	5-152

?disna	5-154
Generalized Symmetric-Definite Eigenvalue Problems.	5-157
?sygst	5-158
?hegst	5-160
?spgst	5-162
?hpgst	5-164
?sbgst	5-166
?hbgst	5-169
?pbstf	5-172
Nonsymmetric Eigenvalue Problems	5-174
?gehrd	5-178
?orghr	5-180
?ormhr	5-182
?unghr	5-185
?unmhr	5-187
?gebal	5-190
?gebak	5-193
?hseqr	5-195
?hsein	5-199
?trevc	5-205
?trsna	5-210
?trexc	5-215
?trsen	5-217
?trsyl	5-222
Generalized Nonsymmetric Eigenvalue Problems	5-225
?gghrd	5-226
?ggbal	5-230
?ggbak	5-233
?hgeqz	5-235
?tgevc	5-242
?tgexc	5-247
?tgsen	5-250

?tgsyl	5-256
?tgsna	5-261
Generalized Singular Value Decomposition	5-266
?ggsvp	5-267
?tgsja	5-271
Driver Routines	5-278
Linear Least Squares (LLS) Problems	5-278
?gels	5-279
?gelsy	5-282
?gelss	5-286
?gelsd	5-289
Generalized LLS Problems.....	5-293
?gglse	5-293
?ggglm	5-296
Symmetric Eigenproblems.....	5-298
?syev	5-299
?heev	5-301
?syevd	5-303
?heevd	5-306
?syevx	5-309
?heevx	5-313
?syevr	5-317
?heevr	5-322
?spev	5-327
?hpev	5-329
?spevd	5-331
?hpevd	5-334
?spevx	5-338
?hpevx	5-342
?sbev	5-346
?hbev	5-348
?sbevd	5-350

?hbevd	5-353
?sbevx	5-357
?hbevx	5-361
?stev	5-365
?stevd	5-367
?stevx	5-370
?stevr	5-374
Nonsymmetric Eigenproblems	5-379
?gees	5-379
?geesx	5-384
?geev	5-390
?geevx	5-394
Singular Value Decomposition	5-400
?gesvd	5-400
?gesdd	5-405
?ggsvd	5-409
Generalized Symmetric Definite Eigenproblems.....	5-415
?sygv	5-416
?hegv	5-419
?sygvd	5-422
?hegvd	5-425
?sygvx	5-429
?hegvx	5-434
?spgv	5-439
?hpgv	5-442
?spgvd	5-445
?hpgvd	5-448
?spgvx	5-452
?hpgvx	5-456
?sbgv	5-460
?hbgv	5-463
?sbgvd	5-466

?hbgvd	5-469
?sbgvx	5-473
?hbgvx	5-477
Generalized Nonsymmetric Eigenproblems	5-482
?gges	5-482
?ggesx	5-489
?ggev	5-497
?ggevx	5-502
References.....	5-509

Chapter 6 LAPACK Auxiliary Routines

?lacgv	6-1
?lacrm.....	6-2
?lacrt.....	6-3
?laesy	6-4
?rot.....	6-6
?spmv	6-7
?spr.....	6-9
?symv	6-11
?syr	6-13
i?max1	6-15
ilaenv	6-16
lsame	6-18
lsamen	6-19
?sum1	6-20
?gbtf2.....	6-21
?gebd2.....	6-23
?gehd2.....	6-26
?gelq2.....	6-29
?geql2.....	6-31
?geqr2.....	6-33
?gerq2.....	6-35
?gesc2	6-37

?getc2	6-39
?getf2	6-40
?gtts2	6-42
?labad	6-43
?labrd	6-45
?lacon	6-48
?lacpy	6-50
?ladiv	6-51
?lae2	6-52
?laebz	6-54
?laed0	6-60
?laed1	6-64
?laed2	6-67
?laed3	6-70
?laed4	6-73
?laed5	6-74
?laed6	6-75
?laed7	6-78
?laed8	6-83
?laed9	6-87
?laeda	6-89
?laein	6-92
?laev2	6-95
?laexc	6-98
?lag2	6-100
?lags2	6-103
?lagtf	6-105
?lagtm	6-108
?lagts	6-110
?lagv2	6-113
?lahqr	6-115
?lahrd	6-118
?laic1	6-121

?lan2	6-124
?lals0	6-127
?lalsa	6-132
?lalsd	6-136
?lamch	6-139
?lamc1	6-140
?lamc2	6-141
?lamc3	6-142
?lamc4	6-143
?lamc5	6-144
?lamrg	6-145
?langb	6-146
?lange	6-148
?langt	6-149
?lanhs	6-151
?lansb	6-152
?lanhb	6-154
?lansp	6-156
?lanhp	6-158
?lanst/?lanht	6-159
?lansy	6-161
?lanhe	6-163
?lantb	6-164
?lantp	6-167
?lantr	6-169
?lanv2	6-171
?lapll	6-172
?lapmt	6-174
?lapy2	6-175
?lapy3	6-176
?laqgb	6-176
?laqge	6-179

?laqp2	6-181
?laqps	6-183
?laqsb	6-185
?laqsp	6-187
?laqsy	6-189
?laqtr	6-191
?lar1v	6-194
?lar2v	6-196
?larf	6-198
?larfb	6-200
?larfg	6-202
?larft	6-204
?larfx	6-207
?largv	6-208
?larnv	6-210
?larrb	6-212
?larre	6-214
?larrf	6-216
?larrv	6-218
?lartg	6-221
?lartv	6-223
?laruv	6-224
?larz	6-225
?larzb	6-227
?larzt	6-230
?las2	6-233
?lascl	6-234
?lasd0	6-236
?lasd1	6-238
?lasd2	6-241
?lasd3	6-246
?lasd4	6-249

?LASD5	6-251
?LASD6	6-252
?lasd7	6-257
?lasd8	6-262
?lasd9	6-264
?lasda	6-267
?lasdq	6-271
?lasdt	6-274
?laset	6-275
?lasq1	6-277
?lasq2	6-278
?lasq3	6-280
?lasq4	6-281
?lasq5	6-283
?lasq6	6-284
?lasr	6-286
?lasrt	6-288
?lassq	6-289
?lasv2	6-291
?laswp	6-292
?lasy2	6-294
?lasyf	6-296
?lahef	6-299
?latbs	6-301
?latdf	6-304
?latps	6-307
?latrd	6-309
?latrs	6-313
?latrz	6-317
?lauu2	6-320
?lauum	6-321
?org2l/?ung2l	6-323

?org2r/?ung2r.....	6-324
?orgl2/?ungl2	6-326
?orgr2/?ungr2.....	6-328
?orm2l/?unm2l	6-330
?orm2r/?unm2r.....	6-332
?orml2/?unml2	6-335
?ormr2/?unmr2.....	6-337
?ormr3/?unmr3.....	6-340
?pbf2	6-343
?potf2	6-345
?ptts2	6-346
?rscl.....	6-348
?sygs2/?hegs2	6-349
?sytd2/?hetd2.....	6-351
?sytf2.....	6-354
?hetf2	6-356
?tgex2.....	6-358
?tgsy2.....	6-361
?trti2	6-366
xerbla.....	6-367

Chapter 7 Vector Mathematical Functions

Data Types and Accuracy Modes	7-2
Function Naming Conventions	7-2
Functions Interface	7-3
Vector Indexing Methods	7-6
Error Diagnostics	7-6
VML Mathematical Functions	7-8
Inv	7-10
Div	7-11
Sqrt	7-12
InvSqrt	7-13
Cbrt	7-15

InvCbrt	7-16
Pow	7-17
Powx	7-18
Exp	7-20
Ln	7-21
Log10	7-22
Cos	7-23
Sin	7-24
SinCos	7-25
Tan	7-27
Acos	7-28
Asin	7-29
Atan	7-30
Atan2	7-31
Cosh	7-32
Sinh	7-34
Tanh	7-35
Acosh	7-36
Asinh	7-37
Atanh	7-39
Erf	7-40
VML Pack/Unpack Functions	7-42
Pack	7-42
Unpack	7-44
VML Service Functions	7-47
SetMode	7-47
GetMode	7-50
SetErrStatus	7-51
GetErrStatus	7-53
ClearErrStatus	7-54
SetErrorCallBack	7-55
GetErrorCallBack	7-57

ClearErrorCallBack	7-58
Chapter 8 Vector Generators of Statistical Distributions	
Conventions	8-2
Mathematical Notation	8-3
Naming Conventions	8-4
Basic Pseudorandom Generators.....	8-6
Random Streams	8-6
Data Types	8-7
Service Subroutines	8-7
NewStream	8-9
NewStreamEx	8-10
DeleteStream	8-11
CopyStream	8-12
LeapfrogStream	8-13
SkipAheadStream	8-16
GetStreamStateBrng	8-19
Pseudorandom Generators	8-20
Continuous Distributions	8-21
Uniform	8-21
Gaussian	8-23
Exponential	8-26
Laplace	8-28
Weibull	8-31
Cauchy	8-33
Rayleigh	8-36
Lognormal	8-38
Gumbel	8-41
Discrete Distributions	8-43
Uniform	8-44
UniformBits	8-46
Bernoulli	8-48
Geometric	8-50

Binomial	8-52
Hypergeometric	8-54
Poisson	8-56
NegBinomial	8-57
Advanced Service Subroutines	8-59
Data types	8-60
RegisterBrng	8-62
GetBrngProperties	8-63
Formats for User-Designed Generators	8-64
iBRng	8-67
sBRng	8-68
dBRng	8-69

Chapter 9 Advanced DFT Functions

Computing DFT	9-2
DFT Interface	9-9
Status Checking Functions	9-10
ErrorClass	9-11
ErrorMessage	9-13
Descriptor Manipulation	9-15
CreateDescriptor	9-15
CommitDescriptor	9-17
CopyDescriptor	9-18
FreeDescriptor	9-20
DFT Computation	9-21
ComputeForward	9-21
ComputeBackward	9-23
Descriptor Configuration	9-26
SetValue	9-31
GetValue	9-33
Configuration Settings	9-35
Precision of transform	9-35
Forward domain of transform	9-36

Transform dimension and lengths	9-36
Number of transforms	9-37
Sign and scale.....	9-37
Placement of result	9-38
Packed formats	9-38
Storage schemes	9-39
Input and output distances	9-48
Strides	9-49
Initialization Effort.....	9-52
Ordering	9-52
Workspace	9-54
Transposition.....	9-54
Appendix A Routine and Function Arguments	
Vector Arguments in BLAS	A-1
Vector Arguments in VML	A-3
Matrix Arguments	A-4
Appendix B Code Examples	
Appendix C CBLAS Interface to the BLAS	
CBLAS Arguments	1
Enumerated Types	2
Level 1 CBLAS	3
Level 2 CBLAS	6
Level 3 CBLAS	14
Sparse CBLAS	18
Glossary	
Index	

Overview

1

The Intel® Math Kernel Library (Intel® MKL) provides Fortran routines and functions that perform a wide variety of operations on vectors and matrices. The library also includes fast Fourier transform functions and new discrete Fourier transform functions, as well as vector mathematical and vector statistical functions with Fortran and C interfaces. The Intel MKL enhances performance of the programs that use it because the library has been optimized for Intel® processors.

This chapter introduces the Intel Math Kernel Library and provides information about the organization of this manual.

About This Software

The Intel Math Kernel Library includes the following groups of routines:

- Basic Linear Algebra Subprograms (BLAS):
 - vector operations
 - matrix-vector operations
 - matrix-matrix operations
- Sparse BLAS (basic vector operations on sparse vectors)
- Fast Fourier transform routines (with Fortran and C interfaces)
- LAPACK routines for solving systems of linear equations
- LAPACK routines for solving least-squares problems, eigenvalue and singular value problems, and Sylvester’s equations
- Auxiliary LAPACK routines
- Vector Mathematical Library (VML) functions for computing core mathematical functions on vector arguments (with Fortran and C interfaces)
- Vector Statistical Library (VSL) functions for generating vectors of pseudorandom numbers with different types of statistical distributions
- Advanced Discrete Fourier Transform Functions (DFT).

For specific issues on using the library, please refer to the *MKL Release Notes*.

Technical Support

Intel MKL provides a product web site that offers timely and comprehensive product information, including product features, white papers, and technical articles. For the latest information, check: <http://developer.intel.com/software/products/>

Intel also provides a support web site that contains a rich repository of self help information, including getting started tips, known product issues, product errata, license information, user forums, and more (visit <http://support.intel.com/support/>).

Registering your product entitles you to one year of technical support and product updates through Intel® Premier Support. Intel Premier Support is an interactive issue management and communication web site providing these services:

- Submit issues and review their status.
- Download product updates anytime of the day.

To register your product, contact Intel, or seek product support, please visit: <http://www.intel.com/software/products/support>

BLAS Routines

BLAS routines and functions are divided into the following groups according to the operations they perform:

- [BLAS Level 1 Routines and Functions](#) perform operations of both addition and reduction on vectors of data. Typical operations include scaling and dot products.
- [BLAS Level 2 Routines](#) perform matrix-vector operations, such as matrix-vector multiplication, rank-1 and rank-2 matrix updates, and solution of triangular systems.
- [BLAS Level 3 Routines](#) perform matrix-matrix operations, such as matrix-matrix multiplication, rank-k update, and solution of triangular systems.

Sparse BLAS Routines

[Sparse BLAS Routines and Functions](#) operate on sparse vectors (that is, vectors in which most of the elements are zeros). These routines perform vector operations similar to BLAS Level 1 routines. Sparse BLAS routines take advantage of vectors' sparsity: they allow you to store only non-zero elements of vectors.

Fast Fourier Transforms

[Fast Fourier Transforms](#) (FFTs) are used in digital signal processing and image processing and in partial differential equation (PDE) solvers. Combined with the BLAS routines, the FFTs contribute to the portability of the programs and provide a simplified interface between your program and the available library. To obtain more functionality and ease of use, consider also using the new DFT functions described in [Chapter 9](#).

LAPACK Routines

The Intel Math Kernel Library covers the full set of the LAPACK computational and driver routines. These routines can be divided into the following groups according to the operations they perform:

- Routines for solving systems of linear equations, factoring and inverting matrices, and estimating condition numbers (see [Chapter 4](#)).
- Routines for solving least-squares problems, eigenvalue and singular value problems, and Sylvester's equations (see [Chapter 5](#)).
- Auxiliary routines used to perform certain subtasks or common low-level computation (see [Chapter 6](#)).

VML Functions

VML functions (see [Chapter 7](#)) include a set of highly optimized implementations of certain computationally expensive core mathematical functions (power, trigonometric, exponential, hyperbolic etc.) that operate on real vector arguments.

VSL Functions

Vector Statistical Library (VSL) functions (see [Chapter 8](#)) include a set of pseudorandom number generator subroutines implementing basic continuous and discrete distributions. To provide best performance, VSL subroutines use calls to highly optimized Basic Random Number Generators and the library of vector mathematical functions, VML.

DFT Functions

The newly developed Discrete Fourier Transform functions (see [Chapter 9](#)) provide uniformity of DFT computation and combine functionality with ease of use. Both Fortran and C interface specification are given. Users are encouraged to migrate to the new interface in their application programs.

Performance Enhancements

The Intel Math Kernel Library has been optimized by exploiting both processor and system features and capabilities. Special care has been given to those routines that most profit from cache-management techniques. These especially include matrix-matrix operation routines such as `dgemm()`.

In addition, code optimization techniques have been applied to minimize dependencies of scheduling integer and floating-point units on the results within the processor.

The major optimization techniques used throughout the library include:

- Loop unrolling to minimize loop management costs.
- Blocking of data to improve data reuse opportunities.
- Copying to reduce chances of data eviction from cache.
- Data prefetching to help hide memory latency.
- Multiple simultaneous operations (for example, dot products in `dgemm`) to eliminate stalls due to arithmetic unit pipelines.
- Use of hardware features such as the SIMD arithmetic units, where appropriate.

These are techniques from which the arithmetic code benefits the most.

Parallelism

In addition to the performance enhancements discussed above, the Intel MKL offers performance gains through parallelism provided by the symmetric multiprocessing performance (SMP) feature. You can obtain improvements from SMP in the following ways:

- One way is based on user-managed threads in the program and further distribution of the operations over the threads based on data decomposition, domain decomposition, control decomposition, or some other parallelizing technique. Each thread can use any of the Intel MKL functions because the library has been designed to be thread-safe.
- Another method is to use the FFT and BLAS level 3 routines. They have been parallelized and require no alterations of your application to gain the performance enhancements of multiprocessing. Performance using multiple processors on the level 3 BLAS shows excellent scaling. Since the threads are called and managed within the library, the application does not need to be recompiled thread-safe (see also [BLAS Level 3 Routines](#) in Chapter 2).
- Yet another method is to use *tuned LAPACK routines*. Currently these include the single- and double precision flavors of routines for *QR* factorization of general matrices, triangular factorization of general and symmetric positive-definite matrices, solving systems of equations with such matrices, as well as solving symmetric eigenvalue problems.

For instructions on setting the number of available processors for the BLAS level 3 and LAPACK routines, see the *Release Notes*.

Platforms Supported

The Intel Math Kernel Library includes Fortran routines and functions optimized for Intel[®] processor-based computers running operating systems that support multiprocessing. In addition to the Fortran interface, the Intel MKL includes a C-language interface for the fast Fourier transform functions, new discrete Fourier transform API, as well as for the Vector Mathematical Library and Vector Statistical Library functions.

About This Manual

This manual describes the routines of the Intel Math Kernel Library. Each reference section describes a routine group consisting of routines used with four basic data types: single-precision real, double-precision real, single-precision complex, and double-precision complex.

Each routine group is introduced by its name, a short description of its purpose, and the calling sequence for each type of data with which each routine of the group is used. The following sections are also included:

Discussion	Describes the operation performed by routines of the group based on one or more equations. The data types of the arguments are defined in general terms for the group.
Input Parameters	Defines the data type for each parameter on entry, for example: a REAL for saxpy DOUBLE PRECISION for daxpy
Output Parameters	Lists resultant parameters on exit.

Audience for This Manual

The manual addresses programmers proficient in computational linear algebra and assumes a working knowledge of linear algebra and Fourier transform principles and vocabulary.

Manual Organization

The manual contains the following chapters and appendixes:

Chapter 1	Overview . Introduces the Intel Math Kernel Library software, provides information on manual organization, and explains notational conventions.
Chapter 2	BLAS and Sparse BLAS Routines . Provides descriptions of BLAS and Sparse BLAS functions and routines.

-
- Chapter 3 [Fast Fourier Transforms](#). Provides descriptions of fast Fourier transforms (FFT).
- Chapter 4 [LAPACK Routines: Linear Equations](#). Provides descriptions of LAPACK routines for solving systems of linear equations and performing a number of related computational tasks: triangular factorization, matrix inversion, estimating the condition number of matrices.
- Chapter 5 [LAPACK Routines: Least Squares and Eigenvalue Problems](#). Provides descriptions of LAPACK routines for solving least-squares problems, standard and generalized eigenvalue problems, singular value problems, and Sylvester's equations.
- Chapter 6 [LAPACK Auxiliary Routines](#). Describes auxiliary LAPACK routines that perform certain subtasks or common low-level computation.
- Chapter 7 [Vector Mathematical Functions](#). Provides descriptions of VML functions for computing elementary mathematical functions on vector arguments.
- Chapter 8 [Vector Generators of Statistical Distributions](#). Provides descriptions of VSL functions for generating vectors of pseudorandom numbers.
- Chapter 9 [Advanced DFT Functions](#). Describes new functions for computing the Discrete Fourier Transform.
- Appendix A [Routine and Function Arguments](#). Describes the major arguments of the BLAS routines and VML functions: vector and matrix arguments.
- Appendix B [Code Examples](#). Provides code examples of calling BLAS functions and routines.
- Appendix C [CBLAS Interface to the BLAS](#). Provides the C interface to the BLAS.

The manual also includes a [Glossary](#) and an [Index](#).

Notational Conventions

This manual uses the following notational conventions:

- Routine name shorthand (`?ungqr` instead of `cungqr/zungqr`).
- Font conventions used for distinction between the text and the code.

Routine Name Shorthand

For shorthand, character codes are represented by a question mark “?” in names of routine groups. The question mark is used to indicate any or all possible varieties of a function; for example:

`?swap` Refers to all four data types of the vector-vector
`?swap` routine: `sswap`, `dswap`, `cswap`, and `zswap`.

Font Conventions

The following font conventions are used:

`UPPERCASE COURIER` Data type used in the discussion of input and output parameters for Fortran interface. For example, `CHARACTER*1`.

`lowercase courier` Code examples:
`a(k+i,j) = matrix(i,j)`
and data types for C interface, for example, `const float*`

`lowercase courier mixed with UpperCase courier` Function names for C interface, for example, `vmlSetMode`

`lowercase courier italic` Variables in arguments and parameters discussion. For example, `incx`.

`*` Used as a multiplication symbol in code examples and equations and where required by the Fortran syntax.

Related Publications

For more information about the BLAS, Sparse BLAS, LAPACK, VML, VSL, and DFT routines, refer to the following publications:

- BLAS Level 1
C. Lawson, R. Hanson, D. Kincaid, and F. Krough. *Basic Linear Algebra Subprograms for Fortran Usage*, ACM Transactions on Mathematical Software, Vol.5, No.3 (September 1979) 308-325.
- BLAS Level 2
J. Dongarra, J. Du Croz, S. Hammarling, and R. Hanson. *An Extended Set of Fortran Basic Linear Algebra Subprograms*, ACM Transactions on Mathematical Software, Vol.14, No.1 (March 1988) 1-32.
- BLAS Level 3
J. Dongarra, J. DuCroz, I. Duff, and S. Hammarling. *A Set of Level 3 Basic Linear Algebra Subprograms*, ACM Transactions on Mathematical Software (December 1989).
- Sparse BLAS
D. Dodson, R. Grimes, and J. Lewis. *Sparse Extensions to the FORTRAN Basic Linear Algebra Subprograms*, ACM Transactions on Mathematical Software, Vol.17, No.2 (June 1991).
D. Dodson, R. Grimes, and J. Lewis. *Algorithm 692: Model Implementation and Test Package for the Sparse Basic Linear Algebra Subprograms*, ACM Transactions on Mathematical Software, Vol.17, No.2 (June 1991).
- LAPACK
E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Donagarrá, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide*, Third Edition, Society for Industrial and Applied Mathematics (SIAM), 1999.
G. Golub and C. Van Loan. *Matrix Computations*, Johns Hopkins University Press, 1989.
- VML
J.M.Muller. *Elementary functions: algorithms and implementation*, Birkhauser Boston, 1997.
IEEE Standard for Binary Floating-Point Arithmetic. ANSI/IEEE Std 754-1985.

- VSL
- [Bratley87] Bratley P., Fox B.L., and Schrage L.E. *A Guide to Simulation*. 2nd edition. Springer-Verlag, New York, 1987.
- [Coddington94] Coddington, P. D. *Analysis of Random Number Generators Using Monte Carlo Simulation*. Int. J. Mod. Phys. C-5, 547, 1994.
- [Gentle98] Gentle, James E. *Random Number Generation and Monte Carlo Methods*, Springer-Verlag New York, Inc., 1998.
- [L'Ecuyer94] L'Ecuyer, Pierre. *Uniform Random Number Generation*. Annals of Operations Research, 53, 77–120, 1994.
- [L'Ecuyer99] L'Ecuyer, Pierre. *Tables of Linear Congruential Generators of Different Sizes and Good Lattice Structure*. Mathematics of Computation, 68, 225, 249-260, 1999.
- [L'Ecuyer99a] L'Ecuyer, Pierre. *Good Parameter Sets for Combined Multiple Recursive Random Number Generators*. Operations Research, 47, 1, 159-164, 1999.
- [L'Ecuyer01] L'Ecuyer, Pierre. *Software for Uniform Random Number Generation: Distinguishing the Good and the Bad*. Proceedings of the 2001 Winter Simulation Conference, IEEE Press, 95–105, Dec. 2001.
- [Kirkpatrick81] Kirkpatrick, S., and Stoll, E. *A Very Fast Shift-Register Sequence Random Number Generator*. Journal of Computational Physics, V. 40. 517–526, 1981.
- [Knuth81] Knuth, Donald E. *The Art of Computer Programming, Volume 2, Seminumerical Algorithms*. 2nd edition, Addison-Wesley Publishing Company, Reading, Massachusetts, 1981.
- [NAG] NAG Numerical Libraries.
http://www.nag.co.uk/numeric/numerical_libraries.asp

- DFT

[1] E. Oran Brigham, *The Fast Fourier Transform and Its Applications*, Prentice Hall, New Jersey, 1988.

[2] Athanasios Papoulis, *The Fourier Integral and its Applications*, 2nd edition, McGraw-Hill, New York, 1984.

[3] Ping Tak Peter Tang, *DFTI, a New API for DFT: Motivation, Design, and Rationale*, July 2002.

[4] Charles Van Loan, *Computational Frameworks for the Fast Fourier Transform*, SIAM, Philadelphia, 1992

For a reference implementation of BLAS, sparse BLAS, and LAPACK packages (without platform-specific optimizations) visit www.netlib.org.

BLAS and Sparse BLAS Routines

2

This chapter contains descriptions of the BLAS and Sparse BLAS routines of the Intel[®] Math Kernel Library. The routine descriptions are arranged in four sections according to the BLAS level of operation:

- [BLAS Level 1 Routines and Functions](#) (vector-vector operations)
- [BLAS Level 2 Routines](#) (matrix-vector operations)
- [BLAS Level 3 Routines](#) (matrix-matrix operations)
- [Sparse BLAS Routines and Functions](#).

Each section presents the routine and function group descriptions in alphabetical order by routine or function group name; for example, the `?asum` group, the `?axpy` group. The question mark in the group name corresponds to different character codes indicating the data type (`s`, `d`, `c`, and `z` or their combination); see *Routine Naming Conventions* on the next page.

When BLAS routines encounter an error, they call the error reporting routine `XERBLA`. To be able to view error reports, you must include `XERBLA` in your code. A copy of the source code for `XERBLA` is included in the library.

In BLAS Level 1 groups `i?amax` and `i?amin`, an “i” is placed before the character code and corresponds to the index of an element in the vector. These groups are placed in the end of the BLAS Level 1 section.

Routine Naming Conventions

BLAS routine names have the following structure:

<character code> <name> <mod> ()

The *<character code>* is a character that indicates the data type:

s real, single precision *c* complex, single precision
d real, double precision *z* complex, double precision

Some routines and functions can have combined character codes, such as *sc* or *dz*. For example, the function *scasum* uses a complex input array and returns a real value.

The *<name>* field, in BLAS level 1, indicates the operation type. For example, the BLAS level 1 routines *?dot*, *?rot*, *?swap* compute a vector dot product, vector rotation, and vector swap, respectively.

In BLAS level 2 and 3, *<name>* reflects the matrix argument type:

ge general matrix
gb general band matrix
sy symmetric matrix
sp symmetric matrix (packed storage)
sb symmetric band matrix
he Hermitian matrix
hp Hermitian matrix (packed storage)
hb Hermitian band matrix
tr triangular matrix
tp triangular matrix (packed storage)
tb triangular band matrix.

The *<mod>* field, if present, provides additional details of the operation. BLAS level 1 names can have the following characters in the *<mod>* field:

c conjugated vector
u unconjugated vector
g Givens rotation.

BLAS level 2 names can have the following characters in the *<mod>* field:

mv matrix-vector product
sv solving a system of linear equations with matrix-vector operations
r rank-1 update of a matrix
r2 rank-2 update of a matrix.

BLAS level 3 names can have the following characters in the *<mod>* field:

- mm** matrix-matrix product
- sm** solving a system of linear equations with matrix-matrix operations
- rk** rank-*k* update of a matrix
- r2k** rank-2*k* update of a matrix.

The examples below illustrate how to interpret BLAS routine names:

- <d> <dot>** **ddot**: double-precision real vector-vector dot product
- <c> <dot> <c>** **cdotc**: complex vector-vector dot product, conjugated
- <sc> <asum>** **scasum**: sum of magnitudes of vector elements, single precision real output and single precision complex input
- <c> <dot> <u>** **cdotu**: vector-vector dot product, unconjugated, complex
- <s> <ge> <mv>** **sgemv**: matrix-vector product, general matrix, single precision
- <z> <tr> <mm>** **ztrmm**: matrix-matrix product, triangular matrix, double-precision complex.

Sparse BLAS naming conventions are similar to those of BLAS level 1.

For more information, see *Naming conventions in Sparse BLAS*.

Matrix Storage Schemes

Matrix arguments of BLAS routines can use the following storage schemes:

- *Full storage*: a matrix *A* is stored in a two-dimensional array **a**, with the matrix element a_{ij} stored in the array element **a(i, j)**.
- *Packed storage* scheme allows you to store symmetric, Hermitian, or triangular matrices more compactly: the upper or lower triangle of the matrix is packed by columns in a one-dimensional array.
- *Band storage*: a band matrix is stored compactly in a two-dimensional array: columns of the matrix are stored in the corresponding columns of the array, and *diagonals* of the matrix are stored in rows of the array.

For more information on matrix storage schemes, see [Matrix Arguments](#) in Appendix A.

BLAS Level 1 Routines and Functions

BLAS Level 1 includes routines and functions, which perform vector-vector operations. Table 2-1 lists the BLAS Level 1 routine and function groups and the data types associated with them.

Table 2-1 BLAS Level 1 Routine Groups and Their Data Types

Routine or Function Group	Data Types	Description
?asum	s, d, sc, dz	Sum of vector magnitudes (functions)
?axpy	s, d, c, z	Scalar-vector product (routines)
?copy	s, d, c, z	Copy vector (routines)
?dot	s, d	Dot product (functions)
?sdot	sd, d	Dot product with extended precision (functions)
?dotc	c, z	Dot product conjugated (functions)
?dotu	c, z	Dot product unconjugated (functions)
?nrm2	s, d, sc, dz	Vector 2-norm (Euclidean norm) a normal or null vector (functions)
?rot	s, d, cs, zd	Plane rotation of points (routines)
?rotg	s, d, c, z	Givens rotation of points (routines)
?rotm	s, d	Modified plane rotation of points
?rotmg	s, d	Givens modified plane rotation of points
?scal	s, d, c, z, cs, zd	Vector scaling (routines)
?swap	s, d, c, z	Vector-vector swap (routines)
i?amax	s, d, c, z	Vector maximum value, absolute largest element of a vector where <i>i</i> is an index to this value in the vector array (functions)
i?amin	s, d, c, z	Vector minimum value, absolute smallest element of a vector where <i>i</i> is an index to this value in the vector array (functions)

?asum

Computes the sum of magnitudes of the vector elements.

```
res = sasum ( n, x, incx )
res = scasum ( n, x, incx )
res = dasum ( n, x, incx )
res = dzasum ( n, x, incx )
```

Discussion

Given a vector x , ?asum functions compute the sum of the magnitudes of its elements or, for complex vectors, the sum of magnitudes of the elements' real parts plus magnitudes of their imaginary parts:

$$res = | \text{Re}x(1) | + | \text{Im}x(1) | + | \text{Re}x(2) | + | \text{Im}x(2) | + \dots + | \text{Re}x(n) | + | \text{Im}x(n) |$$

where x is a vector of order n .

Input Parameters

n **INTEGER**. Specifies the order of vector x .

x **REAL** for sasum
DOUBLE PRECISION for dasum
COMPLEX for scasum
DOUBLE COMPLEX for dzasum

 Array, **DIMENSION** at least $(1 + (n-1)*\text{abs}(incx))$.

$incx$ **INTEGER**. Specifies the increment for the elements of x .

Output Parameters

res **REAL** for sasum
DOUBLE PRECISION for dasum
REAL for scasum
DOUBLE PRECISION for dzasum

 Contains the sum of magnitudes of all elements' real parts plus magnitudes of their imaginary parts.

?axpy

Computes a vector-scalar product and adds the result to a vector.

```
call saxpy ( n, a, x, incx, y, incy )
call daxpy ( n, a, x, incx, y, incy )
call caxpy ( n, a, x, incx, y, incy )
call zaxpy ( n, a, x, incx, y, incy )
```

Discussion

The ?axpy routines perform a vector-vector operation defined as

$$y := a*x + y$$

where:

a is a scalar

x and *y* are vectors of order *n*.

Input Parameters

<i>n</i>	INTEGER. Specifies the order of vectors <i>x</i> and <i>y</i> .
<i>a</i>	REAL for saxpy DOUBLE PRECISION for daxpy COMPLEX for caxpy DOUBLE COMPLEX for zaxpy Specifies the scalar <i>a</i> .
<i>x</i>	REAL for saxpy DOUBLE PRECISION for daxpy COMPLEX for caxpy DOUBLE COMPLEX for zaxpy Array, DIMENSION at least $(1 + (n-1)*abs(incx))$.
<i>incx</i>	INTEGER. Specifies the increment for the elements of <i>x</i> .

y REAL for saxpy
 DOUBLE PRECISION for daxpy
 COMPLEX for caxpy
 DOUBLE COMPLEX for zaxpy
 Array, DIMENSION at least $(1 + (n-1)*abs(incy))$.
 $incy$ INTEGER. Specifies the increment for the elements of y .

Output Parameters

y Contains the updated vector y .

?copy

Copies vector to another vector.

```

call scopy ( n, x, incx, y, incy )
call dcopy ( n, x, incx, y, incy )
call ccopy ( n, x, incx, y, incy )
call zcopy ( n, x, incx, y, incy )
  
```

Discussion

The ?copy routines perform a vector-vector operation defined as

$$y = x$$

where x and y are vectors.

Input Parameters

n INTEGER. Specifies the order of vectors x and y .
 x REAL for scopy
 DOUBLE PRECISION for dcopy
 COMPLEX for ccopy
 DOUBLE COMPLEX for zcopy
 Array, DIMENSION at least $(1 + (n-1)*abs(incx))$.
 $incx$ INTEGER. Specifies the increment for the elements of x .

y REAL for `scopy`
 DOUBLE PRECISION for `dcopy`
 COMPLEX for `ccopy`
 DOUBLE COMPLEX for `zcopy`

Array, DIMENSION at least $(1 + (n-1)*abs(incy))$.

incy INTEGER. Specifies the increment for the elements of *y*.

Output Parameters

y Contains a copy of the vector *x* if *n* is positive.
 Otherwise, parameters are unaltered.

?dot

Computes a vector-vector dot product.

```
res = sdot ( n, x, incx, y, incy )
res = ddot ( n, x, incx, y, incy )
```

Discussion

The `?dot` functions perform a vector-vector reduction operation defined as

$$res = \sum (x*y)$$

where *x* and *y* are vectors.

Input Parameters

n INTEGER. Specifies the order of vectors *x* and *y*.

x REAL for `sdot`
 DOUBLE PRECISION for `ddot`

Array, DIMENSION at least $(1+(n-1)*abs(incx))$.

incx INTEGER. Specifies the increment for the elements of *x*.

y REAL for `sdot`
DOUBLE PRECISION for `ddot`
Array, DIMENSION at least $(1+(n-1)*abs(incy))$.

incy INTEGER. Specifies the increment for the elements of *y*.

Output Parameters

res REAL for `sdot`
DOUBLE PRECISION for `ddot`
Contains the result of the dot product of *x* and *y*, if *n* is positive. Otherwise, *res* contains 0.

?sdot

*Computes a vector-vector dot product
with extended precision.*

```
res = sdsdot ( n, sb, sx, incx, sy, incy )  
res = dsdot ( n, sx, incx, sy, incy )
```

Discussion

The `?sdot` functions compute the inner product of two vectors with extended precision. Both functions use extended precision accumulation of the intermediate results, but the function `sdsdot` outputs the final result in single precision, whereas the function `dsdot` outputs the double precision result. The function `sdsdot` also adds scalar value *sb* to the inner product.

Input Parameters

n INTEGER. Specifies the number of elements in the input vectors *sx* and *sy*.

sb REAL. Single precision scalar to be added to inner product (for the function `sdsdot` only).

<i>sx, sy</i>	REAL. Arrays, DIMENSION at least $(1+(n-1)*abs(incx))$ and $(1+(n-1)*abs(incy))$, respectively. Contain the input single precision vectors.
<i>incx</i>	INTEGER. Specifies the increment for the elements of <i>sx</i> .
<i>incy</i>	INTEGER. Specifies the increment for the elements of <i>sy</i> .

Output Parameters

<i>res</i>	REAL for <i>sdsdot</i> DOUBLE PRECISION for <i>dsdot</i> Contains the result of the dot product of <i>sx</i> and <i>sy</i> (with <i>sb</i> added for <i>sdsdot</i>), if <i>n</i> is positive. Otherwise, <i>res</i> contains <i>sb</i> for <i>sdsdot</i> and 0 for <i>dsdot</i> .
------------	--

?dotc

Computes a dot product of a conjugated vector with another vector.

```
res = cdotc ( n, x, incx, y, incy )
res = zdotc ( n, x, incx, y, incy )
```

Discussion

The *?dotc* functions perform a vector-vector operation defined as

$$res = \sum (conjg(x)*y)$$

where *x* and *y* are *n*-element vectors.

Input Parameters

<i>n</i>	INTEGER. Specifies the order of vectors <i>x</i> and <i>y</i> .
----------	---

`x` `COMPLEX` for `cdotc`
 `DOUBLE COMPLEX` for `zdotc`
 Array, `DIMENSION` at least $(1 + (n-1)*abs(incx))$.

`incx` `INTEGER`. Specifies the increment for the elements of `x`.

`y` `COMPLEX` for `cdotc`
 `DOUBLE COMPLEX` for `zdotc`
 Array, `DIMENSION` at least $(1 + (n-1)*abs(incy))$.

`incy` `INTEGER`. Specifies the increment for the elements of `y`.

Output Parameters

`res` `COMPLEX` for `cdotc`
 `DOUBLE COMPLEX` for `zdotc`
 Contains the result of the dot product of the conjugated `x` and unconjugated `y`, if `n` is positive. Otherwise, `res` contains 0.

?dotu

Computes a vector-vector dot product.

```
res = cdotu ( n, x, incx, y, incy )
res = zdotu ( n, x, incx, y, incy )
```

Discussion

The `?dotu` functions perform a vector-vector reduction operation defined as $res = \sum(x^*y)$ where `x` and `y` are `n`-element complex vectors.

Input Parameters

`n` `INTEGER`. Specifies the order of vectors `x` and `y`.

x COMPLEX for `cdotu`
 DOUBLE COMPLEX for `zdotu`

 Array, DIMENSION at least $(1 + (n-1)*abs(incx))$.

incx INTEGER. Specifies the increment for the elements of *x*.

y COMPLEX for `cdotu`
 DOUBLE COMPLEX for `zdotu`

 Array, DIMENSION at least $(1 + (n-1)*abs(incy))$.

incy INTEGER. Specifies the increment for the elements of *y*.

Output Parameters

res COMPLEX for `cdotu`
 DOUBLE COMPLEX for `zdotu`

 Contains the result of the dot product of *x* and *y*, if *n* is positive. Otherwise, *res* contains 0.

?nrm2

Computes the Euclidean norm of a vector.

```
res = snrm2 ( n, x, incx )  
res = dnrm2 ( n, x, incx )  
res = scnrm2 ( n, x, incx )  
res = dznrm2 ( n, x, incx )
```

Discussion

The ?nrm2 functions perform a vector reduction operation defined as

$$res = ||x||$$

where:

x is a vector

res is a value containing the Euclidean norm of the elements of *x*.

Input Parameters

n **INTEGER**. Specifies the order of vector **x**.

x **REAL** for **snrm2**
DOUBLE PRECISION for **dnorm2**
COMPLEX for **scnorm2**
DOUBLE COMPLEX for **dznrm2**

Array, **DIMENSION** at least $(1 + (n-1)*abs(incx))$.

incx **INTEGER**. Specifies the increment for the elements of **x**.

Output Parameters

res **REAL** for **snrm2**
DOUBLE PRECISION for **dnorm2**
REAL for **scnorm2**
DOUBLE PRECISION for **dznrm2**

Contains the Euclidean norm of the vector **x**.

?rot

Performs rotation of points in the plane.

```
call srot ( n, x, incx, y, incy, c, s )
call drot ( n, x, incx, y, incy, c, s )
call csrot ( n, x, incx, y, incy, c, s )
call zdrot ( n, x, incx, y, incy, c, s )
```

Discussion

Given two complex vectors **x** and **y**, each vector element of these vectors is replaced as follows:

$$x(i) = c*x(i) + s*y(i)$$

$$y(i) = c*y(i) - s*x(i)$$

Input Parameters

<i>n</i>	INTEGER. Specifies the order of vectors <i>x</i> and <i>y</i> .
<i>x</i>	REAL for <code>srot</code> DOUBLE PRECISION for <code>drrot</code> COMPLEX for <code>csrot</code> DOUBLE COMPLEX for <code>zdrot</code> Array, DIMENSION at least $(1 + (n-1)*abs(incx))$.
<i>incx</i>	INTEGER. Specifies the increment for the elements of <i>x</i> .
<i>y</i>	REAL for <code>srot</code> DOUBLE PRECISION for <code>drrot</code> COMPLEX for <code>csrot</code> DOUBLE COMPLEX for <code>zdrot</code> Array, DIMENSION at least $(1 + (n-1)*abs(incy))$.
<i>incy</i>	INTEGER. Specifies the increment for the elements of <i>y</i> .
<i>c</i>	REAL for <code>srot</code> DOUBLE PRECISION for <code>drrot</code> REAL for <code>csrot</code> DOUBLE PRECISION for <code>zdrot</code> A scalar.
<i>s</i>	REAL for <code>srot</code> DOUBLE PRECISION for <code>drrot</code> REAL for <code>csrot</code> DOUBLE PRECISION for <code>zdrot</code> A scalar.

Output Parameters

<i>x</i>	Each element is replaced by $c*x + s*y$.
<i>y</i>	Each element is replaced by $c*y - s*x$.

?rotg

Computes the parameters for a Givens rotation.

```
call srotg ( a, b, c, s )
call drotg ( a, b, c, s )
call crotg ( a, b, c, s )
call zrotg ( a, b, c, s )
```

Discussion

Given the cartesian coordinates (a, b) of a point p , these routines return the parameters a , b , c , and s associated with the Givens rotation that zeros the y -coordinate of the point.

Input Parameters

- a REAL for `srotg`
 DOUBLE PRECISION for `drotg`
 COMPLEX for `crotg`
 DOUBLE COMPLEX for `zrotg`
Provides the x -coordinate of the point p .
- b REAL for `srotg`
 DOUBLE PRECISION for `drotg`
 COMPLEX for `crotg`
 DOUBLE COMPLEX for `zrotg`
Provides the y -coordinate of the point p .

Output Parameters

- a Contains the parameter r associated with the Givens rotation.
- b Contains the parameter z associated with the Givens rotation.

<i>c</i>	REAL for <code>srotg</code> DOUBLE PRECISION for <code>drotg</code> REAL for <code>crotg</code> DOUBLE PRECISION for <code>zrotg</code> Contains the parameter <i>c</i> associated with the Givens rotation.
<i>s</i>	REAL for <code>srotg</code> DOUBLE PRECISION for <code>drotg</code> COMPLEX for <code>crotg</code> DOUBLE COMPLEX for <code>zrotg</code> Contains the parameter <i>s</i> associated with the Givens rotation.

?rotm

Performs rotation of points in the modified plane.

```
call srotm ( n, x, incx, y, incy, param )  
call drotm ( n, x, incx, y, incy, param )
```

Discussion

Given two complex vectors *x* and *y*, each vector element of these vectors is replaced as follows:

$$x(i) = H*x(i) + H*y(i)$$

$$y(i) = H*y(i) - H*x(i)$$

where:

H is a modified Givens transformation matrix whose values are stored in the *param(2)* through *param(5)* array. See discussion on the *param* argument.

Input Parameters

<i>n</i>	INTEGER. Specifies the order of vectors <i>x</i> and <i>y</i> .
<i>x</i>	REAL for srotm DOUBLE PRECISION for drotm Array, DIMENSION at least $(1 + (n-1)*abs(incx))$.
<i>incx</i>	INTEGER. Specifies the increment for the elements of <i>x</i> .
<i>y</i>	REAL for srotm DOUBLE PRECISION for drotm Array, DIMENSION at least $(1 + (n-1)*abs(incy))$.
<i>incy</i>	INTEGER. Specifies the increment for the elements of <i>y</i> .
<i>param</i>	REAL for srotm DOUBLE PRECISION for drotm Array, DIMENSION 5.

The elements of the *param* array are:

param(1) contains a switch, *flag*.

param(2-5) contain *h11*, *h21*, *h12*, and *h22*, respectively, the components of the array *H*.

Depending on the values of *flag*, the components of *H* are set as follows:

$$flag = -1.: H = \begin{bmatrix} h11 & h12 \\ h21 & h22 \end{bmatrix}$$

$$flag = 0.: H = \begin{bmatrix} 1. & h12 \\ h21 & 1. \end{bmatrix}$$

$$flag = 1.: H = \begin{bmatrix} h11 & 1. \\ -1. & h22 \end{bmatrix}$$

$$flag = -2.: H = \begin{bmatrix} 1. & 0. \\ 0. & 1. \end{bmatrix}$$

In the above cases, the matrix entries of 1., -1., and 0. are assumed based on the last three values of *flag* and are not actually loaded into the *param* vector.

Output Parameters

x	Each element is replaced by $h11*x + h12*y$.
y	Each element is replaced by $h21*x + h22*y$.
H	Givens transformation matrix updated.

?rotmg

Computes the modified parameters for a Givens rotation.

```
call srotmg ( d1, d2, x1, y1, param )
call drotmg ( d1, d2, x1, y1, param )
```

Discussion

Given cartesian coordinates $(x1, y1)$ of an input vector, these routines compute the components of a modified Givens transformation matrix H that zeros the y -component of the resulting vector:

$$\begin{bmatrix} x \\ 0 \end{bmatrix} = H \begin{bmatrix} x1 \\ y1 \end{bmatrix}$$

Input Parameters

$d1$	REAL for <code>srotmg</code> DOUBLE PRECISION for <code>drotmg</code> Provides the scaling factor for the x -coordinate of the input vector ($\text{sqrt}(d1)x1$).
$d2$	REAL for <code>srotmg</code> DOUBLE PRECISION for <code>drotmg</code> Provides the scaling factor for the y -coordinate of the input vector ($\text{sqrt}(d2)y1$).

x1 REAL for `srotmg`
 DOUBLE PRECISION for `drotmg`
 Provides the *x*-coordinate of the input vector.

y1 REAL for `srotmg`
 DOUBLE PRECISION for `drotmg`
 Provides the *y*-coordinate of the input vector.

Output Parameters

param REAL for `srotmg`
 DOUBLE PRECISION for `drotmg`
 Array, DIMENSION 5.
 The elements of the *param* array are:
param(1) contains a switch, *flag*.
param(2-5) contain *h11*, *h21*, *h12*, and *h22*,
 respectively, the components of the array *H*.

Depending on the values of *flag*, the components of *H* are set as follows:

$$flag = -1.: H = \begin{bmatrix} h11 & h12 \\ h21 & h22 \end{bmatrix}$$

$$flag = 0.: H = \begin{bmatrix} 1. & h12 \\ h21 & 1. \end{bmatrix}$$

$$flag = 1.: H = \begin{bmatrix} h11 & 1. \\ -1. & h22 \end{bmatrix}$$

$$flag = -2.: H = \begin{bmatrix} 1. & 0. \\ 0. & 1. \end{bmatrix}$$

In the above cases, the matrix entries of 1., -1., and 0. are assumed based on the last three values of *flag* and are not actually loaded into the *param* vector.

?scal

Computes a vector by a scalar product.

```
call sscal ( n, a, x, incx )
call dscal ( n, a, x, incx )
call cscal ( n, a, x, incx )
call zscal ( n, a, x, incx )
call csscal ( n, a, x, incx )
call zdscal ( n, a, x, incx )
```

Discussion

The ?scal routines perform a vector-vector operation defined as

$$x = a*x$$

where:

a is a scalar, *x* is an *n*-element vector.

Input Parameters

<i>n</i>	INTEGER. Specifies the order of vector <i>x</i> .
<i>a</i>	REAL for sscal and csscal DOUBLE PRECISION for dscal and zdscal COMPLEX for cscal DOUBLE COMPLEX for zscal Specifies the scalar <i>a</i> .
<i>x</i>	REAL for sscal DOUBLE PRECISION for dscal COMPLEX for cscal and csscal DOUBLE COMPLEX for zscal and csscal Array, DIMENSION at least $(1 + (n-1)*abs(incx))$.
<i>incx</i>	INTEGER. Specifies the increment for the elements of <i>x</i> .

Output Parameters

x Overwritten by the updated vector **x**.

?swap

Swaps a vector with another vector.

```
call sswap ( n, x, incx, y, incy )
call dswap ( n, x, incx, y, incy )
call cswap ( n, x, incx, y, incy )
call zswap ( n, x, incx, y, incy )
```

Discussion

Given the two complex vectors **x** and **y**, the **?swap** routines return vectors **y** and **x** swapped, each replacing the other.

Input Parameters

n INTEGER. Specifies the order of vectors **x** and **y**.

x REAL for **sswap**
 DOUBLE PRECISION for **dswap**
 COMPLEX for **cswap**
 DOUBLE COMPLEX for **zswap**
 Array, DIMENSION at least $(1 + (n-1)*abs(incx))$.

incx INTEGER. Specifies the increment for the elements of **x**.

y REAL for **sswap**
 DOUBLE PRECISION for **dswap**
 COMPLEX for **cswap**
 DOUBLE COMPLEX for **zswap**
 Array, DIMENSION at least $(1 + (n-1)*abs(incy))$.

incy INTEGER. Specifies the increment for the elements of **y**.

Output Parameters

<code>x</code>	Contains the resultant vector <code>x</code> .
<code>y</code>	Contains the resultant vector <code>y</code> .

i?amax

Finds the element of a vector that has the largest absolute value.

```

index = isamax ( n, x, incx )
index = idamax ( n, x, incx )
index = icamax ( n, x, incx )
index = izamax ( n, x, incx )

```

Discussion

Given a vector `x`, the `i?amax` functions return the position of the vector element `x(i)` that has the largest absolute value or, for complex flavors, the position of the element with the largest sum $| \text{Re } x(i) | + | \text{Im } x(i) |$.

If `n` is not positive, 0 is returned.

If more than one vector element is found with the same largest absolute value, the index of the first one encountered is returned.

Input Parameters

<code>n</code>	INTEGER. Specifies the order of the vector <code>x</code> .
<code>x</code>	REAL for <code>isamax</code> DOUBLE PRECISION for <code>idamax</code> COMPLEX for <code>icamax</code> DOUBLE COMPLEX for <code>izamax</code> Array, DIMENSION at least $(1+(n-1)*\text{abs}(incx))$.
<code>incx</code>	INTEGER. Specifies the increment for the elements of <code>x</code> .

Output Parameters

index **INTEGER**. Contains the position of vector element x that has the largest absolute value.

i?amin

Finds the element of a vector that has the smallest absolute value.

```
index = isamin ( n, x, incx )
index = idamin ( n, x, incx )
index = icamin ( n, x, incx )
index = izamin ( n, x, incx )
```

Discussion

Given a vector x , the *i?amin* functions return the position of the vector element $x(i)$ that has the smallest absolute value or, for complex flavors, the position of the element with the smallest sum $|\text{Re}x(i)| + |\text{Im}x(i)|$.

If n is not positive, 0 is returned.

If more than one vector element is found with the same smallest absolute value, the index of the first one encountered is returned.

Input Parameters

n **INTEGER**. On entry, n specifies the order of the vector x .

x **REAL** for *isamin*
DOUBLE PRECISION for *idamin*
COMPLEX for *icamin*
DOUBLE COMPLEX for *izamin*

Array, **DIMENSION** at least $(1+(n-1)*\text{abs}(\text{incx}))$.

incx **INTEGER**. Specifies the increment for the elements of x .

Output Parameters

index **INTEGER**. Contains the position of vector element **x** that has the smallest absolute value.

BLAS Level 2 Routines

This section describes BLAS Level 2 routines, which perform matrix-vector operations. Table 2-2 lists the BLAS Level 2 routine groups and the data types associated with them.

Table 2-2 BLAS Level 2 Routine Groups and Their Data Types

Routine Groups	Data Types	Description
?gbmv	s, d, c, z	Matrix-vector product using a general band matrix
?gemv	s, d, c, z	Matrix-vector product using a general matrix
?ger	s, d	Rank-1 update of a general matrix
?gerc	c, z	Rank-1 update of a conjugated general matrix
?geru	c, z	Rank-1 update of a general matrix, unconjugated
?hbmV	c, z	Matrix-vector product using a Hermitian band matrix
?hemv	c, z	Matrix-vector product using a Hermitian matrix
?her	c, z	Rank-1 update of a Hermitian matrix
?her2	c, z	Rank-2 update of a Hermitian matrix
?hpmv	c, z	Matrix-vector product using a Hermitian packed matrix
?hpr	c, z	Rank-1 update of a Hermitian packed matrix
?hpr2	c, z	Rank-2 update of a Hermitian packed matrix
?sbmv	s, d	Matrix-vector product using symmetric band matrix
?spmv	s, d	Matrix-vector product using a symmetric packed matrix
?spr	s, d	Rank-1 update of a symmetric packed matrix
?spr2	s, d	Rank-2 update of a symmetric packed matrix
?symv	s, d	Matrix-vector product using a symmetric matrix
?syr	s, d	Rank-1 update of a symmetric matrix
?syr2	s, d	Rank-2 update of a symmetric matrix

continued *

Table 2-2 BLAS Level 2 Routine Groups and Their Data Types (continued)

Routine Groups	Data Types	Description
?tbmv	s, d, c, z	Matrix-vector product using a triangular band matrix
?tbsv	s, d, c, z	Linear solution of a triangular band matrix
?tpmv	s, d, c, z	Matrix-vector product using a triangular packed matrix
?tpsv	s, d, c, z	Linear solution of a triangular packed matrix
?trmv	s, d, c, z	Matrix-vector product using a triangular matrix
?trsv	s, d, c, z	Linear solution of a triangular matrix

[?gbmv](#)

Computes a matrix-vector product using a general band matrix

```
call sgbmv ( trans, m, n, kl, ku, alpha, a, lda, x, incx,
            beta, y, incy )
call dgbmv ( trans, m, n, kl, ku, alpha, a, lda, x, incx,
            beta, y, incy )
call cgbmv ( trans, m, n, kl, ku, alpha, a, lda, x, incx,
            beta, y, incy )
call zgbmv ( trans, m, n, kl, ku, alpha, a, lda, x, incx,
            beta, y, incy )
```

Discussion

The [?gbmv](#) routines perform a matrix-vector operation defined as

$$y := \alpha * a * x + \beta * y$$

or

$$y := \alpha * a' * x + \beta * y,$$

or

$$y := \alpha * \text{conjg}(a') * x + \beta * y,$$

where:

alpha and *beta* are scalars

x and *y* are vectors

a is an *m* by *n* band matrix, with *kl* sub-diagonals and *ku* super-diagonals.

Input Parameters

trans CHARACTER*1. Specifies the operation to be performed, as follows:

<i>trans</i> value	Operation to be Performed
N or n	$y := \alpha * a * x + \beta * y$
T or t	$y := \alpha * a' * x + \beta * y$
C or c	$y := \alpha * \text{conjg}(a') * x + \beta * y$

m INTEGER. Specifies the number of rows of the matrix *a*. The value of *m* must be at least zero.

n INTEGER. Specifies the number of columns of the matrix *a*. The value of *n* must be at least zero.

kl INTEGER. Specifies the number of sub-diagonals of the matrix *a*. The value of *kl* must satisfy $0 \leq kl$.

ku INTEGER. Specifies the number of super-diagonals of the matrix *a*. The value of *ku* must satisfy $0 \leq ku$.

alpha REAL for *sgbmv*
 DOUBLE PRECISION for *dgbmv*
 COMPLEX for *cgbmv*
 DOUBLE COMPLEX for *zgbmv*
 Specifies the scalar *alpha*.

a REAL for *sgbmv*
 DOUBLE PRECISION for *dgbmv*
 COMPLEX for *cgbmv*
 DOUBLE COMPLEX for *zgbmv*

Array, `DIMENSION (lda, n)`. Before entry, the leading $(kl + ku + 1)$ by n part of the array `a` must contain the matrix of coefficients. This matrix must be supplied column-by-column, with the leading diagonal of the matrix in row $(ku + 1)$ of the array, the first super-diagonal starting at position 2 in row ku , the first sub-diagonal starting at position 1 in row $(ku + 2)$, and so on. Elements in the array `a` that do not correspond to elements in the band matrix (such as the top left ku by ku triangle) are not referenced.

The following program segment transfers a band matrix from conventional full matrix storage to band storage:

```
do 20, j = 1, n
  k = ku + 1 - j
  do 10, i = max(1, j-ku), min(m, j+kl)
    a(k+i, j) = matrix(i,j)
  10 continue
20 continue
```

`lda` **INTEGER**. Specifies the first dimension of `a` as declared in the calling (sub)program. The value of `lda` must be at least $(kl + ku + 1)$.

`x` **REAL** for `sgbmv`
DOUBLE PRECISION for `dgbmv`
COMPLEX for `cgbmv`
DOUBLE COMPLEX for `zgbmv`

Array, `DIMENSION` at least $(1 + (n - 1) * \text{abs}(\text{incx}))$ when `trans = 'N'` or `'n'` and at least $(1 + (m - 1) * \text{abs}(\text{incx}))$ otherwise. Before entry, the incremented array `x` must contain the vector `x`.

`incx` **INTEGER**. Specifies the increment for the elements of `x`. `incx` must not be zero.

`beta` **REAL** for `sgbmv`
DOUBLE PRECISION for `dgbmv`
COMPLEX for `cgbmv`
DOUBLE COMPLEX for `zgbmv`

Specifies the scalar beta. When *beta* is supplied as zero, then *y* need not be set on input.

y REAL for *sgbmv*
 DOUBLE PRECISION for *dgbmv*
 COMPLEX for *cgbmv*
 DOUBLE COMPLEX for *zgbmv*

Array, DIMENSION at least $(1 + (m - 1) * \text{abs}(\text{incy}))$ when *trans* = 'N' or 'n' and at least $(1 + (n - 1) * \text{abs}(\text{incy}))$ otherwise. Before entry, the incremented array *y* must contain the vector *y*.

incy INTEGER. Specifies the increment for the elements of *y*. The value of *incy* must not be zero.

Output Parameters

y Overwritten by the updated vector *y*.

?gemv

Computes a matrix-vector product using a general matrix

```
call sgemv ( trans, m, n, alpha, a, lda, x, incx, beta,
            y, incy )
call dgemv ( trans, m, n, alpha, a, lda, x, incx, beta,
            y, incy )
call cgemv ( trans, m, n, alpha, a, lda, x, incx, beta,
            y, incy )
call zgemv ( trans, m, n, alpha, a, lda, x, incx, beta,
            y, incy )
```

Discussion

The ?gemv routines perform a matrix-vector operation defined as

$$y := \text{alpha} * a * x + \text{beta} * y,$$

or

$$y := \alpha * a' * x + \beta * y,$$

or

$$y := \alpha * \text{conjg}(a') * x + \beta * y,$$

where:

α and β are scalars

x and y are vectors

a is an m by n matrix.

Input Parameters

trans CHARACTER*1. Specifies the operation to be performed, as follows:

<i>trans</i> value	Operation to be Performed
N or n	$y := \alpha * a * x + \beta * y$
T or t	$y := \alpha * a' * x + \beta * y$
C or c	$y := \alpha * \text{conjg}(a') * x + \beta * y$

m INTEGER. Specifies the number of rows of the matrix a . m must be at least zero.

n INTEGER. Specifies the number of columns of the matrix a . The value of n must be at least zero.

alpha REAL for *sgemv*
 DOUBLE PRECISION for *dgemv*
 COMPLEX for *cgemv*
 DOUBLE COMPLEX for *zgemv*
 Specifies the scalar α .

a REAL for *sgemv*
 DOUBLE PRECISION for *dgemv*
 COMPLEX for *cgemv*
 DOUBLE COMPLEX for *zgemv*

Array, **DIMENSION** (*lda*, *n*). Before entry, the leading *m* by *n* part of the array *a* must contain the matrix of coefficients.

lda **INTEGER**. Specifies the first dimension of *a* as declared in the calling (sub)program. The value of *lda* must be at least $\max(1, m)$.

x **REAL** for *sgemv*
DOUBLE PRECISION for *dgemv*
COMPLEX for *cgemv*
DOUBLE COMPLEX for *zgemv*

Array, **DIMENSION** at least $(1 + (n - 1) * \text{abs}(\text{incx}))$ when *trans* = 'N' or 'n' and at least $(1 + (m - 1) * \text{abs}(\text{incx}))$ otherwise. Before entry, the incremented array *x* must contain the vector *x*.

incx **INTEGER**. Specifies the increment for the elements of *x*. The value of *incx* must not be zero.

beta **REAL** for *sgemv*
DOUBLE PRECISION for *dgemv*
COMPLEX for *cgemv*
DOUBLE COMPLEX for *zgemv*

Specifies the scalar *beta*. When *beta* is supplied as zero, then *y* need not be set on input.

y **REAL** for *sgemv*
DOUBLE PRECISION for *dgemv*
COMPLEX for *cgemv*
DOUBLE COMPLEX for *zgemv*

Array, **DIMENSION** at least $(1 + (m - 1) * \text{abs}(\text{incy}))$ when *trans* = 'N' or 'n' and at least $(1 + (n - 1) * \text{abs}(\text{incy}))$ otherwise. Before entry with *beta* non-zero, the incremented array *y* must contain the vector *y*.

incy **INTEGER**. Specifies the increment for the elements of *y*. The value of *incy* must not be zero.

Output Parameters

y Overwritten by the updated vector *y*.

?ger

Performs a rank-1 update of a general matrix.

```
call sger ( m, n, alpha, x, incx, y, incy, a, lda )
call dger ( m, n, alpha, x, incx, y, incy, a, lda )
```

Discussion

The ?ger routines perform a matrix-vector operation defined as

$$a := \alpha * x * y' + a,$$

where:

alpha is a scalar

x is an *m*-element vector

y is an *n*-element vector

a is an *m* by *n* matrix.

Input Parameters

m INTEGER. Specifies the number of rows of the matrix *a*. The value of *m* must be at least zero.

n INTEGER. Specifies the number of columns of the matrix *a*. The value of *n* must be at least zero.

alpha REAL for sger
DOUBLE PRECISION for dger
Specifies the scalar *alpha*.

x REAL for sger
DOUBLE PRECISION for dger

	Array, DIMENSION at least $(1 + (m - 1) * \text{abs}(incx))$. Before entry, the incremented array x must contain the m -element vector x .
incx	INTEGER . Specifies the increment for the elements of x . The value of incx must not be zero.
y	REAL for sger DOUBLE PRECISION for dger
	Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(incy))$. Before entry, the incremented array y must contain the n -element vector y .
incy	INTEGER . Specifies the increment for the elements of y . The value of incy must not be zero.
a	REAL for sger DOUBLE PRECISION for dger
	Array, DIMENSION (lda , n). Before entry, the leading m by n part of the array a must contain the matrix of coefficients.
lda	INTEGER . Specifies the first dimension of a as declared in the calling (sub)program. The value of lda must be at least $\text{max}(1, m)$.

Output Parameters

a	Overwritten by the updated matrix.
----------	------------------------------------

?gerc

*Performs a rank-1 update (conjugated)
of a general matrix.*

```
call cgerc ( m, n, alpha, x, incx, y, incy, a, lda )
call zgerc ( m, n, alpha, x, incx, y, incy, a, lda )
```

Discussion

The `?gerc` routines perform a matrix-vector operation defined as

$$a := \alpha * x * \text{conjg}(y') + a,$$

where:

`alpha` is a scalar

`x` is an m -element vector

`y` is an n -element vector

`a` is an m by n matrix.

Input Parameters

<code>m</code>	INTEGER. Specifies the number of rows of the matrix <code>a</code> . The value of <code>m</code> must be at least zero.
<code>n</code>	INTEGER. Specifies the number of columns of the matrix <code>a</code> . The value of <code>n</code> must be at least zero.
<code>alpha</code>	SINGLE PRECISION COMPLEX for <code>cgerc</code> DOUBLE PRECISION COMPLEX for <code>zgerc</code> Specifies the scalar <code>alpha</code> .
<code>x</code>	SINGLE PRECISION COMPLEX for <code>cgerc</code> DOUBLE PRECISION COMPLEX for <code>zgerc</code> Array, DIMENSION at least $(1 + (m - 1) * \text{abs}(\text{incx}))$. Before entry, the incremented array <code>x</code> must contain the m -element vector <code>x</code> .
<code>incx</code>	INTEGER. Specifies the increment for the elements of <code>x</code> . The value of <code>incx</code> must not be zero.
<code>y</code>	COMPLEX for <code>cgerc</code> DOUBLE COMPLEX for <code>zgerc</code> Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(\text{incy}))$. Before entry, the incremented array <code>y</code> must contain the n -element vector <code>y</code> .
<code>incy</code>	INTEGER. Specifies the increment for the elements of <code>y</code> . The value of <code>incy</code> must not be zero.

a **COMPLEX** for *cgerc*
 DOUBLE COMPLEX for *zgerc*
Array, **DIMENSION** (*lda*, *n*). Before entry, the leading *m* by *n* part of the array *a* must contain the matrix of coefficients.

lda **INTEGER**. Specifies the first dimension of *a* as declared in the calling (sub)program. The value of *lda* must be at least $\max(1, m)$.

Output Parameters

a Overwritten by the updated matrix.

?geru

*Performs a rank-1 update
(unconjugated) of a general matrix.*

```
call cgeru ( m, n, alpha, x, incx, y, incy, a, lda )  
call zgeru ( m, n, alpha, x, incx, y, incy, a, lda )
```

Discussion

The ?geru routines perform a matrix-vector operation defined as

*a := alpha*x*y' + a,*

where:

alpha is a scalar

x is an *m*-element vector

y is an *n*-element vector

a is an *m* by *n* matrix.

Input Parameters

<i>m</i>	INTEGER . Specifies the number of rows of the matrix <i>a</i> . The value of <i>m</i> must be at least zero.
<i>n</i>	INTEGER . Specifies the number of columns of the matrix <i>a</i> . The value of <i>n</i> must be at least zero.
<i>alpha</i>	COMPLEX for <i>cgeru</i> DOUBLE COMPLEX for <i>zgeru</i> Specifies the scalar <i>alpha</i> .
<i>x</i>	COMPLEX for <i>cgeru</i> DOUBLE COMPLEX for <i>zgeru</i> Array, DIMENSION at least $(1 + (m - 1) * \text{abs}(\text{incx}))$. Before entry, the incremented array <i>x</i> must contain the <i>m</i> -element vector <i>x</i> .
<i>incx</i>	INTEGER . Specifies the increment for the elements of <i>x</i> . The value of <i>incx</i> must not be zero.
<i>y</i>	COMPLEX for <i>cgeru</i> DOUBLE COMPLEX for <i>zgeru</i> Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(\text{incy}))$. Before entry, the incremented array <i>y</i> must contain the <i>n</i> -element vector <i>y</i> .
<i>incy</i>	INTEGER . Specifies the increment for the elements of <i>y</i> . The value of <i>incy</i> must not be zero.
<i>a</i>	COMPLEX for <i>cgeru</i> DOUBLE COMPLEX for <i>zgeru</i> Array, DIMENSION (<i>lda</i> , <i>n</i>). Before entry, the leading <i>m</i> by <i>n</i> part of the array <i>a</i> must contain the matrix of coefficients.
<i>lda</i>	INTEGER . Specifies the first dimension of <i>a</i> as declared in the calling (sub)program. The value of <i>lda</i> must be at least $\text{max}(1, m)$.

Output Parameters

<i>a</i>	Overwritten by the updated matrix.
----------	------------------------------------

?hbmv

Computes a matrix-vector product using a Hermitian band matrix.

```
call chbmv ( uplo, n, k, alpha, a, lda, x, incx, beta, y,
            incy )
call zhbmv ( uplo, n, k, alpha, a, lda, x, incx, beta, y,
            incy )
```

Discussion

The ?hbmv routines perform a matrix-vector operation defined as

$$y := \alpha * a * x + \beta * y,$$

where:

alpha and *beta* are scalars

x and *y* are *n*-element vectors

a is an *n* by *n* Hermitian band matrix, with *k* super-diagonals.

Input Parameters

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the band matrix *a* is being supplied, as follows:

<i>uplo</i> value	Part of Matrix <i>a</i> Supplied
U or u	The upper triangular part of matrix <i>a</i> is being supplied.
L or l	The lower triangular part of matrix <i>a</i> is being supplied.

n INTEGER. Specifies the order of the matrix *a*. The value of *n* must be at least zero.

k INTEGER. Specifies the number of super-diagonals of the matrix *a*. The value of *k* must satisfy $0 \leq k$.

alpha COMPLEX for `chbmv`
 DOUBLE COMPLEX for `zhbmv`

Specifies the scalar *alpha*.

a COMPLEX for `chbmv`
 DOUBLE COMPLEX for `zhbmv`

Array, DIMENSION (*lda*, *n*). Before entry with *uplo* = 'U' or 'u', the leading (*k* + 1) by *n* part of the array *a* must contain the upper triangular band part of the Hermitian matrix. The matrix must be supplied column-by-column, with the leading diagonal of the matrix in row (*k* + 1) of the array, the first super-diagonal starting at position 2 in row *k*, and so on. The top left *k* by *k* triangle of the array *a* is not referenced.

The following program segment transfers the upper triangular part of a Hermitian band matrix from conventional full matrix storage to band storage:

```
do 20, j = 1, n
  m = k + 1 - j
  do 10, i = max(1, j - k), j
    a(m + i, j) = matrix(i, j)
  10 continue
20 continue
```

Before entry with *uplo* = 'L' or 'l', the leading (*k* + 1) by *n* part of the array *a* must contain the lower triangular band part of the Hermitian matrix, supplied column-by-column, with the leading diagonal of the matrix in row 1 of the array, the first sub-diagonal starting at position 1 in row 2, and so on. The bottom right *k* by *k* triangle of the array *a* is not referenced.

The following program segment transfers the lower triangular part of a Hermitian band matrix from conventional full matrix storage to band storage:

```
do 20, j = 1, n
  m = 1 - j
  do 10, i = j, min( n, j + k )
```

```

      a( m + i, j ) = matrix( i, j )
      10 continue
      20 continue

```

The imaginary parts of the diagonal elements need not be set and are assumed to be zero.

- lda* **INTEGER**. Specifies the first dimension of *a* as declared in the calling (sub)program. The value of *lda* must be at least $(k + 1)$.
- x* **COMPLEX** for **chbmv**
DOUBLE COMPLEX for **zhbmv**
- Array, **DIMENSION** at least $(1 + (n - 1) * \text{abs}(incx))$. Before entry, the incremented array *x* must contain the vector *x*.
- incx* **INTEGER**. Specifies the increment for the elements of *x*. The value of *incx* must not be zero.
- beta* **COMPLEX** for **chbmv**
DOUBLE COMPLEX for **zhbmv**
- Specifies the scalar *beta*.
- y* **COMPLEX** for **chbmv**
DOUBLE COMPLEX for **zhbmv**
- Array, **DIMENSION** at least $(1 + (n - 1) * \text{abs}(incy))$. Before entry, the incremented array *y* must contain the vector *y*.
- incy* **INTEGER**. Specifies the increment for the elements of *y*. The value of *incy* must not be zero.

Output Parameters

- y* Overwritten by the updated vector *y*.

?hemv

Computes a matrix-vector product using a Hermitian matrix.

```
call chemv ( uplo, n, alpha, a, lda, x, incx, beta, y,
             incy )
call zhemv ( uplo, n, alpha, a, lda, x, incx, beta, y,
             incy )
```

Discussion

The ?hemv routines perform a matrix-vector operation defined as

$$y := \alpha * a * x + \beta * y,$$

where:

alpha and *beta* are scalars

x and *y* are *n*-element vectors

a is an *n* by *n* Hermitian matrix.

Input Parameters

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the array *a* is to be referenced, as follows:

<i>uplo</i> value	Part of Array <i>a</i> To Be Referenced
U or u	The upper triangular part of array <i>a</i> is to be referenced.
L or l	The lower triangular part of array <i>a</i> is to be referenced.

n INTEGER. Specifies the order of the matrix *a*. The value of *n* must be at least zero.

<i>alpha</i>	COMPLEX for chemv DOUBLE COMPLEX for zhemv Specifies the scalar <i>alpha</i> .
<i>a</i>	COMPLEX for chemv DOUBLE COMPLEX for zhemv Array, DIMENSION (<i>lda</i> , <i>n</i>). Before entry with <i>uplo</i> = 'U' or 'u', the leading <i>n</i> by <i>n</i> upper triangular part of the array <i>a</i> must contain the upper triangular part of the Hermitian matrix and the strictly lower triangular part of <i>a</i> is not referenced. Before entry with <i>uplo</i> = 'L' or 'l', the leading <i>n</i> by <i>n</i> lower triangular part of the array <i>a</i> must contain the lower triangular part of the Hermitian matrix and the strictly upper triangular part of <i>a</i> is not referenced. The imaginary parts of the diagonal elements need not be set and are assumed to be zero.
<i>lda</i>	INTEGER. Specifies the first dimension of <i>a</i> as declared in the calling (sub)program. The value of <i>lda</i> must be at least $\max(1, n)$.
<i>x</i>	COMPLEX for chemv DOUBLE COMPLEX for zhemv Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(\text{incx}))$. Before entry, the incremented array <i>x</i> must contain the <i>n</i> -element vector <i>x</i> .
<i>incx</i>	INTEGER. Specifies the increment for the elements of <i>x</i> . The value of <i>incx</i> must not be zero.
<i>beta</i>	COMPLEX for chemv DOUBLE COMPLEX for zhemv Specifies the scalar <i>beta</i> . When <i>beta</i> is supplied as zero then <i>y</i> need not be set on input.

y **COMPLEX** for `chemv`
 DOUBLE COMPLEX for `zhemv`

Array, **DIMENSION** at least $(1 + (n - 1) * \text{abs}(incy))$.
Before entry, the incremented array *y* must contain the
n-element vector *y*.

incy **INTEGER**. Specifies the increment for the elements of *y*.
The value of *incy* must not be zero.

Output Parameters

y Overwritten by the updated vector *y*.

?her

*Performs a rank-1 update of a
Hermitian matrix.*

```
call cher ( uplo, n, alpha, x, incx, a, lda )  
call zher ( uplo, n, alpha, x, incx, a, lda )
```

Discussion

The ?her routines perform a matrix-vector operation defined as

$a := \text{alpha} * x * \text{conjg}(x') + a,$

where:

alpha is a real scalar

x is an *n*-element vector

a is an *n* by *n* Hermitian matrix.

Input Parameters

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the array *a* is to be referenced, as follows:

<i>uplo</i> value	Part of Array <i>a</i> To Be Referenced
U or u	The upper triangular part of array <i>a</i> is to be referenced.
L or l	The lower triangular part of array <i>a</i> is to be referenced.

n INTEGER. Specifies the order of the matrix *a*. The value of *n* must be at least zero.

alpha REAL for *cher*
DOUBLE PRECISION for *zher*
Specifies the scalar *alpha*.

x COMPLEX for *cher*
DOUBLE COMPLEX for *zher*
Array, dimension at least $(1 + (n - 1) * \text{abs}(\text{incx}))$. Before entry, the incremented array *x* must contain the *n*-element vector *x*.

incx INTEGER. Specifies the increment for the elements of *x*. The value of *incx* must not be zero.

a COMPLEX for *cher*
DOUBLE COMPLEX for *zher*
Array, DIMENSION (*lda*, *n*). Before entry with *uplo* = 'U' or 'u', the leading *n* by *n* upper triangular part of the array *a* must contain the upper triangular part of the Hermitian matrix and the strictly lower triangular part of *a* is not referenced.
Before entry with *uplo* = 'L' or 'l', the leading *n* by *n* lower triangular part of the array *a* must contain the lower triangular part of the Hermitian matrix and the strictly upper triangular part of *a* is not referenced.

The imaginary parts of the diagonal elements need not be set and are assumed to be zero.

lda **INTEGER**. Specifies the first dimension of *a* as declared in the calling (sub)program. The value of *lda* must be at least $\max(1, n)$.

Output Parameters

a With *uplo* = 'U' or 'u', the upper triangular part of the array *a* is overwritten by the upper triangular part of the updated matrix.

With *uplo* = 'L' or 'l', the lower triangular part of the array *a* is overwritten by the lower triangular part of the updated matrix.

The imaginary parts of the diagonal elements are set to zero.

?her2

Performs a rank-2 update of a Hermitian matrix.

```
call cher2 ( uplo, n, alpha, x, incx, y, incy, a, lda )
call zher2 ( uplo, n, alpha, x, incx, y, incy, a, lda )
```

Discussion

The ?her2 routines perform a matrix-vector operation defined as

$$a := \alpha * x * \text{conjg}(y') + \text{conjg}(\alpha) * y * \text{conjg}(x') + a,$$

where:

alpha is a scalar

x and *y* are *n*-element vectors

a is an *n* by *n* Hermitian matrix.

Input Parameters

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the array *a* is to be referenced, as follows:

<i>uplo</i> value	Part of Array <i>a</i> To Be Referenced
U or u	The upper triangular part of array <i>a</i> is to be referenced.
L or l	The lower triangular part of array <i>a</i> is to be referenced.

n INTEGER. Specifies the order of the matrix *a*. The value of *n* must be at least zero.

alpha COMPLEX for *cher2*
DOUBLE COMPLEX for *zher2*
Specifies the scalar *alpha*.

x COMPLEX for *cher2*
DOUBLE COMPLEX for *zher2*
Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(\text{incx}))$. Before entry, the incremented array *x* must contain the *n*-element vector *x*.

incx INTEGER. Specifies the increment for the elements of *x*. The value of *incx* must not be zero.

y COMPLEX for *cher2*
DOUBLE COMPLEX for *zher2*
Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(\text{incy}))$. Before entry, the incremented array *y* must contain the *n*-element vector *y*.

incy INTEGER. Specifies the increment for the elements of *y*. The value of *incy* must not be zero.

a COMPLEX for *cher2*
DOUBLE COMPLEX for *zher2*

Array, `DIMENSION (lda, n)`. Before entry with `uplo = 'U' or 'u'`, the leading n by n upper triangular part of the array `a` must contain the upper triangular part of the Hermitian matrix and the strictly lower triangular part of `a` is not referenced.

Before entry with `uplo = 'L' or 'l'`, the leading n by n lower triangular part of the array `a` must contain the lower triangular part of the Hermitian matrix and the strictly upper triangular part of `a` is not referenced.

The imaginary parts of the diagonal elements need not be set and are assumed to be zero.

`lda` **INTEGER**. Specifies the first dimension of `a` as declared in the calling (sub)program. The value of `lda` must be at least `max(1, n)`.

Output Parameters

`a` With `uplo = 'U' or 'u'`, the upper triangular part of the array `a` is overwritten by the upper triangular part of the updated matrix.

With `uplo = 'L' or 'l'`, the lower triangular part of the array `a` is overwritten by the lower triangular part of the updated matrix.

The imaginary parts of the diagonal elements are set to zero.

?hpmv

Computes a matrix-vector product using a Hermitian packed matrix.

```
call chpmv ( uplo, n, alpha, ap, x, incx, beta, y, incy )
call zhpmv ( uplo, n, alpha, ap, x, incx, beta, y, incy )
```

Discussion

The `?hpmv` routines perform a matrix-vector operation defined as

$$y := \alpha * a * x + \beta * y,$$

where:

`alpha` and `beta` are scalars

`x` and `y` are n -element vectors

`a` is an n by n Hermitian matrix, supplied in packed form.

Input Parameters

`uplo` CHARACTER*1. Specifies whether the upper or lower triangular part of the matrix `a` is supplied in the packed array `ap`, as follows:

<code>uplo</code> value	Part of Matrix <code>a</code> Supplied
U or u	The upper triangular part of matrix <code>a</code> is supplied in <code>ap</code> .
L or l	The lower triangular part of matrix <code>a</code> is supplied in <code>ap</code> .

`n` INTEGER. Specifies the order of the matrix `a`. The value of `n` must be at least zero.

`alpha` COMPLEX for `chpmv`
DOUBLE COMPLEX for `zhpmv`
Specifies the scalar `alpha`.

`ap` COMPLEX for `chpmv`
DOUBLE COMPLEX for `zhpmv`
Array, DIMENSION at least $((n * (n + 1)) / 2)$. Before entry with `uplo = 'U'` or `'u'`, the array `ap` must contain the upper triangular part of the Hermitian matrix packed sequentially, column-by-column, so that `ap(1)` contains `a(1, 1)`, `ap(2)` and `ap(3)` contain `a(1, 2)` and `a(2, 2)` respectively, and so on. Before entry with `uplo = 'L'` or `'l'`, the array `ap` must contain the lower triangular part of the Hermitian matrix packed

sequentially, column-by-column, so that $ap(1)$ contains $a(1, 1)$, $ap(2)$ and $ap(3)$ contain $a(2, 1)$ and $a(3, 1)$ respectively, and so on.

The imaginary parts of the diagonal elements need not be set and are assumed to be zero.

x	<p>COMPLEX for <code>chpmv</code> DOUBLE PRECISION COMPLEX for <code>zhpmv</code></p> <p>Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(incx))$. Before entry, the incremented array x must contain the n-element vector x.</p>
$incx$	<p>INTEGER. Specifies the increment for the elements of x. The value of $incx$ must not be zero.</p>
$beta$	<p>COMPLEX for <code>chpmv</code> DOUBLE COMPLEX for <code>zhpmv</code></p> <p>Specifies the scalar $beta$. When $beta$ is supplied as zero then y need not be set on input.</p>
y	<p>COMPLEX for <code>chpmv</code> DOUBLE COMPLEX for <code>zhpmv</code></p> <p>Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(incy))$. Before entry, the incremented array y must contain the n-element vector y.</p>
$incy$	<p>INTEGER. Specifies the increment for the elements of y. The value of $incy$ must not be zero.</p>

Output Parameters

y	Overwritten by the updated vector y
-----	---------------------------------------

?hpr

Performs a rank-1 update of a Hermitian packed matrix.

```
call chpr ( uplo, n, alpha, x, incx, ap )
call zhpr ( uplo, n, alpha, x, incx, ap )
```

Discussion

The ?hpr routines perform a matrix-vector operation defined as

$$a := \alpha * x * \text{conjg}(x') + a,$$

where:

alpha is a real scalar

x is an *n*-element vector

a is an *n* by *n* Hermitian matrix, supplied in packed form.

Input Parameters

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the matrix *a* is supplied in the packed array *ap*, as follows:

<i>uplo</i> value	Part of Matrix <i>a</i> Supplied
U or u	The upper triangular part of matrix <i>a</i> is supplied in <i>ap</i> .
L or l	The lower triangular part of matrix <i>a</i> is supplied in <i>ap</i> .

n INTEGER. Specifies the order of the matrix *a*. The value of *n* must be at least zero.

alpha REAL for chpr
DOUBLE PRECISION for zhpr
Specifies the scalar *alpha*.

x **COMPLEX** for *chpr*
DOUBLE COMPLEX for *zhpr*

Array, **DIMENSION** at least $(1 + (n - 1) * \text{abs}(incx))$. Before entry, the incremented array *x* must contain the *n*-element vector *x*.

incx **INTEGER**. Specifies the increment for the elements of *x*. *incx* must not be zero.

ap **COMPLEX** for *chpr*
DOUBLE COMPLEX for *zhpr*

Array, **DIMENSION** at least $((n * (n + 1)) / 2)$. Before entry with *uplo* = 'U' or 'u', the array *ap* must contain the upper triangular part of the Hermitian matrix packed sequentially, column-by-column, so that *ap*(1) contains *a*(1, 1), *ap*(2) and *ap*(3) contain *a*(1, 2) and *a*(2, 2) respectively, and so on.

Before entry with *uplo* = 'L' or 'l', the array *ap* must contain the lower triangular part of the Hermitian matrix packed sequentially, column-by-column, so that *ap*(1) contains *a*(1, 1), *ap*(2) and *ap*(3) contain *a*(2, 1) and *a*(3, 1) respectively, and so on.

The imaginary parts of the diagonal elements need not be set and are assumed to be zero.

Output Parameters

ap With *uplo* = 'U' or 'u', overwritten by the upper triangular part of the updated matrix.

With *uplo* = 'L' or 'l', overwritten by the lower triangular part of the updated matrix.

The imaginary parts of the diagonal elements are set to zero.

?hpr2

Performs a rank-2 update of a Hermitian packed matrix.

```
call chpr2 ( uplo, n, alpha, x, incx, y, incy, ap )
call zhpr2 ( uplo, n, alpha, x, incx, y, incy, ap )
```

Discussion

The ?hpr2 routines perform a matrix-vector operation defined as

$$a := \alpha * x * \text{conjg}(y') + \text{conjg}(\alpha) * y * \text{conjg}(x') + a,$$

where:

alpha is a scalar

x and *y* are *n*-element vectors

a is an *n* by *n* Hermitian matrix, supplied in packed form.

Input Parameters

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the matrix *a* is supplied in the packed array *ap*, as follows

<i>uplo</i> value	Part of Matrix <i>a</i> Supplied
U or u	The upper triangular part of matrix <i>a</i> is supplied in <i>ap</i> .
L or l	The lower triangular part of matrix <i>a</i> is supplied in <i>ap</i> .

n INTEGER. Specifies the order of the matrix *a*. The value of *n* must be at least zero.

alpha COMPLEX for chpr2
DOUBLE COMPLEX for zhpr2
Specifies the scalar *alpha*.

<code>x</code>	<p>COMPLEX for <code>chpr2</code> DOUBLE COMPLEX for <code>zhpr2</code></p> <p>Array, dimension at least $(1 + (n - 1) * \text{abs}(\text{incx}))$. Before entry, the incremented array <code>x</code> must contain the n-element vector <code>x</code>.</p>
<code>incx</code>	<p>INTEGER. Specifies the increment for the elements of <code>x</code>. The value of <code>incx</code> must not be zero.</p>
<code>y</code>	<p>COMPLEX for <code>chpr2</code> DOUBLE COMPLEX for <code>zhpr2</code></p> <p>Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(\text{incy}))$. Before entry, the incremented array <code>y</code> must contain the n-element vector <code>y</code>.</p>
<code>incy</code>	<p>INTEGER. Specifies the increment for the elements of <code>y</code>. The value of <code>incy</code> must not be zero.</p>
<code>ap</code>	<p>COMPLEX for <code>chpr2</code> DOUBLE COMPLEX for <code>zhpr2</code></p> <p>Array, DIMENSION at least $((n * (n + 1)) / 2)$. Before entry with <code>uplo = 'U'</code> or <code>'u'</code>, the array <code>ap</code> must contain the upper triangular part of the Hermitian matrix packed sequentially, column-by-column, so that <code>ap(1)</code> contains <code>a(1, 1)</code>, <code>ap(2)</code> and <code>ap(3)</code> contain <code>a(1, 2)</code> and <code>a(2, 2)</code> respectively, and so on.</p> <p>Before entry with <code>uplo = 'L'</code> or <code>'l'</code>, the array <code>ap</code> must contain the lower triangular part of the Hermitian matrix packed sequentially, column-by-column, so that <code>ap(1)</code> contains <code>a(1, 1)</code>, <code>ap(2)</code> and <code>ap(3)</code> contain <code>a(2, 1)</code> and <code>a(3, 1)</code> respectively, and so on.</p> <p>The imaginary parts of the diagonal elements need not be set and are assumed to be zero.</p>

Output Parameters

ap With *uplo* = 'U' or 'u', overwritten by the upper triangular part of the updated matrix.

With *uplo* = 'L' or 'l', overwritten by the lower triangular part of the updated matrix.

The imaginary parts of the diagonal elements need are set to zero.

?sbmv

Computes a matrix-vector product using a symmetric band matrix.

```
call ssbmv ( uplo, n, k, alpha, a, lda, x, incx, beta, y,
             incy )
call dsbmv ( uplo, n, k, alpha, a, lda, x, incx, beta, y,
             incy )
```

Discussion

The *?sbmv* routines perform a matrix-vector operation defined as

$$y := \alpha * a * x + \beta * y,$$

where:

alpha and *beta* are scalars

x and *y* are *n*-element vectors

a is an *n* by *n* symmetric band matrix, with *k* super-diagonals.

Input Parameters

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the band matrix *a* is being supplied, as follows:

<i>uplo</i> value	Part of Matrix <i>a</i> Supplied
U or u	The upper triangular part of matrix <i>a</i> is supplied.
L or l	The lower triangular part of matrix <i>a</i> is supplied.

n INTEGER. Specifies the order of the matrix *a*. The value of *n* must be at least zero.

k INTEGER. Specifies the number of super-diagonals of the matrix *a*. The value of *k* must satisfy $0 \leq k$.

alpha REAL for *ssbmv*
DOUBLE PRECISION for *dsbmv*
Specifies the scalar *alpha*.

a REAL for *ssbmv*
DOUBLE PRECISION for *dsbmv*
Array, DIMENSION (*lda*, *n*). Before entry with *uplo* = 'U' or 'u', the leading (*k* + 1) by *n* part of the array *a* must contain the upper triangular band part of the symmetric matrix, supplied column-by-column, with the leading diagonal of the matrix in row (*k* + 1) of the array, the first super-diagonal starting at position 2 in row *k*, and so on. The top left *k* by *k* triangle of the array *a* is not referenced.

The following program segment transfers the upper triangular part of a symmetric band matrix from conventional full matrix storage to band storage:

```
do 20, j = 1, n
  m = k + 1 - j
  do 10, i = max( 1, j - k ), j
    a( m + i, j ) = matrix( i, j )
    10 continue
  20 continue
```

Before entry with *uplo* = 'L' or 'l', the leading $(k + 1)$ by n part of the array *a* must contain the lower triangular band part of the symmetric matrix, supplied column-by-column, with the leading diagonal of the matrix in row 1 of the array, the first sub-diagonal starting at position 1 in row 2, and so on. The bottom right k by k triangle of the array *a* is not referenced.

The following program segment transfers the lower triangular part of a symmetric band matrix from conventional full matrix storage to band storage:

```
do 20, j = 1, n
  m = 1 - j
  do 10, i = j, min( n, j + k )
    a( m + i, j ) = matrix( i, j )
  10 continue
20 continue
```

lda **INTEGER**. Specifies the first dimension of *a* as declared in the calling (sub)program. The value of *lda* must be at least $(k + 1)$.

x **REAL** for *ssbmv*
DOUBLE PRECISION for *dsbmv*
 Array, **DIMENSION** at least $(1 + (n - 1) * \text{abs}(incx))$.
 Before entry, the incremented array *x* must contain the vector *x*.

incx **INTEGER**. Specifies the increment for the elements of *x*.
 The value of *incx* must not be zero.

beta **REAL** for *ssbmv*
DOUBLE PRECISION for *dsbmv*
 Specifies the scalar *beta*.

y **REAL** for *ssbmv*
DOUBLE PRECISION for *dsbmv*
 Array, **DIMENSION** at least $(1 + (n - 1) * \text{abs}(incy))$.
 Before entry, the incremented array *y* must contain the vector *y*.

incy **INTEGER**. Specifies the increment for the elements of *y*. The value of *incy* must not be zero.

Output Parameters

y Overwritten by the updated vector *y*.

?spmv

Computes a matrix-vector product using a symmetric packed matrix.

```
call sspmv ( uplo, n, alpha, ap, x, incx, beta, y, incy )
call dspmv ( uplo, n, alpha, ap, x, incx, beta, y, incy )
```

Discussion

The ?spmv routines perform a matrix-vector operation defined as

$$y := \alpha * a * x + \beta * y,$$

where:

alpha and *beta* are scalars

x and *y* are *n*-element vectors

a is an *n* by *n* symmetric matrix, supplied in packed form.

Input Parameters

uplo **CHARACTER*1**. Specifies whether the upper or lower triangular part of the matrix *a* is supplied in the packed array *ap*, as follows:

<i>uplo</i> value	Part of Matrix <i>a</i> Supplied
U or u	The upper triangular part of matrix <i>a</i> is supplied in <i>ap</i> .
L or l	The lower triangular part of matrix <i>a</i> is supplied in <i>ap</i> .

<i>n</i>	INTEGER . Specifies the order of the matrix <i>a</i> . The value of <i>n</i> must be at least zero.
<i>alpha</i>	REAL for sspmv DOUBLE PRECISION for dspmv Specifies the scalar <i>alpha</i> .
<i>ap</i>	REAL for sspmv DOUBLE PRECISION for dspmv Array, DIMENSION at least $((n*(n+1))/2)$. Before entry with <i>uplo</i> = 'U' or 'u', the array <i>ap</i> must contain the upper triangular part of the symmetric matrix packed sequentially, column-by-column, so that <i>ap</i> (1) contains <i>a</i> (1, 1), <i>ap</i> (2) and <i>ap</i> (3) contain <i>a</i> (1, 2) and <i>a</i> (2, 2) respectively, and so on. Before entry with <i>uplo</i> = 'L' or 'l', the array <i>ap</i> must contain the lower triangular part of the symmetric matrix packed sequentially, column-by-column, so that <i>ap</i> (1) contains <i>a</i> (1, 1), <i>ap</i> (2) and <i>ap</i> (3) contain <i>a</i> (2, 1) and <i>a</i> (3, 1) respectively, and so on.
<i>x</i>	REAL for sspmv DOUBLE PRECISION for dspmv Array, DIMENSION at least $(1 + (n - 1)*abs(incx))$. Before entry, the incremented array <i>x</i> must contain the <i>n</i> -element vector <i>x</i> .
<i>incx</i>	INTEGER . Specifies the increment for the elements of <i>x</i> . The value of <i>incx</i> must not be zero.
<i>beta</i>	REAL for sspmv DOUBLE PRECISION for dspmv Specifies the scalar <i>beta</i> . When <i>beta</i> is supplied as zero, then <i>y</i> need not be set on input.
<i>y</i>	REAL for sspmv DOUBLE PRECISION for dspmv Array, DIMENSION at least $(1 + (n - 1)*abs(incy))$. Before entry, the incremented array <i>y</i> must contain the <i>n</i> -element vector <i>y</i> .

incy **INTEGER**. Specifies the increment for the elements of *y*. The value of *incy* must not be zero.

Output Parameters

y Overwritten by the updated vector *y*.

?spr

Performs a rank-1 update of a symmetric packed matrix.

```
call sspr( uplo, n, alpha, x, incx, ap )
call dspr( uplo, n, alpha, x, incx, ap )
```

Discussion

The ?spr routines perform a matrix-vector operation defined as

$$a := \alpha * x * x' + a,$$

where:

alpha is a real scalar

x is an *n*-element vector

a is an *n* by *n* symmetric matrix, supplied in packed form.

Input Parameters

uplo **CHARACTER*1**. Specifies whether the upper or lower triangular part of the matrix *a* is supplied in the packed array *ap*, as follows:

<i>uplo</i> value	Part of Matrix <i>a</i> Supplied
U or u	The upper triangular part of matrix <i>a</i> is supplied in <i>ap</i> .
L or l	The lower triangular part of matrix <i>a</i> is supplied in <i>ap</i> .

<i>n</i>	INTEGER . Specifies the order of the matrix <i>a</i> . The value of <i>n</i> must be at least zero.
<i>alpha</i>	REAL for sspr DOUBLE PRECISION for dspr Specifies the scalar <i>alpha</i> .
<i>x</i>	REAL for sspr DOUBLE PRECISION for dspr Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(\text{incx}))$. Before entry, the incremented array <i>x</i> must contain the <i>n</i> -element vector <i>x</i> .
<i>incx</i>	INTEGER . Specifies the increment for the elements of <i>x</i> . The value of <i>incx</i> must not be zero.
<i>ap</i>	REAL for sspr DOUBLE PRECISION for dspr Array, DIMENSION at least $((n * (n + 1)) / 2)$. Before entry with <i>uplo</i> = 'U' or 'u', the array <i>ap</i> must contain the upper triangular part of the symmetric matrix packed sequentially, column-by-column, so that <i>ap</i> (1) contains <i>a</i> (1,1), <i>ap</i> (2) and <i>ap</i> (3) contain <i>a</i> (1,2) and <i>a</i> (2,2) respectively, and so on. Before entry with <i>uplo</i> = 'L' or 'l', the array <i>ap</i> must contain the lower triangular part of the symmetric matrix packed sequentially, column-by-column, so that <i>ap</i> (1) contains <i>a</i> (1,1), <i>ap</i> (2) and <i>ap</i> (3) contain <i>a</i> (2,1) and <i>a</i> (3,1) respectively, and so on.

Output Parameters

<i>ap</i>	With <i>uplo</i> = 'U' or 'u', overwritten by the upper triangular part of the updated matrix. With <i>uplo</i> = 'L' or 'l', overwritten by the lower triangular part of the updated matrix.
-----------	--

?spr2

Performs a rank-2 update of a symmetric packed matrix.

```
call sspr2( uplo, n, alpha, x, incx, y, incy, ap )
call dspr2( uplo, n, alpha, x, incx, y, incy, ap )
```

Discussion

The ?spr2 routines perform a matrix-vector operation defined as

$$a := \alpha * x * y' + \alpha * y * x' + a,$$

where:

alpha is a scalar

x and *y* are *n*-element vectors

a is an *n* by *n* symmetric matrix, supplied in packed form.

Input Parameters

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the matrix *a* is supplied in the packed array *ap*, as follows:

<i>uplo</i> value	Part of Matrix <i>a</i> Supplied
U or u	The upper triangular part of matrix <i>a</i> is supplied in <i>ap</i> .
L or l	The lower triangular part of matrix <i>a</i> is supplied in <i>ap</i> .

n INTEGER. Specifies the order of the matrix *a*. The value of *n* must be at least zero.

alpha REAL for *sspr2*
DOUBLE PRECISION for *dspr2*
Specifies the scalar *alpha*.

x REAL for `sspr2`
DOUBLE PRECISION for `dspr2`
Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(incx))$.
Before entry, the incremented array *x* must contain the *n*-element vector *x*.

incx INTEGER. Specifies the increment for the elements of *x*.
The value of *incx* must not be zero.

y REAL for `sspr2`
DOUBLE PRECISION for `dspr2`
Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(incy))$.
Before entry, the incremented array *y* must contain the *n*-element vector *y*.

incy INTEGER. Specifies the increment for the elements of *y*.
The value of *incy* must not be zero.

ap REAL for `sspr2`
DOUBLE PRECISION for `dspr2`
Array, DIMENSION at least $((n * (n + 1)) / 2)$. Before entry with *uplo* = 'U' or 'u', the array *ap* must contain the upper triangular part of the symmetric matrix packed sequentially, column-by-column, so that *ap*(1) contains *a*(1,1), *ap*(2) and *ap*(3) contain *a*(1,2) and *a*(2,2) respectively, and so on.
Before entry with *uplo* = 'L' or 'l', the array *ap* must contain the lower triangular part of the symmetric matrix packed sequentially, column-by-column, so that *ap*(1) contains *a*(1,1), *ap*(2) and *ap*(3) contain *a*(2,1) and *a*(3,1) respectively, and so on.

Output Parameters

ap With *uplo* = 'U' or 'u', overwritten by the upper triangular part of the updated matrix.
With *uplo* = 'L' or 'l', overwritten by the lower triangular part of the updated matrix.

?symv

Computes a matrix-vector product for a symmetric matrix.

```
call ssymv ( uplo, n, alpha, a, lda, x, incx, beta, y,
             incy )
call dsymv ( uplo, n, alpha, a, lda, x, incx, beta, y,
             incy )
```

Discussion

The ?symv routines perform a matrix-vector operation defined as

$$y := \alpha * a * x + \beta * y,$$

where:

alpha and *beta* are scalars

x and *y* are *n*-element vectors

a is an *n* by *n* symmetric matrix.

Input Parameters

uplo **CHARACTER*1**. Specifies whether the upper or lower triangular part of the array *a* is to be referenced, as follows:

<i>uplo</i> value	Part of Array <i>a</i> To Be Referenced
U or u	The upper triangular part of array <i>a</i> is to be referenced.
L or l	The lower triangular part of array <i>a</i> is to be referenced.

n **INTEGER**. Specifies the order of the matrix *a*. The value of *n* must be at least zero.

<i>alpha</i>	REAL for <i>ssymv</i> DOUBLE PRECISION for <i>dsymv</i> Specifies the scalar <i>alpha</i> .
<i>a</i>	REAL for <i>ssymv</i> DOUBLE PRECISION for <i>dsymv</i> Array, DIMENSION (<i>lda</i> , <i>n</i>). Before entry with <i>uplo</i> = 'U' or 'u', the leading <i>n</i> by <i>n</i> upper triangular part of the array <i>a</i> must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of <i>a</i> is not referenced. Before entry with <i>uplo</i> = 'L' or 'l', the leading <i>n</i> by <i>n</i> lower triangular part of the array <i>a</i> must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of <i>a</i> is not referenced.
<i>lda</i>	INTEGER. Specifies the first dimension of <i>a</i> as declared in the calling (sub)program. The value of <i>lda</i> must be at least $\max(1, n)$.
<i>x</i>	REAL for <i>ssymv</i> DOUBLE PRECISION for <i>dsymv</i> Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(\text{incx}))$. Before entry, the incremented array <i>x</i> must contain the <i>n</i> -element vector <i>x</i> .
<i>incx</i>	INTEGER. Specifies the increment for the elements of <i>x</i> . The value of <i>incx</i> must not be zero.
<i>beta</i>	REAL for <i>ssymv</i> DOUBLE PRECISION for <i>dsymv</i> Specifies the scalar <i>beta</i> . When <i>beta</i> is supplied as zero, then <i>y</i> need not be set on input.
<i>y</i>	REAL for <i>ssymv</i> DOUBLE PRECISION for <i>dsymv</i> Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(\text{incy}))$. Before entry, the incremented array <i>y</i> must contain the <i>n</i> -element vector <i>y</i> .

incy **INTEGER**. Specifies the increment for the elements of *y*. The value of *incy* must not be zero.

Output Parameters

y Overwritten by the updated vector *y*.

?syr

Performs a rank-1 update of a symmetric matrix.

```
call ssyr( uplo, n, alpha, x, incx, a, lda )
call dsyr( uplo, n, alpha, x, incx, a, lda )
```

Discussion

The *?syr* routines perform a matrix-vector operation defined as

$$a := \alpha * x * x' + a,$$

where:

alpha is a real scalar

x is an *n*-element vector

a is an *n* by *n* symmetric matrix.

Input Parameters

uplo **CHARACTER*1**. Specifies whether the upper or lower triangular part of the array *a* is to be referenced, as follows:

<i>uplo</i> value	Part of Array <i>a</i> To Be Referenced
U or u	The upper triangular part of array <i>a</i> is to be referenced.
L or l	The lower triangular part of array <i>a</i> is to be referenced.

<i>n</i>	INTEGER . Specifies the order of the matrix <i>a</i> . The value of <i>n</i> must be at least zero.
<i>alpha</i>	REAL for <i>ssyr</i> DOUBLE PRECISION for <i>dsyr</i> Specifies the scalar <i>alpha</i> .
<i>x</i>	REAL for <i>ssyr</i> DOUBLE PRECISION for <i>dsyr</i> Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(incx))$. Before entry, the incremented array <i>x</i> must contain the <i>n</i> -element vector <i>x</i> .
<i>incx</i>	INTEGER . Specifies the increment for the elements of <i>x</i> . The value of <i>incx</i> must not be zero.
<i>a</i>	REAL for <i>ssyr</i> DOUBLE PRECISION for <i>dsyr</i> Array, DIMENSION (<i>lda</i> , <i>n</i>). Before entry with <i>uplo</i> = 'U' or 'u', the leading <i>n</i> by <i>n</i> upper triangular part of the array <i>a</i> must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of <i>a</i> is not referenced. Before entry with <i>uplo</i> = 'L' or 'l', the leading <i>n</i> by <i>n</i> lower triangular part of the array <i>a</i> must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of <i>a</i> is not referenced.
<i>lda</i>	INTEGER . Specifies the first dimension of <i>a</i> as declared in the calling (sub)program. The value of <i>lda</i> must be at least $\max(1, n)$.

Output Parameters

<i>a</i>	With <i>uplo</i> = 'U' or 'u', the upper triangular part of the array <i>a</i> is overwritten by the upper triangular part of the updated matrix. With <i>uplo</i> = 'L' or 'l', the lower triangular part of the array <i>a</i> is overwritten by the lower triangular part of the updated matrix.
----------	--

?syr2

Performs a rank-2 update of symmetric matrix.

```
call ssyr2( uplo, n, alpha, x, incx, y, incy, a, lda )
call dsyr2( uplo, n, alpha, x, incx, y, incy, a, lda )
```

Discussion

The ?syr2 routines perform a matrix-vector operation defined as

$$a := \alpha * x * y' + \alpha * y * x' + a,$$

where:

alpha is a scalar

x and *y* are *n*-element vectors

a is an *n* by *n* symmetric matrix.

Input Parameters

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the array *a* is to be referenced, as follows:

<i>uplo</i> value	Part of Array <i>a</i> To Be Referenced
U or u	The upper triangular part of array <i>a</i> is to be referenced.
L or l	The lower triangular part of array <i>a</i> is to be referenced.

n INTEGER. Specifies the order of the matrix *a*. The value of *n* must be at least zero.

alpha REAL for ssyr2
DOUBLE PRECISION for dsyr2
Specifies the scalar *alpha*.

<i>x</i>	REAL for <i>ssyr2</i> DOUBLE PRECISION for <i>dsyr2</i> Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(incx))$. Before entry, the incremented array <i>x</i> must contain the <i>n</i> -element vector <i>x</i> .
<i>incx</i>	INTEGER. Specifies the increment for the elements of <i>x</i> . The value of <i>incx</i> must not be zero.
<i>y</i>	REAL for <i>ssyr2</i> DOUBLE PRECISION for <i>dsyr2</i> Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(incy))$. Before entry, the incremented array <i>y</i> must contain the <i>n</i> -element vector <i>y</i> .
<i>incy</i>	INTEGER. Specifies the increment for the elements of <i>y</i> . The value of <i>incy</i> must not be zero.
<i>a</i>	REAL for <i>ssyr2</i> DOUBLE PRECISION for <i>dsyr2</i> Array, DIMENSION (<i>lda</i> , <i>n</i>). Before entry with <i>uplo</i> = 'U' or 'u', the leading <i>n</i> by <i>n</i> upper triangular part of the array <i>a</i> must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of <i>a</i> is not referenced. Before entry with <i>uplo</i> = 'L' or 'l', the leading <i>n</i> by <i>n</i> lower triangular part of the array <i>a</i> must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of <i>a</i> is not referenced.
<i>lda</i>	INTEGER. Specifies the first dimension of <i>a</i> as declared in the calling (sub)program. The value of <i>lda</i> must be at least $\max(1, n)$.

Output Parameters

a With *uplo* = 'U' or 'u', the upper triangular part of the array *a* is overwritten by the upper triangular part of the updated matrix.

With *uplo* = 'L' or 'l', the lower triangular part of the array *a* is overwritten by the lower triangular part of the updated matrix.

?tbmv

Computes a matrix-vector product using a triangular band matrix.

```
call stbmv ( uplo, trans, diag, n, k, a, lda, x, incx )
call dtbmv ( uplo, trans, diag, n, k, a, lda, x, incx )
call ctbmv ( uplo, trans, diag, n, k, a, lda, x, incx )
call ztbmv ( uplo, trans, diag, n, k, a, lda, x, incx )
```

Discussion

The ?tbmv routines perform one of the matrix-vector operations defined as

$x := a*x$, or $x := a'*x$, or $x := \text{conjg}(a')*x$,

where:

x is an *n*-element vector

a is an *n* by *n* unit, or non-unit, upper or lower triangular band matrix, with (*k* + 1) diagonals.

Input Parameters

uplo CHARACTER*1. Specifies whether the matrix is an upper or lower triangular matrix, as follows:

<i>uplo</i> value	Matrix <i>a</i>
U or u	An upper triangular matrix.
L or l	A lower triangular matrix.

trans CHARACTER*1. Specifies the operation to be performed, as follows:

<i>trans</i> value	Operation to be Performed
N or n	$x := a*x$
T or t	$x := a'*x$
C or c	$x := \text{conjg}(a')*x$

diag CHARACTER*1. Specifies whether or not *a* is unit triangular, as follows:

<i>diag</i> value	Matrix <i>a</i>
U or u	Matrix <i>a</i> is assumed to be unit triangular.
N or n	Matrix <i>a</i> is not assumed to be unit triangular.

n INTEGER. Specifies the order of the matrix *a*. The value of *n* must be at least zero.

k INTEGER. On entry with *uplo* = 'U' or 'u', *k* specifies the number of super-diagonals of the matrix *a*. On entry with *uplo* = 'L' or 'l', *k* specifies the number of sub-diagonals of the matrix *a*. The value of *k* must satisfy $0 \leq k$.

a REAL for *stbmv*
 DOUBLE PRECISION for *dtbmv*
 COMPLEX for *ctbmv*
 DOUBLE COMPLEX for *ztbmv*

Array, DIMENSION (*lda*, *n*). Before entry with *uplo* = 'U' or 'u', the leading (*k* + 1) by *n* part of the array *a* must contain the upper triangular band part of the matrix of coefficients, supplied column-by-column, with the leading diagonal of the matrix in row (*k* + 1)

of the array, the first super-diagonal starting at position 2 in row k , and so on. The top left k by k triangle of the array a is not referenced. The following program segment transfers an upper triangular band matrix from conventional full matrix storage to band storage:

```
do 20, j = 1, n
  m = k + 1 - j
  do 10, i = max(1, j - k), j
    a(m + i, j) = matrix(i, j)
  10 continue
20 continue
```

Before entry with $uplo = 'L'$ or $'l'$, the leading $(k + 1)$ by n part of the array a must contain the lower triangular band part of the matrix of coefficients, supplied column-by-column, with the leading diagonal of the matrix in row 1 of the array, the first sub-diagonal starting at position 1 in row 2, and so on. The bottom right k by k triangle of the array a is not referenced. The following program segment transfers a lower triangular band matrix from conventional full matrix storage to band storage:

```
do 20, j = 1, n
  m = 1 - j
  do 10, i = j, min(n, j + k)
    a(m + i, j) = matrix(i, j)
  10 continue
20 continue
```

Note that when $diag = 'U'$ or $'u'$, the elements of the array a corresponding to the diagonal elements of the matrix are not referenced, but are assumed to be unity.

lda

INTEGER. Specifies the first dimension of a as declared in the calling (sub)program. The value of lda must be at least $(k + 1)$.

x REAL for `stbmv`
 DOUBLE PRECISION for `dtbmv`
 COMPLEX for `ctbmv`
 DOUBLE COMPLEX for `ztbmv`

Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(\text{incx}))$.
 Before entry, the incremented array **x** must contain the *n*-element vector **x**.

incx INTEGER. Specifies the increment for the elements of **x**.
 The value of **incx** must not be zero.

Output Parameters

x Overwritten with the transformed vector **x**.

?tbsv

*Solves a system of linear equations
 whose coefficients are in a triangular
 band matrix.*

```
call stbsv ( uplo, trans, diag, n, k, a, lda, x, incx )
call dtbsv ( uplo, trans, diag, n, k, a, lda, x, incx )
call ctbsv ( uplo, trans, diag, n, k, a, lda, x, incx )
call ztbsv ( uplo, trans, diag, n, k, a, lda, x, incx )
```

Discussion

The `?tbsv` routines solve one of the following systems of equations:

$a*x = b$, or $a'*x = b$, or $\text{conjg}(a')*x = b$,

where:

b and **x** are *n*-element vectors

a is an *n* by *n* unit, or non-unit, upper or lower triangular band matrix, with $(k + 1)$ diagonals.

The routine does not test for singularity or near-singularity. Such tests must be performed before calling this routine.

Input Parameters

uplo CHARACTER*1. Specifies whether the matrix is an upper or lower triangular matrix, as follows:

<i>uplo</i> value	Matrix <i>a</i>
U or u	An upper triangular matrix.
L or l	A lower triangular matrix.

trans CHARACTER*1. Specifies the operation to be performed, as follows:

<i>trans</i> value	Operation to be Performed
N or n	$a*x = b$
T or t	$a'*x = b$
C or c	$conjg(a')*x = b$

diag CHARACTER*1. Specifies whether or not *a* is unit triangular, as follows:

<i>diag</i> value	Matrix <i>a</i>
U or u	Matrix <i>a</i> is assumed to be unit triangular.
N or n	Matrix <i>a</i> is not assumed to be unit triangular.

n INTEGER. Specifies the order of the matrix *a*. The value of *n* must be at least zero.

k INTEGER. On entry with *uplo* = 'U' or 'u', *k* specifies the number of super-diagonals of the matrix *a*. On entry with *uplo* = 'L' or 'l', *k* specifies the number of sub-diagonals of the matrix *a*. The value of *k* must satisfy $0 \leq k$.

a

REAL for `stbsv`
 DOUBLE PRECISION for `dtbsv`
 COMPLEX for `ctbsv`
 DOUBLE COMPLEX for `ztbsv`

Array, DIMENSION (*lda*, *n*). Before entry with *uplo* = 'U' or 'u', the leading (*k* + 1) by *n* part of the array *a* must contain the upper triangular band part of the matrix of coefficients, supplied column-by-column, with the leading diagonal of the matrix in row (*k* + 1) of the array, the first super-diagonal starting at position 2 in row *k*, and so on. The top left *k* by *k* triangle of the array *a* is not referenced.

The following program segment transfers an upper triangular band matrix from conventional full matrix storage to band storage:

```
do 20, j = 1, n
  m = k + 1 - j
  do 10, i = max(1, j - k), j
    a(m + i, j) = matrix (i, j)
  10 continue
20 continue
```

Before entry with *uplo* = 'L' or 'l', the leading (*k* + 1) by *n* part of the array *a* must contain the lower triangular band part of the matrix of coefficients, supplied column-by-column, with the leading diagonal of the matrix in row 1 of the array, the first sub-diagonal starting at position 1 in row 2, and so on. The bottom right *k* by *k* triangle of the array *a* is not referenced.

The following program segment transfers a lower triangular band matrix from conventional full matrix storage to band storage:

```
do 20, j = 1, n
  m = 1 - j
  do 10, i = j, min(n, j + k)
    a(m + i, j) = matrix (i, j)
  10 continue
20 continue
```

When *diag* = 'U' or 'u', the elements of the array *a* corresponding to the diagonal elements of the matrix are not referenced, but are assumed to be unity.

lda **INTEGER**. Specifies the first dimension of *a* as declared in the calling (sub)program. The value of *lda* must be at least $(k + 1)$.

x **REAL** for *stbsv*
DOUBLE PRECISION for *dtbsv*
COMPLEX for *ctbsv*
DOUBLE COMPLEX for *ztbsv*

Array, **DIMENSION** at least $(1 + (n - 1) * \text{abs}(\text{incx}))$. Before entry, the incremented array *x* must contain the *n*-element right-hand side vector *b*.

incx **INTEGER**. Specifies the increment for the elements of *x*. The value of *incx* must not be zero.

Output Parameters

x Overwritten with the solution vector *x*.

?tpmv

Computes a matrix-vector product using a triangular packed matrix.

```
call stpmv ( uplo, trans, diag, n, ap, x, incx )
call dtpmv ( uplo, trans, diag, n, ap, x, incx )
call ctpmv ( uplo, trans, diag, n, ap, x, incx )
call ztpmv ( uplo, trans, diag, n, ap, x, incx )
```

Discussion

The ?tpmv routines perform one of the matrix-vector operations defined as $x := a*x$, or $x := a'*x$, or $x := \text{conjg}(a')*x$,

where:

x is an n -element vector

a is an n by n unit, or non-unit, upper or lower triangular matrix, supplied in packed form.

Input Parameters

uplo CHARACTER*1. Specifies whether the matrix a is an upper or lower triangular matrix, as follows:

<i>uplo</i> value	Matrix a
U or u	An upper triangular matrix.
L or l	A lower triangular matrix.

trans CHARACTER*1. Specifies the operation to be performed, as follows:

<i>trans</i> value	Operation To Be Performed
N or n	$x := a*x$
T or t	$x := a'*x$
C or c	$x := \text{conjg}(a')*x$

diag CHARACTER*1. Specifies whether or not a is unit triangular, as follows:

<i>diag</i> value	Matrix a
U or u	Matrix a is assumed to be unit triangular.
N or n	Matrix a is not assumed to be unit triangular.

n INTEGER. Specifies the order of the matrix a . The value of n must be at least zero.

ap REAL for *stpmv*
 DOUBLE PRECISION for *dtpmv*
 COMPLEX for *ctpmv*
 DOUBLE COMPLEX for *ztpmv*

Array, `DIMENSION` at least $((n*(n+1))/2)$. Before entry with `uplo = 'U'` or `'u'`, the array `ap` must contain the upper triangular matrix packed sequentially, column-by-column, so that `ap(1)` contains `a(1,1)`, `ap(2)` and `ap(3)` contain `a(1,2)` and `a(2,2)` respectively, and so on. Before entry with `uplo = 'L'` or `'l'`, the array `ap` must contain the lower triangular matrix packed sequentially, column-by-column, so that `ap(1)` contains `a(1,1)`, `ap(2)` and `ap(3)` contain `a(2,1)` and `a(3,1)` respectively, and so on. When `diag = 'U'` or `'u'`, the diagonal elements of `a` are not referenced, but are assumed to be unity.

`x` `REAL` for `stpmv`
 `DOUBLE PRECISION` for `dtpmv`
 `COMPLEX` for `ctpmv`
 `DOUBLE COMPLEX` for `ztpmv`

Array, `DIMENSION` at least $(1 + (n - 1)*abs(incx))$. Before entry, the incremented array `x` must contain the `n`-element vector `x`.

`incx` `INTEGER`. Specifies the increment for the elements of `x`. The value of `incx` must not be zero.

Output Parameters

`x` Overwritten with the transformed vector `x`.

?tpsv

Solves a system of linear equations whose coefficients are in a triangular packed matrix.

```
call stpsv ( uplo, trans, diag, n, ap, x, incx )
call dtpsv ( uplo, trans, diag, n, ap, x, incx )
call ctpsv ( uplo, trans, diag, n, ap, x, incx )
call ztpsv ( uplo, trans, diag, n, ap, x, incx )
```

Discussion

The ?tpsv routines solve one of the following systems of equations

$a*x = b$, or $a'*x = b$, or $\text{conjg}(a')*x = b$,

where:

b and x are n -element vectors

a is an n by n unit, or non-unit, upper or lower triangular matrix, supplied in packed form.

This routine does not test for singularity or near-singularity. Such tests must be performed before calling this routine.

Input Parameters

uplo CHARACTER*1. Specifies whether the matrix a is an upper or lower triangular matrix, as follows:

<i>uplo</i> value	Matrix a
U or u	An upper triangular matrix.
L or l	A lower triangular matrix.

trans CHARACTER*1. Specifies the operation to be performed, as follows:

<i>trans</i> value	Operation To Be Performed
N or n	$a*x = b$
T or t	$a'*x = b$
C or c	$\text{conjg}(a')*x = b$

diag CHARACTER*1. Specifies whether or not *a* is unit triangular, as follows:

<i>diag</i> value	Matrix <i>a</i>
U or u	Matrix <i>a</i> is assumed to be unit triangular.
N or n	Matrix <i>a</i> is not assumed to be unit triangular.

n INTEGER. Specifies the order of the matrix *a*. The value of *n* must be at least zero.

ap REAL for *stpsv*
 DOUBLE PRECISION for *dtpsv*
 COMPLEX for *ctpsv*
 DOUBLE COMPLEX for *ztpsv*

Array, DIMENSION at least $((n*(n+1))/2)$. Before entry with *uplo* = 'U' or 'u', the array *ap* must contain the upper triangular matrix packed sequentially, column-by-column, so that *ap*(1) contains *a*(1, 1), *ap*(2) and *ap*(3) contain *a*(1, 2) and *a*(2, 2) respectively, and so on. Before entry with *uplo* = 'L' or 'l', the array *ap* must contain the lower triangular matrix packed sequentially, column-by-column, so that *ap*(1) contains *a*(1, 1), *ap*(2) and *ap*(3) contain *a*(2, 1) and *a*(3, 1) respectively, and so on. When *diag* = 'U' or 'u', the diagonal elements of *a* are not referenced, but are assumed to be unity.

x REAL for `stpsv`
DOUBLE PRECISION for `dtpsv`
COMPLEX for `ctpsv`
DOUBLE COMPLEX for `ztpsv`

Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(\text{incx}))$.
Before entry, the incremented array **x** must contain the *n*-element right-hand side vector *b*.

incx INTEGER. Specifies the increment for the elements of **x**.
The value of **incx** must not be zero.

Output Parameters

x Overwritten with the solution vector **x**.

?trmv

*Computes a matrix-vector product using
a triangular matrix.*

```
call strmv ( uplo, trans, diag, n, a, lda, x, incx )
call dtrmv ( uplo, trans, diag, n, a, lda, x, incx )
call ctrmv ( uplo, trans, diag, n, a, lda, x, incx )
call ztrmv ( uplo, trans, diag, n, a, lda, x, incx )
```

Discussion

The `?trmv` routines perform one of the following matrix-vector operations defined as

`x := a*x` or `x := a'*x` or `x := conjg(a')*x`,

where:

x is an *n*-element vector

a is an *n* by *n* unit, or non-unit, upper or lower triangular matrix.

Input Parameters

uplo CHARACTER*1. Specifies whether the matrix *a* is an upper or lower triangular matrix, as follows:

<i>uplo</i> value	Matrix <i>a</i>
U or u	An upper triangular matrix.
L or l	A lower triangular matrix.

trans CHARACTER*1. Specifies the operation to be performed, as follows:

<i>trans</i> value	Operation To Be Performed
N or n	$x := a*x$
T or t	$x := a'*x$
C or c	$x := \text{conjg}(a')*x$

diag CHARACTER*1. Specifies whether or not *a* is unit triangular, as follows:

<i>diag</i> value	Matrix <i>a</i>
U or u	Matrix <i>a</i> is assumed to be unit triangular.
N or n	Matrix <i>a</i> is not assumed to be unit triangular.

n INTEGER. Specifies the order of the matrix *a*. The value of *n* must be at least zero.

a REAL for *strmv*
 DOUBLE PRECISION for *dtrmv*
 COMPLEX for *ctrmv*
 DOUBLE COMPLEX for *ztrmv*

Array, DIMENSION (*lda*, *n*). Before entry with *uplo* = 'U' or 'u', the leading *n* by *n* upper triangular part of the array *a* must contain the upper triangular matrix and the strictly lower triangular part of *a* is not referenced. Before entry with *uplo* = 'L' or 'l', the leading *n* by *n* lower triangular part of the array *a* must

contain the lower triangular matrix and the strictly upper triangular part of a is not referenced. When $diag = 'U'$ or $'u'$, the diagonal elements of a are not referenced either, but are assumed to be unity.

lda **INTEGER**. Specifies the first dimension of a as declared in the calling (sub)program. The value of lda must be at least $\max(1, n)$.

x **REAL** for `strmv`
DOUBLE PRECISION for `dtrmv`
COMPLEX for `ctrmv`
DOUBLE COMPLEX for `ztrmv`

Array, **DIMENSION** at least $(1 + (n - 1) * \text{abs}(\text{incx}))$. Before entry, the incremented array x must contain the n -element vector x .

$incx$ **INTEGER**. Specifies the increment for the elements of x . The value of $incx$ must not be zero.

Output Parameters

x Overwritten with the transformed vector x .

?trsv

Solves a system of linear equations whose coefficients are in a triangular matrix.

```
call strsv ( uplo, trans, diag, n, a, lda, x, incx )
call dtrsv ( uplo, trans, diag, n, a, lda, x, incx )
call ctrsv ( uplo, trans, diag, n, a, lda, x, incx )
call ztrsv ( uplo, trans, diag, n, a, lda, x, incx )
```

Discussion

The `?trsv` routines solve one of the systems of equations:

$$a*x = b \text{ or } a'*x = b, \text{ or } \text{conjg}(a')*x = b,$$

where:

b and x are n -element vectors

a is an n by n unit, or non-unit, upper or lower triangular matrix.

The routine does not test for singularity or near-singularity. Such tests must be performed before calling this routine.

Input Parameters

uplo CHARACTER*1. Specifies whether the matrix is an upper or lower triangular matrix, as follows:

<i>uplo</i> value	Matrix a
U or u	An upper triangular matrix.
L or l	A lower triangular matrix.

trans CHARACTER*1. Specifies the operation to be performed, as follows:

<i>trans</i> value	Operation To Be Performed
N or n	$a*x = b$
T or t	$a'*x = b$
C or c	$\text{conjg}(a')*x = b$

diag CHARACTER*1. Specifies whether or not a is unit triangular, as follows:

<i>diag</i> value	Matrix a
U or u	Matrix a is assumed to be unit triangular.
N or n	Matrix a is not assumed to be unit triangular.

n INTEGER. Specifies the order of the matrix a . The value of n must be at least zero.

a	REAL for strsv DOUBLE PRECISION for dtrsv COMPLEX for ctrsv DOUBLE COMPLEX for ztrsv Array, DIMENSION (<i>lda</i> , <i>n</i>). Before entry with <i>uplo</i> = 'U' or 'u', the leading <i>n</i> by <i>n</i> upper triangular part of the array a must contain the upper triangular matrix and the strictly lower triangular part of a is not referenced. Before entry with <i>uplo</i> = 'L' or 'l', the leading <i>n</i> by <i>n</i> lower triangular part of the array a must contain the lower triangular matrix and the strictly upper triangular part of a is not referenced. When <i>diag</i> = 'U' or 'u', the diagonal elements of a are not referenced either, but are assumed to be unity.
lda	INTEGER . Specifies the first dimension of a as declared in the calling (sub)program. The value of <i>lda</i> must be at least $\max(1, n)$.
x	REAL for strsv DOUBLE PRECISION for dtrsv COMPLEX for ctrsv DOUBLE COMPLEX for ztrsv Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(\text{incx}))$. Before entry, the incremented array x must contain the <i>n</i> -element right-hand side vector b .
incx	INTEGER . Specifies the increment for the elements of x . The value of <i>incx</i> must not be zero.

Output Parameters

x	Overwritten with the solution vector x .
----------	---

BLAS Level 3 Routines

BLAS Level 3 routines perform matrix-matrix operations. Table 2-3 lists the BLAS Level 3 routine groups and the data types associated with them.

Table 2-3 BLAS Level 3 Routine Groups and Their Data Types

Routine Group	Data Types	Description
?gemm	s, d, c, z	Matrix-matrix product of general matrices
?hemm	c, z	Matrix-matrix product of Hermitian matrices
?herk	c, z	Rank-k update of Hermitian matrices
?her2k	c, z	Rank-2k update of Hermitian matrices
?symm	s, d, c, z	Matrix-matrix product of symmetric matrices
?syrk	s, d, c, z	Rank-k update of symmetric matrices
?syr2k	s, d, c, z	Rank-2k update of symmetric matrices
?trmm	s, d, c, z	Matrix-matrix product of triangular matrices
?trsm	s, d, c, z	Linear matrix-matrix solution for triangular matrices

Symmetric Multiprocessing Version of Intel® MKL

Many applications spend considerable time for executing BLAS level 3 routines. This time can be scaled by the number of processors available on the system through using the symmetric multiprocessing (SMP) feature built into the Intel MKL Library. The performance enhancements based on the parallel use of the processors are available without any programming effort on your part.

To enhance performance, the library uses the following methods:

- The operation of BLAS level 3 matrix-matrix functions permits to restructure the code in a way which increases the localization of data reference, enhances cache memory use, and reduces the dependency on the memory bus.

- Once the code has been effectively blocked as described above, one of the matrices is distributed across the processors to be multiplied by the second matrix. Such distribution ensures effective cache management which reduces the dependency on the memory bus performance and brings good scaling results.

?gemm

Computes a scalar-matrix-matrix product and adds the result to a scalar-matrix product.

```
call sgemm ( transa, transb, m, n, k, alpha, a, lda,
             b, ldb, beta, c, ldc )
call dgemm ( transa, transb, m, n, k, alpha, a, lda,
             b, ldb, beta, c, ldc )
call cgemm ( transa, transb, m, n, k, alpha, a, lda,
             b, ldb, beta, c, ldc )
call zgemm ( transa, transb, m, n, k, alpha, a, lda,
             b, ldb, beta, c, ldc )
```

Discussion

The ?gemm routines perform a matrix-matrix operation with general matrices. The operation is defined as

$$c := \alpha * \text{op}(a) * \text{op}(b) + \beta * c,$$

where:

$\text{op}(x)$ is one of $\text{op}(x) = x$ or $\text{op}(x) = x'$ or $\text{op}(x) = \text{conjg}(x')$,

α and β are scalars

a, b and c are matrices:

$\text{op}(a)$ is an m by k matrix

$\text{op}(b)$ is a k by n matrix

c is an m by n matrix.

Input Parameters

transa CHARACTER*1. Specifies the form of $\text{op}(a)$ to be used in the matrix multiplication as follows:

<i>transa</i> value	Form of $\text{op}(a)$
N or n	$\text{op}(a) = a$
T or t	$\text{op}(a) = a'$
C or c	$\text{op}(a) = \text{conjg}(a')$

transb CHARACTER*1. Specifies the form of $\text{op}(b)$ to be used in the matrix multiplication as follows:

<i>transb</i> value	Form of $\text{op}(b)$
N or n	$\text{op}(b) = b$
T or t	$\text{op}(b) = b'$
C or c	$\text{op}(b) = \text{conjg}(b')$

m INTEGER. Specifies the number of rows of the matrix $\text{op}(a)$ and of the matrix *c*. The value of *m* must be at least zero.

n INTEGER. Specifies the number of columns of the matrix $\text{op}(b)$ and the number of columns of the matrix *c*. The value of *n* must be at least zero.

k INTEGER. Specifies the number of columns of the matrix $\text{op}(a)$ and the number of rows of the matrix $\text{op}(b)$. The value of *k* must be at least zero.

alpha REAL for *sgemm*
 DOUBLE PRECISION for *dgemm*
 COMPLEX for *cgemm*
 DOUBLE COMPLEX for *zgemm*
 Specifies the scalar *alpha*.

a REAL for *sgemm*
DOUBLE PRECISION for *dgemm*
COMPLEX for *cgemm*
DOUBLE COMPLEX for *zgemm*

Array, DIMENSION (*lda*, *ka*), where *ka* is *k* when *transa* = 'N' or 'n', and is *m* otherwise. Before entry with *transa* = 'N' or 'n', the leading *m* by *k* part of the array *a* must contain the matrix *a*, otherwise the leading *k* by *m* part of the array *a* must contain the matrix *a*.

lda INTEGER. Specifies the first dimension of *a* as declared in the calling (sub)program. When *transa* = 'N' or 'n', then *lda* must be at least $\max(1, m)$, otherwise *lda* must be at least $\max(1, k)$.

b REAL for *sgemm*
DOUBLE PRECISION for *dgemm*
COMPLEX for *cgemm*
DOUBLE COMPLEX for *zgemm*

Array, DIMENSION (*ldb*, *kb*), where *kb* is *n* when *transb* = 'N' or 'n', and is *k* otherwise. Before entry with *transb* = 'N' or 'n', the leading *k* by *n* part of the array *b* must contain the matrix *b*, otherwise the leading *n* by *k* part of the array *b* must contain the matrix *b*.

ldb INTEGER. Specifies the first dimension of *b* as declared in the calling (sub)program. When *transb* = 'N' or 'n', then *ldb* must be at least $\max(1, k)$, otherwise *ldb* must be at least $\max(1, n)$.

beta REAL for *sgemm*
DOUBLE PRECISION for *dgemm*
COMPLEX for *cgemm*
DOUBLE COMPLEX for *zgemm*

Specifies the scalar *beta*. When *beta* is supplied as zero, then *c* need not be set on input.

c REAL for `sgemm`
DOUBLE PRECISION for `dgemm`
COMPLEX for `cgemm`
DOUBLE COMPLEX for `zgemm`

Array, DIMENSION (*ldc*, *n*). Before entry, the leading *m* by *n* part of the array *c* must contain the matrix *c*, except when *beta* is zero, in which case *c* need not be set on entry.

ldc INTEGER. Specifies the first dimension of *c* as declared in the calling (sub)program. The value of *ldc* must be at least $\max(1, m)$.

Output Parameters

c Overwritten by the *m* by *n* matrix
($\alpha * \text{op}(a) * \text{op}(b) + \beta * c$).

?hemm

Computes a scalar-matrix-matrix product (either one of the matrices is Hermitian) and adds the result to scalar-matrix product.

```
call chemm ( side, uplo, m, n, alpha, a, lda, b,  
            ldb, beta, c, ldc )  
call zhemm ( side, uplo, m, n, alpha, a, lda, b,  
            ldb, beta, c, ldc )
```

Discussion

The ?hemm routines perform a matrix-matrix operation using Hermitian matrices. The operation is defined as

$c := \alpha * a * b + \beta * c$

or

$c := \alpha * b * a + \beta * c,$

where:

α and β are scalars

a is an Hermitian matrix

b and c are m by n matrices.

Input Parameters

$side$ CHARACTER*1. Specifies whether the Hermitian matrix a appears on the left or right in the operation as follows:

$side$ value	Operation To Be Performed
L or l	$c := \alpha * a * b + \beta * c$
R or r	$c := \alpha * b * a + \beta * c$

$uplo$ CHARACTER*1. Specifies whether the upper or lower triangular part of the Hermitian matrix a is to be referenced as follows:

$uplo$ value	Part of Matrix a To Be Referenced
U or u	Only the upper triangular part of the Hermitian matrix is to be referenced.
L or l	Only the lower triangular part of the Hermitian matrix is to be referenced.

m INTEGER. Specifies the number of rows of the matrix c . The value of m must be at least zero.

n INTEGER. Specifies the number of columns of the matrix c . The value of n must be at least zero.

α COMPLEX for chemm
DOUBLE COMPLEX for zhemm
Specifies the scalar α .

a **COMPLEX** for `chemm`
DOUBLE COMPLEX for `zhemm`

Array, **DIMENSION** (*lda*, *ka*), where *ka* is *m* when *side* = 'L' or 'l' and is *n* otherwise. Before entry with *side* = 'L' or 'l', the *m* by *m* part of the array *a* must contain the Hermitian matrix, such that when *uplo* = 'U' or 'u', the leading *m* by *m* upper triangular part of the array *a* must contain the upper triangular part of the Hermitian matrix and the strictly lower triangular part of *a* is not referenced, and when *uplo* = 'L' or 'l', the leading *m* by *m* lower triangular part of the array *a* must contain the lower triangular part of the Hermitian matrix, and the strictly upper triangular part of *a* is not referenced. Before entry with *side* = 'R' or 'r', the *n* by *n* part of the array *a* must contain the Hermitian matrix, such that when *uplo* = 'U' or 'u', the leading *n* by *n* upper triangular part of the array *a* must contain the upper triangular part of the Hermitian matrix and the strictly lower triangular part of *a* is not referenced, and when *uplo* = 'L' or 'l', the leading *n* by *n* lower triangular part of the array *a* must contain the lower triangular part of the Hermitian matrix, and the strictly upper triangular part of *a* is not referenced. The imaginary parts of the diagonal elements need not be set, they are assumed to be zero.

lda **INTEGER**. Specifies the first dimension of *a* as declared in the calling (sub) program. When *side* = 'L' or 'l' then *lda* must be at least `max(1, m)`, otherwise *lda* must be at least `max(1, n)`.

b **COMPLEX** for `chemm`
DOUBLE COMPLEX for `zhemm`

Array, **DIMENSION** (*ldb*, *n*). Before entry, the leading *m* by *n* part of the array *b* must contain the matrix *b*.

ldb **INTEGER**. Specifies the first dimension of *b* as declared in the calling (sub)program. The value of *ldb* must be at least `max(1, m)`.

<i>beta</i>	COMPLEX for chemm DOUBLE COMPLEX for zhemm Specifies the scalar <i>beta</i> . When <i>beta</i> is supplied as zero, then <i>c</i> need not be set on input.
<i>c</i>	COMPLEX for chemm DOUBLE COMPLEX for zhemm Array, DIMENSION (<i>c</i> , <i>n</i>). Before entry, the leading <i>m</i> by <i>n</i> part of the array <i>c</i> must contain the matrix <i>c</i> , except when <i>beta</i> is zero, in which case <i>c</i> need not be set on entry.
<i>ldc</i>	INTEGER. Specifies the first dimension of <i>c</i> as declared in the calling (sub)program. The value of <i>ldc</i> must be at least $\max(1, m)$.

Output Parameters

<i>c</i>	Overwritten by the <i>m</i> by <i>n</i> updated matrix.
----------	---

?herk

Performs a rank-n update of a Hermitian matrix.

```
call cherk ( uplo, trans, n, k, alpha, a, lda, beta, c,
             ldc )
call zherk ( uplo, trans, n, k, alpha, a, lda, beta, c,
             ldc )
```

Discussion

The ?herk routines perform a matrix-matrix operation using Hermitian matrices. The operation is defined as

$$c := \alpha * a * \text{conjg}(a') + \beta * c,$$

or

$$c := \alpha * \text{conjg}(a') * a + \beta * c,$$

where:

alpha and *beta* are real scalars

c is an *n* by *n* Hermitian matrix

a is an *n* by *k* matrix in the first case and a *k* by *n* matrix in the second case.

Input Parameters

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the array *c* is to be referenced as follows:

<i>uplo</i> value	Part of Array <i>c</i> To Be Referenced
U or u	Only the upper triangular part of <i>C</i> is to be referenced.
L or l	Only the lower triangular part of <i>C</i> is to be referenced.

trans CHARACTER*1. Specifies the operation to be performed as follows:

<i>trans</i> value	Operation to be Performed
N or n	$c := \alpha * a * \text{conjg}(a') + \beta * c$
C or c	$c := \alpha * \text{conjg}(a') * a + \beta * c$

n INTEGER. Specifies the order of the matrix *c*. The value of *n* must be at least zero.

k INTEGER. With *trans* = 'N' or 'n', *k* specifies the number of columns of the matrix *a*, and with *trans* = 'C' or 'c', *k* specifies the number of rows of the matrix *a*. The value of *k* must be at least zero.

alpha REAL for *cherk*
 DOUBLE PRECISION for *zherk*
 Specifies the scalar *alpha*.

<i>a</i>	<p>COMPLEX for <i>cherk</i> DOUBLE COMPLEX for <i>zherk</i></p> <p>Array, DIMENSION (<i>lda</i>, <i>ka</i>), where <i>ka</i> is <i>k</i> when <i>trans</i> = 'N' or 'n', and is <i>n</i> otherwise. Before entry with <i>trans</i> = 'N' or 'n', the leading <i>n</i> by <i>k</i> part of the array <i>a</i> must contain the matrix <i>a</i>, otherwise the leading <i>k</i> by <i>n</i> part of the array <i>a</i> must contain the matrix <i>a</i>.</p>
<i>lda</i>	<p>INTEGER. Specifies the first dimension of <i>a</i> as declared in the calling (sub)program. When <i>trans</i> = 'N' or 'n', then <i>lda</i> must be at least $\max(1, n)$, otherwise <i>lda</i> must be at least $\max(1, k)$.</p>
<i>beta</i>	<p>REAL for <i>cherk</i> DOUBLE PRECISION for <i>zherk</i></p> <p>Specifies the scalar <i>beta</i>.</p>
<i>c</i>	<p>COMPLEX for <i>cherk</i> DOUBLE COMPLEX for <i>zherk</i></p> <p>Array, DIMENSION (<i>ldc</i>, <i>n</i>). Before entry with <i>uplo</i> = 'U' or 'u', the leading <i>n</i> by <i>n</i> upper triangular part of the array <i>c</i> must contain the upper triangular part of the Hermitian matrix and the strictly lower triangular part of <i>c</i> is not referenced.</p> <p>Before entry with <i>uplo</i> = 'L' or 'l', the leading <i>n</i> by <i>n</i> lower triangular part of the array <i>c</i> must contain the lower triangular part of the Hermitian matrix and the strictly upper triangular part of <i>c</i> is not referenced.</p> <p>The imaginary parts of the diagonal elements need not be set, they are assumed to be zero.</p>
<i>ldc</i>	<p>INTEGER. Specifies the first dimension of <i>c</i> as declared in the calling (sub)program. The value of <i>ldc</i> must be at least $\max(1, n)$.</p>

Output Parameters

c With *uplo* = 'U' or 'u', the upper triangular part of the array *c* is overwritten by the upper triangular part of the updated matrix.

With *uplo* = 'L' or 'l', the lower triangular part of the array *c* is overwritten by the lower triangular part of the updated matrix.

The imaginary parts of the diagonal elements are set to zero.

?her2k

Performs a rank-2k update of a Hermitian matrix.

```
call cher2k ( uplo, trans, n, k, alpha, a, lda, b, ldb,  
             beta, c, ldc )  
call zher2k ( uplo, trans, n, k, alpha, a, lda, b, ldb,  
             beta, c, ldc )
```

Discussion

The ?her2k routines perform a rank-2k matrix-matrix operation using Hermitian matrices. The operation is defined as

$$c := \alpha * a * \text{conjg}(b') + \text{conjg}(\alpha) * b * \text{conjg}(a') + \text{beta} * c,$$

or

$$c := \alpha * \text{conjg}(b') * a + \text{conjg}(\alpha) * \text{conjg}(a') * b + \text{beta} * c,$$

where:

alpha is a scalar and *beta* is a real scalar

c is an *n* by *n* Hermitian matrix

a and *b* are *n* by *k* matrices in the first case and *k* by *n* matrices in the second case.

Input Parameters

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the array *c* is to be referenced as follows:

<i>uplo</i> value	Part of Array <i>c</i> To Be Referenced
U or u	Only the upper triangular part of <i>C</i> is to be referenced.
L or l	Only the lower triangular part of <i>C</i> is to be referenced.

trans CHARACTER*1. Specifies the operation to be performed as follows:

<i>trans</i> value	Operation to be Performed
N or n	$c := \alpha * a * \text{conjg}(b') + \alpha * b * \text{conjg}(a') + \beta * c$
C or c	$c := \alpha * \text{conjg}(a') * b + \alpha * \text{conjg}(b') * a + \beta * c$

n INTEGER. Specifies the order of the matrix *c*. The value of *n* must be at least zero.

k INTEGER. With *trans* = 'N' or 'n', *k* specifies the number of columns of the matrix *a*, and with *trans* = 'C' or 'c', *k* specifies the number of rows of the matrix *a*. The value of *k* must be at least zero.

alpha COMPLEX for cher2k
 DOUBLE COMPLEX for zher2k
 Specifies the scalar *alpha*.

<i>a</i>	<p>COMPLEX for <code>cher2k</code> DOUBLE COMPLEX for <code>zher2k</code></p> <p>Array, DIMENSION (<i>lda</i>, <i>ka</i>), where <i>ka</i> is <i>k</i> when <i>trans</i> = 'N' or 'n', and is <i>n</i> otherwise. Before entry with <i>trans</i> = 'N' or 'n', the leading <i>n</i> by <i>k</i> part of the array <i>a</i> must contain the matrix <i>a</i>, otherwise the leading <i>k</i> by <i>n</i> part of the array <i>a</i> must contain the matrix <i>a</i>.</p>
<i>lda</i>	<p>INTEGER. Specifies the first dimension of <i>a</i> as declared in the calling (sub)program. When <i>trans</i> = 'N' or 'n', then <i>lda</i> must be at least $\max(1, n)$, otherwise <i>lda</i> must be at least $\max(1, k)$.</p>
<i>beta</i>	<p>REAL for <code>cher2k</code> DOUBLE PRECISION for <code>zher2k</code></p> <p>Specifies the scalar <i>beta</i>.</p>
<i>b</i>	<p>COMPLEX for <code>cher2k</code> DOUBLE COMPLEX for <code>zher2k</code></p> <p>Array, DIMENSION (<i>ldb</i>, <i>kb</i>), where <i>kb</i> is <i>k</i> when <i>trans</i> = 'N' or 'n', and is <i>n</i> otherwise. Before entry with <i>trans</i> = 'N' or 'n', the leading <i>n</i> by <i>k</i> part of the array <i>b</i> must contain the matrix <i>b</i>, otherwise the leading <i>k</i> by <i>n</i> part of the array <i>b</i> must contain the matrix <i>b</i>.</p>
<i>ldb</i>	<p>INTEGER. Specifies the first dimension of <i>b</i> as declared in the calling (sub)program. When <i>trans</i> = 'N' or 'n', then <i>ldb</i> must be at least $\max(1, n)$, otherwise <i>ldb</i> must be at least $\max(1, k)$.</p>
<i>c</i>	<p>COMPLEX for <code>cher2k</code> DOUBLE COMPLEX for <code>zher2k</code></p> <p>Array, DIMENSION (<i>ldc</i>, <i>n</i>). Before entry with <i>uplo</i> = 'U' or 'u', the leading <i>n</i> by <i>n</i> upper triangular part of the array <i>c</i> must contain the upper triangular part of the Hermitian matrix and the strictly lower triangular part of <i>c</i> is not referenced.</p>

Before entry with `uplo = 'L'` or `'l'`, the leading n by n lower triangular part of the array `c` must contain the lower triangular part of the Hermitian matrix and the strictly upper triangular part of `c` is not referenced.

The imaginary parts of the diagonal elements need not be set, they are assumed to be zero.

`ldc` **INTEGER**. Specifies the first dimension of `c` as declared in the calling (sub)program. The value of `ldc` must be at least $\max(1, n)$.

Output Parameters

`c` With `uplo = 'U'` or `'u'`, the upper triangular part of the array `c` is overwritten by the upper triangular part of the updated matrix.

With `uplo = 'L'` or `'l'`, the lower triangular part of the array `c` is overwritten by the lower triangular part of the updated matrix.

The imaginary parts of the diagonal elements are set to zero.

?symm

Performs a scalar-matrix-matrix product (one matrix operand is symmetric) and adds the result to a scalar-matrix product.

```
call ssymm ( side, uplo, m, n, alpha, a, lda, b, ldb,  
            beta, c, ldc )  
call dsymm ( side, uplo, m, n, alpha, a, lda, b, ldb,  
            beta, c, ldc )  
call csymm ( side, uplo, m, n, alpha, a, lda, b, ldb,  
            beta, c, ldc )  
call zsymm ( side, uplo, m, n, alpha, a, lda, b, ldb,  
            beta, c, ldc )
```

Discussion

The ?symm routines perform a matrix-matrix operation using symmetric matrices. The operation is defined as

$$c := \alpha * a * b + \beta * c,$$

or

$$c := \alpha * b * a + \beta * c,$$

where:

α and β are scalars

a is a symmetric matrix

b and c are m by n matrices.

Input Parameters

side CHARACTER*1. Specifies whether the symmetric matrix *a* appears on the left or right in the operation as follows:

<i>side</i> value	Operation to be Performed
L or l	$c := \alpha * a * b + \beta * c$
R or r	$c := \alpha * b * a + \beta * c$

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the symmetric matrix *a* is to be referenced as follows:

<i>uplo</i> value	Part of Array <i>a</i> To Be Referenced
U or u	Only the upper triangular part of the symmetric matrix is to be referenced.
L or l	Only the lower triangular part of the symmetric matrix is to be referenced.

m INTEGER. Specifies the number of rows of the matrix *c*. The value of *m* must be at least zero.

n INTEGER. Specifies the number of columns of the matrix *c*. The value of *n* must be at least zero.

alpha REAL for *ssymm*
 DOUBLE PRECISION for *dsymm*
 COMPLEX for *csymm*
 DOUBLE COMPLEX for *zsymm*
 Specifies the scalar *alpha*.

a REAL for *ssymm*
 DOUBLE PRECISION for *dsymm*
 COMPLEX for *csymm*
 DOUBLE COMPLEX for *zsymm*
 Array, DIMENSION (*lda*, *ka*), where *ka* is *m* when *side* = 'L' or 'l' and is *n* otherwise. Before entry with *side* = 'L' or 'l', the *m* by *m* part of the array *a* must contain the symmetric matrix, such that when

`uplo = 'U'` or `'u'`, the leading m by m upper triangular part of the array `a` must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of `a` is not referenced, and when `uplo = 'L'` or `'l'`, the leading m by m lower triangular part of the array `a` must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of `a` is not referenced.

Before entry with `side = 'R'` or `'r'`, the n by n part of the array `a` must contain the symmetric matrix, such that when `uplo = 'U'` or `'u'`, the leading n by n upper triangular part of the array `a` must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of `a` is not referenced, and when `uplo = 'L'` or `'l'`, the leading n by n lower triangular part of the array `a` must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of `a` is not referenced.

`lda` **INTEGER**. Specifies the first dimension of `a` as declared in the calling (sub)program. When `side = 'L'` or `'l'` then `lda` must be at least $\max(1, m)$, otherwise `lda` must be at least $\max(1, n)$.

`b` **REAL** for `ssymm`
DOUBLE PRECISION for `dsymm`
COMPLEX for `csymm`
DOUBLE COMPLEX for `zsymm`

Array, **DIMENSION** (`ldb, n`). Before entry, the leading m by n part of the array `b` must contain the matrix `b`.

`ldb` **INTEGER**. Specifies the first dimension of `b` as declared in the calling (sub)program. The value of `ldb` must be at least $\max(1, m)$.

<i>beta</i>	REAL for <i>ssymm</i> DOUBLE PRECISION for <i>dsymm</i> COMPLEX for <i>csymm</i> DOUBLE COMPLEX for <i>zsymm</i> Specifies the scalar <i>beta</i> . When <i>beta</i> is supplied as zero, then <i>c</i> need not be set on input.
<i>c</i>	REAL for <i>ssymm</i> DOUBLE PRECISION for <i>dsymm</i> COMPLEX for <i>csymm</i> DOUBLE COMPLEX for <i>zsymm</i> Array, DIMENSION (<i>ldc</i> , <i>n</i>). Before entry, the leading <i>m</i> by <i>n</i> part of the array <i>c</i> must contain the matrix <i>c</i> , except when <i>beta</i> is zero, in which case <i>c</i> need not be set on entry.
<i>ldc</i>	INTEGER. Specifies the first dimension of <i>c</i> as declared in the calling (sub)program. The value of <i>ldc</i> must be at least $\max(1, m)$.

Output Parameters

<i>c</i>	Overwritten by the <i>m</i> by <i>n</i> updated matrix.
----------	---

?syrk

Performs a rank- n update of a symmetric matrix.

```
call ssyrk ( uplo, trans, n, k, alpha, a, lda, beta, c,
             ldc )
call dsyrk ( uplo, trans, n, k, alpha, a, lda, beta, c,
             ldc )
call csyrk ( uplo, trans, n, k, alpha, a, lda, beta, c,
             ldc )
call zsyrk ( uplo, trans, n, k, alpha, a, lda, beta, c,
             ldc )
```

Discussion

The ?syrk routines perform a matrix-matrix operation using symmetric matrices. The operation is defined as

$$c := \alpha * a * a' + \beta * c,$$

or

$$c := \alpha * a' * a + \beta * c,$$

where:

α and β are scalars

c is an n by n symmetric matrix

a is an n by k matrix in the first case and a k by n matrix in the second case.

Input Parameters

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the array *c* is to be referenced as follows:

<i>uplo</i> value	Part of Array <i>c</i> To Be Referenced
U or u	Only the upper triangular part of <i>c</i> is to be referenced.
L or l	Only the lower triangular part of <i>c</i> is to be referenced.

trans CHARACTER*1. Specifies the operation to be performed as follows:

<i>trans</i> value	Operation to be Performed
N or n	$c := \alpha * a * a' + \beta * c$
T or t	$c := \alpha * a' * a + \beta * c$
C or c	$c := \alpha * a' * a + \beta * c$

n INTEGER. Specifies the order of the matrix *c*. The value of *n* must be at least zero.

k INTEGER. On entry with *trans* = 'N' or 'n', *k* specifies the number of columns of the matrix *a*, and on entry with *trans* = 'T' or 't' or 'C' or 'c', *k* specifies the number of rows of the matrix *a*. The value of *k* must be at least zero.

alpha REAL for *ssyrk*
 DOUBLE PRECISION for *dsyrk*
 COMPLEX for *csyrk*
 DOUBLE COMPLEX for *zsyrk*
 Specifies the scalar *alpha*.

<i>a</i>	<p>REAL for <i>ssyrk</i> DOUBLE PRECISION for <i>dsyrk</i> COMPLEX for <i>csyrk</i> DOUBLE COMPLEX for <i>zsyrk</i></p> <p>Array, DIMENSION (<i>lda</i>, <i>ka</i>), where <i>ka</i> is <i>k</i> when <i>trans</i> = 'N' or 'n', and is <i>n</i> otherwise. Before entry with <i>trans</i> = 'N' or 'n', the leading <i>n</i> by <i>k</i> part of the array <i>a</i> must contain the matrix <i>a</i>, otherwise the leading <i>k</i> by <i>n</i> part of the array <i>a</i> must contain the matrix <i>a</i>.</p>
<i>lda</i>	<p>INTEGER. Specifies the first dimension of <i>a</i> as declared in the calling (sub)program. When <i>trans</i> = 'N' or 'n', then <i>lda</i> must be at least $\max(1, n)$, otherwise <i>lda</i> must be at least $\max(1, k)$.</p>
<i>beta</i>	<p>REAL for <i>ssyrk</i> DOUBLE PRECISION for <i>dsyrk</i> COMPLEX for <i>csyrk</i> DOUBLE COMPLEX for <i>zsyrk</i></p> <p>Specifies the scalar <i>beta</i>.</p>
<i>c</i>	<p>REAL for <i>ssyrk</i> DOUBLE PRECISION for <i>dsyrk</i> COMPLEX for <i>csyrk</i> DOUBLE COMPLEX for <i>zsyrk</i></p> <p>Array, DIMENSION (<i>ldc</i>, <i>n</i>). Before entry with <i>uplo</i> = 'U' or 'u', the leading <i>n</i> by <i>n</i> upper triangular part of the array <i>c</i> must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of <i>c</i> is not referenced.</p> <p>Before entry with <i>uplo</i> = 'L' or 'l', the leading <i>n</i> by <i>n</i> lower triangular part of the array <i>c</i> must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of <i>c</i> is not referenced.</p>
<i>ldc</i>	<p>INTEGER. Specifies the first dimension of <i>c</i> as declared in the calling (sub)program. The value of <i>ldc</i> must be at least $\max(1, n)$.</p>

Output Parameters

c With *uplo* = 'U' or 'u', the upper triangular part of the array *c* is overwritten by the upper triangular part of the updated matrix.

With *uplo* = 'L' or 'l', the lower triangular part of the array *c* is overwritten by the lower triangular part of the updated matrix.

?syr2k

Performs a rank-2k update of a symmetric matrix.

```
call ssyr2k ( uplo, trans, n, k, alpha, a, lda, b, ldb,
             beta, c, ldc )
call dsyr2k ( uplo, trans, n, k, alpha, a, lda, b, ldb,
             beta, c, ldc )
call csyr2k ( uplo, trans, n, k, alpha, a, lda, b, ldb,
             beta, c, ldc )
call zsyr2k ( uplo, trans, n, k, alpha, a, lda, b, ldb,
             beta, c, ldc )
```

Discussion

The ?syr2k routines perform a rank-2k matrix-matrix operation using symmetric matrices. The operation is defined as

$$c := \alpha * a * b' + \alpha * b * a' + \beta * c,$$

or

$$c := \alpha * a' * b + \alpha * b' * a + \beta * c,$$

where:

alpha and *beta* are scalars

c is an *n* by *n* symmetric matrix

a and b are n by k matrices in the first case and k by n matrices in the second case.

Input Parameters

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the array c is to be referenced as follows:

<i>uplo</i> value	Part of Array c To Be Referenced
U or u	Only the upper triangular part of c is to be referenced.
L or l	Only the lower triangular part of c is to be referenced.

trans CHARACTER*1. Specifies the operation to be performed as follows:

<i>trans</i> value	Operation to be Performed
N or n	$c := \alpha * a * b' + \alpha * b * a' + \beta * c$
T or t	$c := \alpha * a' * b + \alpha * b' * a + \beta * c$
C or c	$c := \alpha * a' * b + \alpha * b' * a + \beta * c$

n INTEGER. Specifies the order of the matrix c . The value of n must be at least zero.

k INTEGER. On entry with $trans = 'N'$ or $'n'$, k specifies the number of columns of the matrices a and b , and on entry with $trans = 'T'$ or $'t'$ or $'C'$ or $'c'$, k specifies the number of rows of the matrices a and b . The value of k must be at least zero.

alpha REAL for `ssyr2k`
 DOUBLE PRECISION for `dsyr2k`
 COMPLEX for `csyr2k`
 DOUBLE COMPLEX for `zsyr2k`
 Specifies the scalar α .

a REAL for `ssyr2k`
DOUBLE PRECISION for `dsyr2k`
COMPLEX for `csyr2k`
DOUBLE COMPLEX for `zsyr2k`

Array, DIMENSION (*lda*, *ka*), where *ka* is *k* when *trans* = 'N' or 'n', and is *n* otherwise. Before entry with *trans* = 'N' or 'n', the leading *n* by *k* part of the array *a* must contain the matrix *a*, otherwise the leading *k* by *n* part of the array *a* must contain the matrix *a*.

lda INTEGER. Specifies the first dimension of *a* as declared in the calling (sub)program. When *trans* = 'N' or 'n', then *lda* must be at least $\max(1, n)$, otherwise *lda* must be at least $\max(1, k)$.

b REAL for `ssyr2k`
DOUBLE PRECISION for `dsyr2k`
COMPLEX for `csyr2k`
DOUBLE COMPLEX for `zsyr2k`

Array, DIMENSION (*ldb*, *kb*) where *kb* is *k* when *trans* = 'N' or 'n' and is 'n' otherwise. Before entry with *trans* = 'N' or 'n', the leading *n* by *k* part of the array *b* must contain the matrix *b*, otherwise the leading *k* by *n* part of the array *b* must contain the matrix *b*.

ldb INTEGER. Specifies the first dimension of *a* as declared in the calling (sub)program. When *trans* = 'N' or 'n', then *ldb* must be at least $\max(1, n)$, otherwise *ldb* must be at least $\max(1, k)$.

beta REAL for `ssyr2k`
DOUBLE PRECISION for `dsyr2k`
COMPLEX for `csyr2k`
DOUBLE COMPLEX for `zsyr2k`

Specifies the scalar *beta*.

c REAL for `ssyr2k`
DOUBLE PRECISION for `dsyr2k`
COMPLEX for `csyr2k`
DOUBLE COMPLEX for `zsyr2k`

Array, DIMENSION (*ldc*, *n*). Before entry with *uplo* = 'U' or 'u', the leading *n* by *n* upper triangular part of the array *c* must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of *c* is not referenced.

Before entry with *uplo* = 'L' or 'l', the leading *n* by *n* lower triangular part of the array *c* must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of *c* is not referenced.

ldc INTEGER. Specifies the first dimension of *c* as declared in the calling (sub)program. The value of *ldc* must be at least `max(1, n)`.

Output Parameters

c With *uplo* = 'U' or 'u', the upper triangular part of the array *c* is overwritten by the upper triangular part of the updated matrix.

With *uplo* = 'L' or 'l', the lower triangular part of the array *c* is overwritten by the lower triangular part of the updated matrix.

?trmm

Computes a scalar-matrix-matrix product (one matrix operand is triangular).

```
call strmm ( side, uplo, transa, diag, m, n, alpha, a,  
            lda, b, ldb )  
call dtrmm ( side, uplo, transa, diag, m, n, alpha, a,  
            lda, b, ldb )  
call ctrmm ( side, uplo, transa, diag, m, n, alpha, a,  
            lda, b, ldb )  
call ztrmm ( side, uplo, transa, diag, m, n, alpha, a,  
            lda, b, ldb )
```

Discussion

The `?trmm` routines perform a matrix-matrix operation using triangular matrices. The operation is defined as

`b := alpha*op(a)*b`

or

`b := alpha*b*op(a)`

where:

`alpha` is a scalar

`b` is an m by n matrix

`a` is a unit, or non-unit, upper or lower triangular matrix

`op(a)` is one of `op(a) = a` or `op(a) = a'` or `op(a) = conjg(a')`.

Input Parameters

side CHARACTER*1. Specifies whether $\text{op}(a)$ multiplies b from the left or right in the operation as follows:

<i>side</i> value	Operation To Be Performed
L or l	$b := \text{alpha} * \text{op}(a) * b$
R or r	$b := \text{alpha} * b * \text{op}(a)$

uplo CHARACTER*1. Specifies whether the matrix a is an upper or lower triangular matrix as follows:

<i>uplo</i> value	Matrix a
U or u	Matrix a is an upper triangular matrix.
L or l	Matrix a is a lower triangular matrix.

transa CHARACTER*1. Specifies the form of $\text{op}(a)$ to be used in the matrix multiplication as follows:

<i>transa</i> value	Form of $\text{op}(a)$
N or n	$\text{op}(a) = a$
T or t	$\text{op}(a) = a'$
C or c	$\text{op}(a) = \text{conjg}(a')$

diag CHARACTER*1. Specifies whether or not a is unit triangular as follows:

<i>diag</i> value	Matrix a
U or u	Matrix a is assumed to be unit triangular.
N or n	Matrix a is not assumed to be unit triangular.

m INTEGER. Specifies the number of rows of b . The value of m must be at least zero.

n INTEGER. Specifies the number of columns of b . The value of n must be at least zero.

alpha REAL for *strmm*
DOUBLE PRECISION for *dtrmm*
COMPLEX for *ctrmm*
DOUBLE COMPLEX for *ztrmm*

Specifies the scalar *alpha*. When *alpha* is zero, then *a* is not referenced and *b* need not be set before entry.

a REAL for *strmm*
DOUBLE PRECISION for *dtrmm*
COMPLEX for *ctrmm*
DOUBLE COMPLEX for *ztrmm*

Array, DIMENSION (*lda*,*k*), where *k* is *m* when *side* = 'L' or 'l' and is *n* when *side* = 'R' or 'r'. Before entry with *uplo* = 'U' or 'u', the leading *k* by *k* upper triangular part of the array *a* must contain the upper triangular matrix and the strictly lower triangular part of *a* is not referenced.

Before entry with *uplo* = 'L' or 'l', the leading *k* by *k* lower triangular part of the array *a* must contain the lower triangular matrix and the strictly upper triangular part of *a* is not referenced. When *diag* = 'U' or 'u', the diagonal elements of *a* are not referenced either, but are assumed to be unity.

lda INTEGER. Specifies the first dimension of *a* as declared in the calling (sub)program. When *side* = 'L' or 'l', then *lda* must be at least $\max(1, m)$, when *side* = 'R' or 'r', then *lda* must be at least $\max(1, n)$.

b REAL for *strmm*
DOUBLE PRECISION for *dtrmm*
COMPLEX for *ctrmm*
DOUBLE COMPLEX for *ztrmm*

Array, DIMENSION (*ldb*,*n*). Before entry, the leading *m* by *n* part of the array *b* must contain the matrix *b*.

ldb INTEGER. Specifies the first dimension of *b* as declared in the calling (sub)program. The value of *ldb* must be at least $\max(1, m)$.

Output Parameters

b Overwritten by the transformed matrix.

?trsm

Solves a matrix equation (one matrix operand is triangular).

```
call strsm ( side, uplo, transa, diag, m, n, alpha, a,
            lda, b, ldb )
call dtrsm ( side, uplo, transa, diag, m, n, alpha, a,
            lda, b, ldb )
call ctrsm ( side, uplo, transa, diag, m, n, alpha, a,
            lda, b, ldb )
call ztrsm ( side, uplo, transa, diag, m, n, alpha, a,
            lda, b, ldb )
```

Discussion

The ?trsm routines solve one of the following matrix equations:

$op(a)*x = alpha*b,$

or

$x*op(a) = alpha*b,$

where:

alpha is a scalar

x and *b* are *m* by *n* matrices

a is a unit, or non-unit, upper or lower triangular matrix

op(a) is one of $op(a) = a$ or $op(a) = a'$ or

$op(a) = conjg(a')$.

The matrix *x* is overwritten on *b*.

Input Parameters

side CHARACTER*1. Specifies whether $\text{op}(a)$ appears on the left or right of x for the operation to be performed as follows:

<i>side</i> value	Operation To Be Performed
L or l	$\text{op}(a)*x = \alpha*b$
R or r	$x*\text{op}(a) = \alpha*b$

uplo CHARACTER*1. Specifies whether the matrix a is an upper or lower triangular matrix as follows:

<i>uplo</i> value	Matrix a
U or u	Matrix a is an upper triangular matrix.
L or l	Matrix a is a lower triangular matrix.

transa CHARACTER*1. Specifies the form of $\text{op}(a)$ to be used in the matrix multiplication as follows:

<i>transa</i> value	Form of $\text{op}(a)$
N or n	$\text{op}(a) = a$
T or t	$\text{op}(a) = a'$
C or c	$\text{op}(a) = \text{conjg}(a')$

diag CHARACTER*1. Specifies whether or not a is unit triangular as follows:

<i>diag</i> value	Matrix a
U or u	Matrix a is assumed to be unit triangular.
N or n	Matrix a is not assumed to be unit triangular.

m INTEGER. Specifies the number of rows of b . The value of m must be at least zero.

n INTEGER. Specifies the number of columns of b . The value of n must be at least zero.

<i>alpha</i>	<p>REAL for <i>strsm</i> DOUBLE PRECISION for <i>dtrsm</i> COMPLEX for <i>ctrsm</i> DOUBLE COMPLEX for <i>ztrsm</i></p> <p>Specifies the scalar <i>alpha</i>. When <i>alpha</i> is zero, then <i>a</i> is not referenced and <i>b</i> need not be set before entry.</p>
<i>a</i>	<p>REAL for <i>strsm</i> DOUBLE PRECISION for <i>dtrsm</i> COMPLEX for <i>ctrsm</i> DOUBLE COMPLEX for <i>ztrsm</i></p> <p>Array, DIMENSION (<i>lda</i>, <i>k</i>), where <i>k</i> is <i>m</i> when <i>side</i> = 'L' or 'l' and is <i>n</i> when <i>side</i> = 'R' or 'r'. Before entry with <i>uplo</i> = 'U' or 'u', the leading <i>k</i> by <i>k</i> upper triangular part of the array <i>a</i> must contain the upper triangular matrix and the strictly lower triangular part of <i>a</i> is not referenced.</p> <p>Before entry with <i>uplo</i> = 'L' or 'l', the leading <i>k</i> by <i>k</i> lower triangular part of the array <i>a</i> must contain the lower triangular matrix and the strictly upper triangular part of <i>a</i> is not referenced. When <i>diag</i> = 'U' or 'u', the diagonal elements of <i>a</i> are not referenced either, but are assumed to be unity.</p>
<i>lda</i>	<p>INTEGER. Specifies the first dimension of <i>a</i> as declared in the calling (sub)program. When <i>side</i> = 'L' or 'l', then <i>lda</i> must be at least $\max(1, m)$, when <i>side</i> = 'R' or 'r', then <i>lda</i> must be at least $\max(1, n)$.</p>
<i>b</i>	<p>REAL for <i>strsm</i> DOUBLE PRECISION for <i>dtrsm</i> COMPLEX for <i>ctrsm</i> DOUBLE COMPLEX for <i>ztrsm</i></p> <p>Array, DIMENSION (<i>ldb</i>, <i>n</i>). Before entry, the leading <i>m</i> by <i>n</i> part of the array <i>b</i> must contain the right-hand side matrix <i>b</i>.</p>

ldb **INTEGER**. Specifies the first dimension of *b* as declared in the calling (sub)program. The value of *ldb* must be at least $\max(1, m)$.

Output Parameters

b Overwritten by the solution matrix *x*.

Sparse BLAS Routines and Functions

This section describes Sparse BLAS, an extension of BLAS Level 1 included in Intel® Math Kernel Library beginning with Intel MKL release 2.1. Sparse BLAS is a group of routines and functions that perform a number of common vector operations on sparse vectors stored in compressed form.

Sparse vectors are those in which the majority of elements are zeros. Sparse BLAS routines and functions are specially implemented to take advantage of vector sparsity. This allows you to achieve large savings in computer time and memory. If *nz* is the number of non-zero vector elements, the computer time taken by Sparse BLAS operations will be $O(nz)$.

Vector Arguments in Sparse BLAS

Compressed sparse vectors. Let *a* be a vector stored in an array, and assume that the only non-zero elements of *a* are the following:

$$a(k_1), a(k_2), a(k_3) \dots a(k_{nz}),$$

where *nz* is the total number of non-zero elements in *a*.

In Sparse BLAS, this vector can be represented in compressed form by two FORTRAN arrays, *x* (values) and *indx* (indices). Each array has *nz* elements:

$$x(1)=a(k_1), x(2)=a(k_2), \dots x(nz)=a(k_{nz}),$$

$$indx(1)=k_1, indx(2)=k_2, \dots indx(nz)=k_{nz}.$$

Thus, a sparse vector is fully determined by the triple (*nz*, *x*, *indx*). If you pass a negative or zero value of *nz* to Sparse BLAS, the subroutines do not modify any arrays or variables.

Full-storage vectors. Sparse BLAS routines can also use a vector argument fully stored in a single FORTRAN array (a full-storage vector). If *y* is a full-storage vector, its elements must be stored contiguously: the first element in *y*(1), the second in *y*(2), and so on. This corresponds to an increment *incy* = 1 in BLAS Level 1. No increment value for full-storage vectors is passed as an argument to Sparse BLAS routines or functions.

Naming Conventions in Sparse BLAS

Similar to BLAS, the names of Sparse BLAS subprograms have prefixes that determine the data type involved: **s** and **d** for single- and double-precision real; **c** and **z** for single- and double-precision complex.

If a Sparse BLAS routine is an extension of a “dense” one, the subprogram name is formed by appending the suffix **i** (standing for *indexed*) to the name of the corresponding “dense” subprogram. For example, the Sparse BLAS routine **saxpyi** corresponds to the BLAS routine **saxpy**, and the Sparse BLAS function **cdotci** corresponds to the BLAS function **cdotc**.

Routines and Data Types in Sparse BLAS

Routines and data types supported in the Intel MKL implementation of Sparse BLAS are listed in Table 2-4.

Table 2-4 Sparse BLAS Routines and Their Data Types

Routine/ Function	Data Types	Description
<u>?axpyi</u>	s, d, c, z	Scalar-vector product plus vector (routines)
<u>?doti</u>	s, d	Dot product (functions)
<u>?dotci</u>	c, z	Complex dot product conjugated (functions)
<u>?dotui</u>	c, z	Complex dot product unconjugated (functions)
<u>?gthr</u>	s, d, c, z	Gathering a full-storage sparse vector into compressed form: nz , x , indx (routines)
<u>?gthrz</u>	s, d, c, z	Gathering a full-storage sparse vector into compressed form and assigning zeros to gathered elements in the full-storage vector (routines)
<u>?roti</u>	s, d	Givens rotation (routines)
<u>?sctr</u>	s, d, c, z	Scattering a vector from compressed form to full-storage form (routines)

BLAS Routines That Can Work With Sparse Vectors

The following BLAS Level 1 routines will give correct results when you pass to them a compressed-form array x (with the increment $incx = 1$):

- `?asum` sum of absolute values of vector elements
- `?copy` copying a vector
- `?nrm2` Euclidean norm of a vector
- `?scal` scaling a vector
- `i?amax` index of the element with the largest absolute value or,
for complex flavors, the largest sum $| \text{Re}x(i) | + | \text{Im}x(i) |$.
- `i?amin` index of the element with the smallest absolute value or,
for complex flavors, the smallest sum $| \text{Re}x(i) | + | \text{Im}x(i) |$.

The result i returned by `i?amax` and `i?amin` should be interpreted as index in the compressed-form array, so that the largest (smallest) value is $x(i)$; the corresponding index in full-storage array is $indx(i)$.

You can also call `?rotg` to compute the parameters of Givens rotation and then pass these parameters to the Sparse BLAS routines `?roti`.

?axpyi

Adds a scalar multiple of compressed sparse vector to a full-storage vector.

```
call saxpyi ( nz, a, x, indx, y )
call daxpyi ( nz, a, x, indx, y )
call caxpyi ( nz, a, x, indx, y )
call zaxpyi ( nz, a, x, indx, y )
```

Discussion

The `?axpyi` routines perform a vector-vector operation defined as

$$y := a*x + y$$

where:

a is a scalar

(*nz*, *x*, *indx*) is a sparse vector stored in compressed form

y is a vector in full storage form.

The ?axpyi routines reference or modify only the elements of *y* whose indices are listed in the array *indx*. The values in *indx* must be distinct.

Input Parameters

<i>nz</i>	INTEGER. The number of elements in <i>x</i> and <i>indx</i> .
<i>a</i>	REAL for saxpyi DOUBLE PRECISION for daxpyi COMPLEX for caxpyi DOUBLE COMPLEX for zaxpyi Specifies the scalar <i>a</i> .
<i>x</i>	REAL for saxpyi DOUBLE PRECISION for daxpyi COMPLEX for caxpyi DOUBLE COMPLEX for zaxpyi Array, DIMENSION at least <i>nz</i> .
<i>indx</i>	INTEGER. Specifies the indices for the elements of <i>x</i> . Array, DIMENSION at least <i>nz</i> .
<i>y</i>	REAL for saxpyi DOUBLE PRECISION for daxpyi COMPLEX for caxpyi DOUBLE COMPLEX for zaxpyi Array, DIMENSION at least $\max_i (\text{indx}(i))$.

Output Parameters

<i>y</i>	Contains the updated vector <i>y</i> .
----------	--

?doti

Computes the dot product of a compressed sparse real vector by a full-storage real vector.

```
res = sdoti ( nz, x, indx, y )
res = ddoti ( nz, x, indx, y )
```

Discussion

The `?doti` functions return the dot product of x and y defined as $x(1)*y(indx(1)) + x(2)*y(indx(2)) + \dots + x(nz)*y(indx(nz))$ where the triple $(nz, x, indx)$ defines a sparse real vector stored in compressed form, and y is a real vector in full storage form. The functions reference only the elements of y whose indices are listed in the array $indx$. The values in $indx$ must be distinct.

Input Parameters

<code>nz</code>	INTEGER. The number of elements in x and $indx$.
<code>x</code>	REAL for <code>sdoti</code> DOUBLE PRECISION for <code>ddoti</code> Array, DIMENSION at least nz .
<code>indx</code>	INTEGER. Specifies the indices for the elements of x . Array, DIMENSION at least nz .
<code>y</code>	REAL for <code>sdoti</code> DOUBLE PRECISION for <code>ddoti</code> Array, DIMENSION at least $\max_i(indx(i))$.

Output Parameters

<code>res</code>	REAL for <code>sdoti</code> DOUBLE PRECISION for <code>ddoti</code> Contains the dot product of x and y , if nz is positive. Otherwise, <code>res</code> contains 0.
------------------	---

?dotci

*Computes the conjugated dot product of
a compressed sparse complex vector
with a full-storage complex vector.*

```
res = cdotci ( nz, x, indx, y )  
res = zdotci ( nz, x, indx, y )
```

Discussion

The ?dotci functions return the dot product of x and y defined as $\text{conjg}(x(1))*y(\text{indx}(1)) + \dots + \text{conjg}(x(\text{nz}))*y(\text{indx}(\text{nz}))$ where the triple (nz, x, indx) defines a sparse complex vector stored in compressed form, and y is a real vector in full storage form. The functions reference only the elements of y whose indices are listed in the array indx . The values in indx must be distinct.

Input Parameters

nz **INTEGER**. The number of elements in x and indx .

x **COMPLEX** for **cdotci**
DOUBLE COMPLEX for **zdotci**
Array, **DIMENSION** at least nz .

indx **INTEGER**. Specifies the indices for the elements of x .
Array, **DIMENSION** at least nz .

y **COMPLEX** for **cdotci**
DOUBLE COMPLEX for **zdotci**
Array, **DIMENSION** at least $\max_i(\text{indx}(i))$.

Output Parameters

res **COMPLEX** for **cdotci**
DOUBLE COMPLEX for **zdotci**
Contains the conjugated dot product of x and y ,
if nz is positive. Otherwise, res contains 0.

?dotui

Computes the dot product of a compressed sparse complex vector by a full-storage complex vector.

```
res = cdotui ( nz, x, indx, y )
res = zdotui ( nz, x, indx, y )
```

Discussion

The ?dotui functions return the dot product of x and y defined as $x(1)*y(indx(1)) + x(2)*y(indx(2)) + \dots + x(nz)*y(indx(nz))$ where the triple $(nz, x, indx)$ defines a sparse complex vector stored in compressed form, and y is a real vector in full storage form. The functions reference only the elements of y whose indices are listed in the array $indx$. The values in $indx$ must be distinct.

Input Parameters

nz	INTEGER. The number of elements in x and $indx$.
x	COMPLEX for cdotui DOUBLE COMPLEX for zdotui Array, DIMENSION at least nz .
$indx$	INTEGER. Specifies the indices for the elements of x . Array, DIMENSION at least nz .
y	COMPLEX for cdotui DOUBLE COMPLEX for zdotui Array, DIMENSION at least $\max_i(indx(i))$.

Output Parameters

res	COMPLEX for cdotui DOUBLE COMPLEX for zdotui Contains the dot product of x and y , if nz is positive. Otherwise, res contains 0.
-------	---

?gthr

Gathers a full-storage sparse vector's elements into compressed form.

```
call sgthr ( nz, y, x, indx )
call dgthr ( nz, y, x, indx )
call cgthr ( nz, y, x, indx )
call zgthr ( nz, y, x, indx )
```

Discussion

The `?gthr` routines gather the specified elements of a full-storage sparse vector y into compressed form $(nz, x, indx)$. The routines reference only the elements of y whose indices are listed in the array $indx$:

$x(i) = y(indx(i))$, for $i=1,2,\dots,nz$.

Input Parameters

nz **INTEGER**. The number of elements of y to be gathered.

$indx$ **INTEGER**. Specifies indices of elements to be gathered.
Array, **DIMENSION** at least nz .

y **REAL** for `sgthr`
DOUBLE PRECISION for `dgthr`
COMPLEX for `cgthr`
DOUBLE COMPLEX for `zgthr`
Array, **DIMENSION** at least $\max_i (indx(i))$.

Output Parameters

x **REAL** for `sgthr`
DOUBLE PRECISION for `dgthr`
COMPLEX for `cgthr`
DOUBLE COMPLEX for `zgthr`
Array, **DIMENSION** at least nz .
Contains the vector converted to the compressed form.

?gthrz

Gathers a sparse vector's elements into compressed form, replacing them by zeros.

```
call sgthrz ( nz, y, x, indx )
call dgthrz ( nz, y, x, indx )
call cgthrz ( nz, y, x, indx )
call zgthrz ( nz, y, x, indx )
```

Discussion

The `?gthrz` routines gather the elements with indices specified by the array `indx` from a full-storage vector `y` into compressed form (`nz`, `x`, `indx`) and overwrite the gathered elements of `y` by zeros. Other elements of `y` are not referenced or modified (see also `?gthr`).

Input Parameters

`nz` **INTEGER**. The number of elements of `y` to be gathered.

`indx` **INTEGER**. Specifies indices of elements to be gathered. Array, **DIMENSION** at least `nz`.

`y` **REAL** for `sgthrz`
DOUBLE PRECISION for `dgthrz`
COMPLEX for `cgthrz`
DOUBLE COMPLEX for `zgthrz`
Array, **DIMENSION** at least $\max_i (\text{indx}(i))$.

Output Parameters

`x` **REAL** for `sgthrz`
DOUBLE PRECISION for `dgthrz`
COMPLEX for `cgthrz`
DOUBLE COMPLEX for `zgthrz`
Array, **DIMENSION** at least `nz`.
Contains the vector converted to the compressed form.

`y` The updated vector `y`.

?roti

Applies Givens rotation to sparse vectors
one of which is in compressed form.

```
call sroti ( nz, x, indx, y, c, s )
call droti ( nz, x, indx, y, c, s )
```

Discussion

The ?roti routines apply the Givens rotation to elements of two real vectors, x (in compressed form $nz, x, indx$) and y (in full storage form):

$$x(i) = c*x(i) + s*y(indx(i))$$

$$y(indx(i)) = c*y(indx(i)) - s*x(i)$$

The routines reference only the elements of y whose indices are listed in the array $indx$. The values in $indx$ must be distinct.

Input Parameters

nz	INTEGER. The number of elements in x and $indx$.
x	REAL for sroti DOUBLE PRECISION for droti Array, DIMENSION at least nz .
$indx$	INTEGER. Specifies the indices for the elements of x . Array, DIMENSION at least nz .
y	REAL for sroti DOUBLE PRECISION for droti Array, DIMENSION at least $\max_i(indx(i))$.
c	A scalar: REAL for sroti DOUBLE PRECISION for droti.
s	A scalar: REAL for sroti DOUBLE PRECISION for droti.

Output Parameters

x and y The updated arrays.

?sctr

Converts compressed sparse vectors into full storage form.

```
call ssctr ( nz, x, indx, y )
call dsctr ( nz, x, indx, y )
call csctr ( nz, x, indx, y )
call zsctr ( nz, x, indx, y )
```

Discussion

The `?sctr` routines scatter the elements of the compressed sparse vector $(nz, x, indx)$ to a full-storage vector y . The routines modify only the elements of y whose indices are listed in the array $indx$:
 $y(indx(i)) = x(i)$, for $i=1,2,\dots,nz$.

Input Parameters

nz **INTEGER**. The number of elements of x to be scattered.

$indx$ **INTEGER**. Specifies indices of elements to be scattered.
Array, **DIMENSION** at least nz .

x **REAL** for `ssctr`
DOUBLE PRECISION for `dsctr`
COMPLEX for `csctr`
DOUBLE COMPLEX for `zsctr`
Array, **DIMENSION** at least nz .
Contains the vector to be converted to full-storage form.

Output Parameters

y **REAL** for `ssctr`
DOUBLE PRECISION for `dsctr`
COMPLEX for `csctr`
DOUBLE COMPLEX for `zsctr`
Array, **DIMENSION** at least $\max_i (indx(i))$.
Contains the vector y with updated elements.

Fast Fourier Transforms

3

This chapter describes the fast Fourier transform (FFT) routines implemented in Intel[®] MKL. The FFT routines included consist of two classes: one-dimensional and two-dimensional. Both one-dimensional and two-dimensional routines have been optimized to effectively use cache memory. All routines work with transforms of a power of 2 length.

For a more general set of Discrete Fourier Transform functions in Intel MKL, refer to [Advanced DFT Functions](#) in this manual.

Although Intel MKL still supports the FFT interface described later in this chapter, users are encouraged to migrate to the new DFT functions in their application programs. Unlike the FFT routines, the DFT routines support transforms of up to 7D, and transform lengths of other than powers of 2 mixed radix.

This chapter contains these major sections:

- One-dimensional FFTs
- Two-dimensional FFTs

Each of the major sections contains the description of three groups of the FFTs.

One-dimensional FFTs

The one-dimensional FFTs include the following groups:

- Complex-to-Complex Transforms
- Real-to-Complex Transforms
- Complex-to-Real Transforms.

All one-dimensional FFTs are in-place. The transform length must be a power of 2. The complex-to-complex transform routines perform both forward and inverse transforms of a complex vector. The real-to-complex transform routines perform forward transforms of a real vector. The complex-to-real transform routines perform inverse transforms of a complex conjugate-symmetric vector, which is packed in a real array.

Data Storage Types

Each FFT group contains two sets of FFTs having the similar functionality: one set is used for the Fortran-interface and the other for the C-interface. The former set stores the complex data as a Fortran complex data type, while the latter stores the complex data as float arrays of real and imaginary parts separately. These sets are distinguished by naming the FFTs within each set. The names of the FFTs used for the C-interface have the letter “c” added to the end of the FFTs’ Fortran names. For example, the names of the `cfft1d/zfft1d` FFTs for the corresponding C-interface routines are `cfft1dc/zfft1dc`. All names of the C-type data items are lower case.

[Table 3-1](#) lists the one-dimensional FFT routine groups and the data types associated with them.

Table 3-1 One-dimensional FFTs: Names and Data Types

Group	Stored as Fortran Complex Data	Stored as C Real Data	Data Types	Description
Complex-to-Complex	cfft1d/zfft1d	cfft1dc/zfft1dc	c, z	Transform complex data to complex data.
Real-to-Complex	scfft1d/dzfft1d	scfft1dc/dzfft1dc	sc, dz	Transform forward real-to-complex data. Complement csfft1d/zdfft1d and csfft1dc/zdfft1dc FFTs.
Complex-to-Real	csfft1d/zdfft1d	csfft1dc/zdfft1dc	cs, zd	Transform inverse complex-to-real data. Complement scfft1d/dzfft1d and scfft1dc/dzfft1dc FFTs.

Data Structure Requirements

For C-interface, storage of the complex-to-complex transform routines data requires separate float arrays for the real and imaginary parts. The real-to-complex and complex-to-real pairs require a single float input/output array.

The C-interface requires scalar values to be passed by value.

All transforms require additional memory to store the transform coefficients. When performing multiple FFTs of the same dimension, the table of coefficients should be created only once and then used on all the FFTs afterwards. Using the same table rather than creating it repeatedly for each FFT produces an obvious performance gain.

Complex-to-Complex One-dimensional FFTs

Each of the complex-to-complex routines computes a forward or inverse FFT of a complex vector.

The forward FFT is computed according to the mathematical equation

$$z_j = \sum_{k=0}^{n-1} r_k * w^{-j*k}, \quad 0 \leq j \leq n-1$$

The inverse FFT is computed according to the mathematical equation

$$r_j = \frac{1}{n} \sum_{k=0}^{n-1} z_k * w^{j*k}, \quad 0 \leq j \leq n-1$$

where $w = \exp\left[\frac{2\pi i}{n}\right]$, i being the imaginary unit.

The operation performed by the complex-to-complex routines is determined by the value of the *isign* parameter used by each of these routines.

If *isign* = -1, perform the forward FFT where input and output are in normal order.

If *isign* = +1, perform the inverse FFT where input and output are in normal order.

If *isign* = -2, perform the forward FFT where input is in normal order and output is in bit-reversed order.

If *isign* = +2, perform the inverse FFT where input is in bit-reversed order

and output is in normal order.

If `isign = 0`, initialize FFT coefficients for both the forward and inverse FFTs.

The above equations apply to all FFTs with all data types indicated in [Table 3-1](#).

To compute a forward or inverse FFT of a given length, first initialize the coefficients by calling the function with `isign = 0`. Thereafter, any number of transforms of the same length can be computed by calling the function with `isign = +1, -1, +2, -2`.

cfft1d/zfft1d

Fortran-interface routines. Compute the forward or inverse FFT of a complex vector (in-place)

```
call cfft1d ( r, n, isign, wsave )
call zfft1d ( r, n, isign, wsave )
```

Discussion

The operation performed by the `cfft1d/zfft1d` routines is determined by the value of `isign`. See the equations of the operations for the [Complex-to-Complex One-dimensional FFTs](#) above.

Input Parameters

<code>r</code>	<code>COMPLEX</code> for <code>cfft1d</code> <code>DOUBLE COMPLEX</code> for <code>zfft1d</code> Array, <code>DIMENSION</code> at least (<code>n</code>). Contains the complex vector on which the transform is to be performed. Not referenced if <code>isign = 0</code> .
<code>n</code>	<code>INTEGER</code> . Transform length; <code>n</code> must be a power of 2.
<code>isign</code>	<code>INTEGER</code> . Flag indicating the type of operation to be performed: if <code>isign = 0</code> , initialize the coefficients <code>wsave</code> ;

if `isign = -1`, perform the forward FFT where input and output are in normal order;
 if `isign = +1`, perform the inverse FFT where input and output are in normal order;
 if `isign = -2`, perform the forward FFT where input is in normal order and output is in bit-reversed order;
 if `isign = +2`, perform the inverse FFT where input is in bit-reversed order and output is in normal order.

`wsave` `COMPLEX` for `cfft1d`
 `DOUBLE COMPLEX` for `zfft1d`
 Array, `DIMENSION` at least $((3*n)/2)$. If `isign = 0`, then `wsave` is an output parameter. Otherwise, `wsave` contains the FFT coefficients initialized on a previous call with `isign = 0`.

Output Parameters

`r` Contains the complex result of the transform depending on `isign`. Does not change if `isign = 0`.

`wsave` If `isign = 0`, `wsave` contains the initialized FFT coefficients. Otherwise, `wsave` does not change.

cfft1dc/zfft1dc

C-interface routines. Compute the forward or inverse FFT of a complex vector (in-place).

```
void cfft1dc (float* r, float* i, int n, int isign, float* wsave)
void zfft1dc (double* r, double* i, int n, int isign, double* wsave)
```

Discussion

The operation performed by the `cfft1dc/zfft1dc` routines is determined by the value of `isign`. See the equations of the operations for the [Complex-to-Complex One-dimensional FFTs](#).

Input Parameters

<i>r</i>	<p><code>float*</code> for <code>cffft1dc</code> <code>double*</code> for <code>zffft1dc</code></p> <p>Pointer to an array of size at least (<i>n</i>). Contains the real parts of complex vector to be transformed. Not referenced if <i>isign</i> = 0.</p>
<i>i</i>	<p><code>float*</code> for <code>cffft1dc</code> <code>double*</code> for <code>zffft1dc</code></p> <p>Pointer to an array of size at least (<i>n</i>). Contains the imaginary parts of complex vector to be transformed. Not referenced if <i>isign</i> = 0.</p>
<i>n</i>	<code>int</code> . Transform length; <i>n</i> must be a power of 2.
<i>isign</i>	<p><code>int</code>. Flag indicating the type of operation to be performed:</p> <ul style="list-style-type: none"> if <i>isign</i> = 0, initialize the coefficients <i>wsave</i>; if <i>isign</i> = -1, perform the forward FFT where input and output are in normal order; if <i>isign</i> = +1, perform the inverse FFT where input and output are in normal order; if <i>isign</i> = -2, perform the forward FFT where input is in normal order and output is in bit-reversed order; if <i>isign</i> = +2, perform the inverse FFT where input is in bit-reversed order and output is in normal order.
<i>wsave</i>	<p><code>float*</code> for <code>cffft1dc</code> <code>double*</code> for <code>zffft1dc</code></p> <p>Pointer to an array of size at least ($3*n$). If <i>isign</i> = 0, then <i>wsave</i> is an output parameter. Otherwise, <i>wsave</i> contains the FFT coefficients initialized on a previous call with <i>isign</i> = 0.</p>

Output Parameters

<i>r</i>	Contains the real part of the transform depending on <i>isign</i> . Does not change if <i>isign</i> = 0.
----------	--

<code>i</code>	Contains the imaginary part of the transform depending on <code>isign</code> .. Does not change if <code>isign = 0</code> .
<code>wsave</code>	If <code>isign = 0</code> , <code>wsave</code> contains the initialized FFT coefficients. Otherwise, <code>wsave</code> does not change.

Real-to-Complex One-dimensional FFTs

Each of the real-to-complex routines computes forward FFT of a real input vector according to the mathematical equation

$$z_j = \sum_{k=0}^{n-1} t_k * w^{-j*k}, \quad 0 \leq j \leq n-1$$

for $t_k = \text{cplx}(r_k, 0)$, where r_k is the real input vector, $0 \leq k \leq n-1$. The mathematical result z_j , $0 \leq j \leq n-1$, is the complex conjugate-symmetric vector, where $z(n/2+i) = \text{conjg}(z(n/2-i))$, $1 \leq i \leq n/2-1$, and moreover $z(0)$ and $z(n/2)$ are real values.

This complex conjugate-symmetric (CCS) vector can be stored in the complex array of size $(n/2+1)$ or in the real array of size $(n+2)$. The data storage of the CCS format is defined later for Fortran-interface and C-interface routines separately.

[Table 3-2](#) shows a comparison of the effects of performing the `cffft1d/zffft1d` complex-to-complex FFT on a vector of length $n=8$ in which all the imaginary elements are zeros, with the real-to-complex `scffft1d/zdffft1d` FFT applied to the same vector. The advantage of the latter approach is that only half of the input data storage is required and there is no need to zero the imaginary part. The last two columns are stored in the real array of size $(n+2)$ containing the complex conjugate-symmetric vector in CCS format.

To compute a forward FFT of a given length, first initialize the coefficients by calling the routine you are going to use with *isign* = 0. Thereafter, any number of real-to-complex and complex-to-real transforms of the same length can be computed by calling that routine with the *isign* value other than 0.

Table 3-2 Comparison of the Storage Effects of Complex-to-Complex and Real-to-Complex FFTs

Input Vectors			Output Vectors			
cfft1d		scfft1d	cfft1d		scfft1d	
Complex Data		Real Data	Complex Data		Real Data	
Real	Imaginary		Real	Imaginary	(Real)	(Imaginary)
0.841471	0.000000	0.841471	1.543091	0.000000	1.543091	0.000000
0.909297	0.000000	0.909297	3.875664	0.910042	3.875664	0.910042
0.141120	0.000000	0.141120	-0.915560	-0.397326	-0.915560	-0.397326
-0.756802	0.000000	-0.756802	-0.274874	-0.121691	-0.274874	-0.121691
-0.958924	0.000000	-0.958924	-0.181784	0.000000	-0.181784	0.000000
-0.279415	0.000000	-0.279415	-0.274874	0.121691		
0.656987	0.000000	0.656987	-0.915560	0.397326		
0.989358	0.000000	0.989358	3.875664	-0.910042		

scfft1d/dzfft1d

Fortran-interface routines. Compute forward FFT of a real vector and represent the complex conjugate-symmetric result in CCS format (in-place).

```
call scfft1d ( r, n, isign, wsave )
call dzfft1d ( r, n, isign, wsave )
```

Discussion

The operation performed by the `scfft1d/dzfft1d` routines is determined by the value of `isign`. See the equations of the operations for [Real-to-Complex One-dimensional FFTs](#) above. These routines are complementary to the complex-to-real transform routines [csfft1d/zdfft1d](#).

Input Parameters

<code>r</code>	REAL for <code>scfft1d</code> DOUBLE PRECISION for <code>dzfft1d</code> Array, DIMENSION at least $(n+2)$. First n elements contain the input vector to be transformed. The elements $r(n+1)$ and $r(n+2)$ are used on output. The array <code>r</code> is not referenced if <code>isign = 0</code> .
<code>n</code>	INTEGER. Transform length; n must be a power of 2.
<code>isign</code>	INTEGER. Flag indicating the type of operation to be performed: if <code>isign</code> is 0, initialize the coefficients <code>wsave</code> ; if <code>isign</code> is not 0, perform the forward FFT.
<code>wsave</code>	REAL for <code>scfft1d</code> DOUBLE PRECISION for <code>dzfft1d</code> Array, DIMENSION at least $(2*n+4)$. If <code>isign = 0</code> , then <code>wsave</code> contains output data. Otherwise, <code>wsave</code> contains coefficients required to perform the FFT that has been initialized on a previous call to this routine or the complementary complex-to-real FFT routine.

Output Parameters

<code>r</code>	If <code>isign = 0</code> , <code>r</code> does not change. If <code>isign</code> is not 0, the output real-valued array <code>r(1:n+2)</code> contains the complex conjugate-symmetric vector <code>z(1:n)</code> packed in CCS format for Fortran interface. The table below shows the relationship between them.
----------------	--

$r(1)$	$r(2)$	$r(3)$	$r(4)$...	$r(n-1)$	$r(n)$	$r(n+1)$	$r(n+2)$
$z(1)$	0	REz(2)	IMz(2)	...	REz(n/2)	IMz(n/2)	$z(n/2+1)$	0

The full complex vector $z(1:n)$ is defined by

$$z(i) = \text{cplx}(r(2*i-1), r(2*i)),$$

$$1 \leq i \leq n/2+1,$$

$$z(n/2+i) = \text{conjg}(z(n/2+2-i)),$$

$$2 \leq i \leq n/2.$$

Then, $z(1:n)$ is the forward FFT of a real input vector $r(1:n)$.

wsave

If $isign = 0$, *wsave* contains the coefficients required by the called routine. Otherwise *wsave* does not change.

scfft1dc/dzfft1dc

C-interface routines. Compute forward FFT of a real vector and represent the complex conjugate-symmetric result in CCS format (in-place).

```
void scfft1dc ( float* r, int n, int isign, float* wsave );
void dzfft1dc ( double* r, int n, int isign, double* wsave );
```

Discussion

The operation performed by the `scfft1dc/dzfft1dc` routines is determined by the value of *isign*. See the equations of the operations for the [Real-to-Complex One-dimensional FFTs](#) above.

These routines are complementary to the complex-to-real transform routines [csfft1dc/zdfft1dc](#).

Input Parameters

<i>r</i>	float* for <code>scfft1dc</code> double* for <code>dzfft1dc</code>
	Pointer to an array of size at least $(n+2)$. First n elements contain the input vector to be transformed. The array <i>r</i> is not referenced if <i>isign</i> = 0.
<i>n</i>	int. Transform length; <i>n</i> must be a power of 2.
<i>isign</i>	int. Flag indicating the type of operation to be performed: if <i>isign</i> is 0, initialize the coefficients <i>wsave</i> ; if <i>isign</i> is not 0, perform the forward FFT.
<i>wsave</i>	float* for <code>scfft1dc</code> double* for <code>dzfft1dc</code>
	Pointer to an array of size at least $(2*n+4)$. If <i>isign</i> = 0, then <i>wsave</i> contains output data. Otherwise, <i>wsave</i> contains coefficients required to perform the FFT that has been initialized on a previous call to this routine or the complementary complex-to-real FFT routine.

Output Parameters

<i>r</i>	If <i>isign</i> = 0, <i>r</i> does not change. If <i>isign</i> is not 0, the output real-valued array $r(0:n+1)$ contains the complex conjugate-symmetric vector $z(0:n-1)$ packed in CCS format for C-interface. The table below shows the relationship between them.
----------	---

$r(0)$	$r(1)$	$r(2)$...	$r(n/2)$	$r(n/2+1)$	$r(n/2+2)$...	$r(n)$	$r(n+1)$
$z(0)$	REz(1)	REz(2)	...	$z(n/2)$	0	IMz(1)	...	IMz($n/2-1$)	0

The full complex vector $z(0:n-1)$ is defined by
 $z(i) = \text{cmplx}(r(i), r(n/2+1+i))$, $0 \leq i \leq n/2$,

$z(n/2+i) = \text{conjg}(z(n/2-i)), 1 \leq i \leq n/2-1.$

Then, $z(0:n-1)$ is the forward FFT of the real input vector of length n .

wsave If *isign* = 0, *wsave* contains the coefficients required by the called routine. Otherwise *wsave* does not change.

Complex-to-Real One-dimensional FFTs

Each of the complex-to-real routines computes a one-dimensional inverse FFT according to the mathematical equation

$$t_j = \frac{1}{n} \sum_{k=0}^{n-1} z_k * w^{j*k}, \quad 0 \leq j \leq n-1$$

The mathematical input is the complex conjugate-symmetric vector z_j , $0 \leq j \leq n-1$, , where $z(n/2+i) = \text{conjg}(z(n/2-i)), 1 \leq i \leq n/2-1$, and moreover $z(0)$ and $z(n/2)$ are real values.

The mathematical result is $t_j = \text{cplx}(r_j, 0)$, where r_j is a real vector, $0 \leq j \leq n-1$.

Input to the complex-to-real transform routines is a real array of size $(n+2)$, which contains the complex conjugate-symmetric vector $z(0:n-1)$ in CCS format (see [Real-to-Complex One-dimensional FFTs](#) above).

Output of the complex-to-real routines is a real vector of size n .

[Table 3-3](#) is identical to [Table 3-2](#), except for reversing the input and output vectors. In the complex-to-real routines the last two columns are stored in the input real array of size $(n+2)$ containing the complex conjugate-symmetric vector in CCS format.

To compute an inverse FFT of a given length, first initialize the coefficients by calling the routine you are going to use with *isign* = 0. Thereafter, any number of real-to-complex and complex-to-real transforms of the same length can be computed by calling the appropriate routine with the *isign* value other than 0.

Table 3-3 Comparison of the Storage Effects of Complex-to-Real and Complex-to-Complex FFTs

Output Vectors			Input Vectors			
cfft1d		csfft1d	cfft1d		csfft1d	
Complex Data		Real Data	Complex Data		Real Data	
Real	Imaginary		Real	Imaginary	(Real)	(Imaginary)
0.841471	0.000000	0.841471	1.543091	0.000000	1.543091	0.000000
0.909297	0.000000	0.909297	3.875664	0.910042	3.875664	0.910042
0.141120	0.000000	0.141120	-0.915560	-0.397326	-0.915560	-0.397326
-0.756802	0.000000	-0.756802	-0.274874	-0.121691	-0.274874	-0.121691
-0.958924	0.000000	-0.958924	-0.181784	0.000000	-0.181784	0.000000
-0.279415	0.000000	-0.279415	-0.274874	0.121691		
0.656987	0.000000	0.656987	-0.915560	0.397326		
0.989358	0.000000	0.989358	3.875664	-0.910042		

csfft1d/zdfft1d

Fortran-interface routines.

Compute inverse FFT of a complex conjugate-symmetric vector packed in CCS format (in-place).

```
call csfft1d ( r, n, isign, wsave )
call zdfft1d ( r, n, isign, wsave )
```

Discussion

The operation performed by the `csfft1d/zdfft1d` routines is determined by the value of `isign`. See the equations of the operations for the [Complex-to-Real One-dimensional FFTs](#) above.

These routines are complementary to the real-to-complex transform routines [scfft1d/dzfft1d](#).

Input Parameters

r REAL for `csfft1d`
 DOUBLE PRECISION for `zdfft1d`
 Array, DIMENSION at least $(n+2)$.
 Not referenced if $isign = 0$.
 If $isign$ is not 0, then $r(1:n+2)$ contains the complex conjugate-symmetric vector packed in CCS format for Fortran-interface.
 The table below shows the relationship between them

$r(1)$	$r(2)$	$r(3)$	$r(4)$...	$r(n-1)$	$r(n)$	$r(n+1)$	$r(n+2)$
$z(1)$	0	RE $z(2)$	IM $z(2)$...	RE $z(n/2)$	IM $z(n/2)$	$z(n/2+1)$	0

The full complex vector $z(1:n)$ is defined by

$$z(i) = \text{cplx}(r(2*i-1), r(2*i)),$$

$$1 \leq i \leq n/2+1,$$

$$z(n/2+i) = \text{conjg}(z(n/2+2-i)),$$

$$2 \leq i \leq n/2.$$

After the transform, $r(1:n)$ contains the inverse FFT of the complex conjugate-symmetric vector $z(1:n)$.

n INTEGER. Transform length; n must be a power of 2.

isign INTEGER. Flag indicating the type of operation to be performed:
 if $isign$ is 0, initialize the coefficients *wsave*;
 if $isign$ is not 0, perform the inverse FFT.

wsave REAL for `csfft1d`
 DOUBLE PRECISION for `zdfft1d`
 Array, DIMENSION at least $(2*n+4)$. If $isign = 0$, then *wsave* contains output data. Otherwise, *wsave*

contains coefficients required to perform the FFT that has been initialized on a previous call to this routine or the complementary real-to-complex FFT routine.

Output Parameters

r If *isign* is not 0, then *r(1:n)* is the real result of the inverse FFT of the complex conjugate-symmetric vector *z(1:n)*. Does not change if *isign* = 0.

wsave If *isign* = 0, *wsave* contains the coefficients required by the called routine. Otherwise *wsave* does not change.

csfft1dc/zdffft1dc

C-interface routines. Compute inverse FFT of a complex conjugate-symmetric vector packed in CCS format (in-place).

```
void csfft1dc ( float* r, int n, int isign, float* wsave )
void zdffft1dc ( double* r, int n, int isign, double* wsave )
```

Discussion

The operation performed by the `csfft1dc/zdffft1dc` routines is determined by the value of *isign*. See the equations of the operations for the [Complex-to-Real One-dimensional FFTs](#) above.

These routines are complementary to the real-to-complex transform routines [scfft1dc/dzffft1dc](#).

Input Parameters

r float* for `csfft1dc`
double* for `zdffft1dc`

Pointer to an array of size at least $(n+2)$. Not referenced if $isign = 0$.

If $isign$ is not 0, then $r(0:n+1)$ contains the complex conjugate-symmetric vector packed in CCS format for C-interface.

The table below shows the relationship between them.

$r(0)$	$r(1)$	$r(2)$...	$r(n/2)$	$r(n/2+1)$	$r(n/2+2)$...	$r(n)$	$r(n+1)$
$z(0)$	REz(1)	REz(2)	...	$z(n/2)$	0	IMz(1)	...	IMz(n/2-1)	0

The full complex vector $z(0:n-1)$ is defined by

$$z(i) = \text{cplx}(r(i), r(n/2+1+i)), \quad 0 \leq i \leq n/2,$$

$$z(n/2+i) = \text{conjg}(z(n/2-i)), \quad 1 \leq i \leq n/2-1.$$

After the transform, $r(0:n-1)$ is the inverse FFT of the complex conjugate-symmetric vector $z(0:n-1)$.

- n **int.** Transform length; n must be a power of 2.
- $isign$ **int.** Flag indicating the type of operation to be performed:
 if $isign = 0$, initialize the coefficients $wsave$;
 if $isign$ is not 0, perform the inverse FFT.
- $wsave$ **float*** for `csfft1dc`
double* for `zdf1dc`
 Pointer to an array of size at least $(2*n+4)$.
 If $isign = 0$, then $wsave$ contains output data.
 Otherwise, $wsave$ contains coefficients required to perform the FFT that has been initialized on a previous call to this routine or the complementary real-to-complex FFT routine.

Output Parameters

- r If $isign$ is not 0, then $r(0:n-1)$ is the real result of the inverse FFT of the complex conjugate-symmetric vector $z(0:n-1)$. Does not change if $isign = 0$.
- $wsave$ If $isign = 0$, $wsave$ contains the coefficients required by the called routine. Otherwise $wsave$ does not change.

Two-dimensional FFTs

The two-dimensional FFTs are functionally the same as one-dimensional FFTs. They contain the following groups:

- Complex-to-Complex Transforms
- Real-to-Complex Transforms
- Complex-to-Real Transforms.

All two-dimensional FFTs are in-place. Transform lengths must be a power of 2. The complex-to-complex transform routines perform both forward and inverse transforms of a complex matrix. The real-to-complex transform routines perform forward transforms of a real matrix. The complex-to-real transform routines perform inverse transforms of a complex conjugate-symmetric matrix, which is packed in a real array.

The naming conventions are also the same as those for one-dimensional FFTs, with “2d” replacing “1d” in all cases. [Table 3-4](#) lists the two-dimensional FFT routine groups and the data types associated with them.

Table 3-4 Two-dimensional FFTs: Names and Data Types

Group	Stored as FORTRAN Complex Data	Stored as C Real Data	Data Types	Description
Complex-to-Complex	cfft2d/ zfft2d	cfft2dc/ zfft2dc	c, z	Transform complex data to complex data.
Real-to-Complex	scfft2d/ dzfft2d	scfft2dc/ dzfft2dc	sc, dz	Transform forward real-to-complex data. Complement csfft2d/zdfft2d and csfft2dc/zdfft2dc FFTs.
Complex-to-Real	csfft2d/ zdfft2d	csfft2dc/ zdfft2dc	cs, zd	Transform inverse complex-to-real data. Complement scfft2d/dzfft2d and scfft2dc/dzfft2dc FFTs.

The C-interface requires scalar values to be passed by value. The major difference between the one-dimensional and two-dimensional FFTs is that your application does not need to provide storage for transform coefficients.

The data storage types and data structure requirements are the same as for one-dimensional FFTs. For more information, see the [Data Storage Types](#) and [Data Structure Requirements](#) sections at the beginning of this chapter.

Complex-to-Complex Two-dimensional FFTs

Each of the complex-to-complex routines computes a forward or inverse FFT of a complex matrix in-place.

The forward FFT is computed according to the mathematical equation

$$z_{i,j} = \sum_{k=0}^{m-1} \sum_{l=0}^{n-1} r_{k,l} * w_m^{-i*k} * w_n^{-j*l}, \quad 0 \leq i \leq m-1, \quad 0 \leq j \leq n-1$$

The inverse FFT is computed according to the mathematical equation

$$r_{i,j} = \frac{1}{m*n} \sum_{k=0}^{m-1} \sum_{l=0}^{n-1} z_{k,l} * w_m^{i*k} * w_n^{j*l}, \quad 0 \leq i \leq m-1, \quad 0 \leq j \leq n-1$$

where $w_m = \exp\left[\frac{2\pi i}{m}\right]$, $w_n = \exp\left[\frac{2\pi i}{n}\right]$, i being the imaginary unit.

The operation performed by the complex-to-complex routines is determined by the value of the *isign* parameter.

If *isign* = -1, perform the forward FFT where input and output are in normal order.

If *isign* = +1, perform the inverse FFT where input and output are in normal order.

If *isign* = -2, perform the forward FFT where input is in normal order and output is in bit-reversed order.

If *isign* = +2, perform the inverse FFT where input is in bit-reversed order and output is in normal order.

The above equations apply to all FFTs with all data types indicated in [Table 3-4](#).

cfft2d/zfft2d

Fortran-interface routines. Compute the forward or inverse FFT of a complex matrix (in-place).

```
call cfft2d ( r, m, n, isign )
call zfft2d ( r, m, n, isign )
```

Discussion

The operation performed by the `cfft2d/zfft2d` routines is determined by the value of `isign`. See the equations of the operations for [Complex-to-Complex Two-dimensional FFTs](#).

Input Parameters

<code>r</code>	<code>COMPLEX</code> for <code>cfft2d</code> <code>DOUBLE COMPLEX</code> for <code>zfft2d</code> Array, <code>DIMENSION</code> at least (m,n) , with its leading dimension equal to m . This array contains the complex matrix to be transformed.
<code>m</code>	<code>INTEGER</code> . Column transform length (number of rows); m must be a power of 2.
<code>n</code>	<code>INTEGER</code> . Row transform length (number of columns); n must be a power of 2.
<code>isign</code>	<code>INTEGER</code> . Flag indicating the type of operation to be performed: if <code>isign = -1</code> , perform the forward FFT where input and output are in normal order; if <code>isign = +1</code> , perform the inverse FFT where input and output are in normal order; if <code>isign = -2</code> , perform the forward FFT where input is in normal order and output is in bit-reversed order; if <code>isign = +2</code> , perform the inverse FFT where input is in bit-reversed order and output is in normal order.

Output Parameters

r Contains the complex result of the transform depending on *isign*.

cfft2dc/zfft2dc

C-interface routines. Compute the forward or inverse FFT of a complex matrix (in-place).

```
void cfft2dc ( float* r, float* i, int m, int n, int isign )
void zfft2dc ( double* r, double* i, int m, int n, int isign )
```

Discussion

The operation performed by the `cfft2dc/zfft2dc` routines is determined by the value of *isign*. See the equations of the operations for the [Complex-to-Complex Two-dimensional FFTs](#) above.

Input Parameters

r `float*` for `cfft2dc`
`double*` for `zfft2dc`

Pointer to a two-dimensional array of size at least (m, n) , with its leading dimension equal to *n*. The array contains the real parts of a complex matrix to be transformed.

i `float*` for `cfft2dc`
`double*` for `zfft2dc`

Pointer to a two-dimensional array of size at least (m, n) , with its leading dimension equal to *n*. The array contains the imaginary parts of a complex matrix to be transformed.

m `int`. Column transform length (number of rows); *m* must be a power of 2.

n **int.** Row transform length (number of columns); *n* must be a power of 2.

isign **int.** Flag indicating the type of operation to be performed:
 if *isign* = -1, perform the forward FFT where input and output are in normal order;
 if *isign* = +1, perform the inverse FFT where input and output are in normal order;
 if *isign* = -2, perform the forward FFT where input is in normal order and output is in bit-reversed order;
 if *isign* = +2, perform the inverse FFT where input is in bit-reversed order and output is in normal order.

Output Parameters

r Contains the real parts of the complex result depending on *isign*.

i Contains the imaginary parts of the complex depending on *isign*.

Real-to-Complex Two-dimensional FFTs

Each of the real-to-complex routines computes the forward FFT of a real matrix according to the mathematical equation

$$z_{i,j} = \sum_{k=0}^{m-1} \sum_{l=0}^{n-1} t_{k,l} * w_m^{-i*k} * w_n^{-j*l}, \quad 0 \leq i \leq m-1, \quad 0 \leq j \leq n-1$$

$t_{k,l} = \text{cplx}(r_{k,l}, 0)$, where $r_{k,l}$ is a real input matrix, $0 \leq k \leq m-1$, $0 \leq l \leq n-1$.
 The mathematical result $z_{i,j}$, $0 \leq i \leq m-1$, $0 \leq j \leq n-1$, is the complex matrix of size (m, n) .
 Each column is the complex conjugate-symmetric vector as follows:

for $0 \leq j \leq n-1$,

$$z(m/2+i, j) = \text{conjg}(z(m/2-i, j)), \quad 1 \leq i \leq m/2-1.$$

Moreover, $z(0, j)$ and $z(m/2, j)$ are real values for $j=0$ and $j=n/2$.

This mathematical result can be stored in the complex two-dimensional array of size $(m/2+1, n/2+1)$ or in the real two-dimensional array of size $(m+2, n+2)$. The data storage of CCS format is defined later for Fortran-interface and C-interface routines separately.

scfft2d/dzfft2d

Fortran-interface routines. Compute forward FFT of a real matrix and represent the complex conjugate-symmetric result in CCS format (in-place).

```
call scfft2d ( r, m, n )
call dzfft2d ( r, m, n )
```

Discussion

See the equations of the operations for the [Real-to-Complex Two-dimensional FFTs](#) above.

These routines are complementary to the complex-to-real transform routines [csfft2d/zdfft2d](#).

Input Parameters

<i>r</i>	REAL for <code>scfft2d</code> DOUBLE PRECISION for <code>dzfft2d</code> Array, DIMENSION at least $(m+2, n+2)$, with its leading dimension equal to $(m+2)$. The first m rows and n columns of this array contain the real matrix to be transformed. Table 3-5 presents the input data layout.
<i>m</i>	INTEGER . Column transform length (number of rows); m must be a power of 2.
<i>n</i>	INTEGER . Row transform length (number of columns); n must be a power of 2.

Table 3-5 Fortran-interface Real Data Storage for the Real-to-Complex and Complex-to-Real Two-dimensional FFTs

$r(1, 1)$	$r(1, 2)$...	$r(1, n-1)$	$r(1, n)$	n/u	n/u
$r(2, 1)$	$r(2, 2)$...	$r(2, n-1)$	$r(2, n)$	n/u	n/u
$r(3, 1)$	$r(3, 2)$...	$r(3, n-1)$	$r(3, n)$	n/u	n/u
$r(4, 1)$	$r(4, 2)$...	$r(4, n-1)$	$r(4, n)$	n/u	n/u
...
$r(m-1, 1)$	$r(m-1, 2)$...	$r(m-1, n-1)$	$r(m-1, n)$	n/u	n/u
$r(m, 1)$	$r(m, 2)$...	$r(m, n-1)$	$r(m, n)$	n/u	n/u
n/u	n/u	...	n/u	n/u	n/u	n/u
n/u	n/u	...	n/u	n/u	n/u	n/u

* n/u - not used

Output Parameters

- r The output real array $r(1:m+2, 1:n+2)$ contains the complex conjugate-symmetric matrix $z(1:m, 1:n)$ packed in CCS format for Fortran-interface as follows:
- Rows 1 and $m+1$ contain in $n+2$ locations the complex conjugate-symmetric vectors $z(1, j)$ and $z(m/2+1, j)$ packed in CCS format (see [Real-to-Complex One-dimensional FFTs](#) above).
The full complex vector $z(1, j)$ is defined by:

$$z(1, j) = \text{cplx}(r(1, 2*j-1), r(1, 2*j)), \quad 1 \leq j \leq n/2+1,$$

$$z(1, n/2+1+j) = \text{conjg}(z(1, n/2+1-j)), \quad 1 \leq j \leq n/2-1.$$
The full complex vector $z(m/2+1, j)$ is defined by:

$$z(m/2+1, j) = \text{cplx}(r(m+1, 2*j-1), r(m+1, 2*j)),$$

$$1 \leq j \leq n/2+1,$$

$$z(m/2+1, n/2+1+j) = \text{conjg}(z(m/2+1, n/2+1-j)),$$

$$1 \leq j \leq n/2-1;$$
 - Rows from 3 to m contain in n locations complex vectors represented as

$$z(i+1, j) = \text{cplx}(r(2*i+1, j), r(2*i+2, j)),$$

$$1 \leq i \leq m/2-1, \quad 1 \leq j \leq n.$$

- The rest matrix elements can be obtained from

$$z(m/2+1+i, j) = \text{conjg}(z(m/2+1-i, j)),$$

$$1 \leq i \leq m/2-1, 1 \leq j \leq n.$$

The storage of the complex conjugate-symmetric matrix z for Fortran-interface is shown in [Table 3-6](#).

Table 3-6 Fortran-interface Data Storage of CCS Format for the Real-to-Complex and Complex-to-Real Two-Dimensional FFTs

$z(1,1)$	0	REz(1,2)	IMz(1,2)	...	REz(1,n/2)	IMz(1,n/2)	$z(1, n/2+1)$	0
0	0	0	0	...	0	0	0	0
REz(2,1)	REz(2,2)	REz(2,3)	REz(2,4)	...	REz(2,n-1)	REz(2,n)	n/u	n/u
IMz(2,1)	IMz(2,2)	IMz(2,3)	IMz(2,4)	...	IMz(2,n-1)	IMz(2,n)	n/u	n/u
...	n/u	n/u
REz(m/2,1)	REz(m/2,2)	REz(m/2,3)	REz(m/2,4)	...	REz(m/2, n-1)	REz(m/2, n)	n/u	n/u
IMz(m/2,1)	IMz(m/2,2)	IMz(m/2,3)	IMz(m/2,4)	...	IMz(m/2, n-1)	IMz(m/2, n)	n/u	n/u
$z(m/2+1,1)$	0	REz(m/2+1,2)	IMz(m/2+1,2)	...	REz(m/2+1, n/2)	IMz(m/2+1, n/2)	$z(m/2+1, n/2+1)$	0
0	0	0	0	...	0	0	n/u	n/u

* n/u - not used

scfft2dc/dzfft2dc

C-interface routine. Compute forward FFT of a real matrix and represent the complex conjugate-symmetric result in CCS format (in-place).

```
void scfft2dc ( float* r, int m, int n )
void dzfft2dc ( double* r, int m, int n )
```

Discussion

See the equations of the operations for the [Real-to-Complex Two-dimensional FFTs](#) above.

These routines are complementary to the complex-to-real transform routines [csfft2dc/zdfft2dc](#).

Input Parameters

- r* `float*` for `scfft2dc`
 `double*` for `dzfft2dc`
- Pointer to an array of size at least $(m+2, n+2)$, with its leading dimension equal to $(n+2)$. The first m rows and n columns of this array contain the real matrix to be transformed.
- [Table 3-7](#) presents the input data layout.
- m* `int`. Column transform length;
 m must be a power of 2.
- n* `int`. Row transform length;
 n must be a power of 2.

Table 3-7 C-interface Real Data Storage for a Real-to-Complex and Complex-to-Real Two-dimensional FFTs

<code>r(0,0)</code>	<code>r(0,1)</code>	...	<code>r(0,n-2)</code>	<code>r(0,n-1)</code>	n/u	n/u
<code>r(1,0)</code>	<code>r(1,1)</code>	...	<code>r(1,n-2)</code>	<code>r(1,n-1)</code>	n/u	n/u
<code>r(2,0)</code>	<code>r(2,1)</code>	...	<code>r(2,n-2)</code>	<code>r(2,n-1)</code>	n/u	n/u
<code>r(3,0)</code>	<code>r(3,1)</code>	...	<code>r(3,n-2)</code>	<code>r(3,n-1)</code>	n/u	n/u
...
<code>r(m-2,0)</code>	<code>r(m-2,1)</code>	...	<code>r(m-2,n-2)</code>	<code>r(m-2,n-1)</code>	n/u	n/u
<code>r(m-1,0)</code>	<code>r(m-1,1)</code>	...	<code>r(m-1,n-2)</code>	<code>r(m-1,n-1)</code>	n/u	n/u
n/u	n/u	...	n/u	n/u	n/u	n/u
n/u	n/u	...	n/u	n/u	n/u	n/u

Output Parameters

r The output real array $r(0:m+1, 0:n+1)$ contains the complex conjugate-symmetric matrix $z(0:m-1, 0:n-1)$ packed in CCS format for C-interface as follows:

- Columns 0 and $n/2$ contain in $m+2$ locations the complex conjugate-symmetric vectors $z(i, 0)$ and $z(i, n/2)$ in CCS format (see [Real-to-Complex One-dimensional FFTs](#) above).

The full complex vector $z(i, 0)$ is defined by:

$$z(i, 0) = \text{cplx}(r(i, 0), r(m/2+i+1, 0)), \quad 0 \leq i \leq m/2,$$

$$z(m/2+i, 0) = \text{conjg}(z(m/2-i, 0)), \quad 1 \leq i \leq m/2-1.$$

The full complex vector $z(i, n/2)$ is defined by:

$$z(i, n/2) = \text{cplx}(r(i, n/2), r(m/2+i+1, n/2)), \quad 0 \leq i \leq m/2,$$

$$z(m/2+i, n/2) = \text{conjg}(z(m/2-i, n/2)), \quad 1 \leq i \leq m/2-1.$$

- Columns from 1 to $n/2-1$ contain real parts, and columns from $n/2+2$ to n contain imaginary parts of complex vectors. These values for each vector are stored in m locations represented as follows

$$z(i, j) = \text{cplx}(r(i, j), r(i, n/2+1+j)),$$

$$0 \leq i \leq m-1, \quad 1 \leq j \leq n/2-1.$$

- The rest matrix elements can be obtained from

$$z(i, n/2+j) = \text{conjg}(z(i, n/2-j)),$$

$$0 \leq i \leq m-1, \quad 1 \leq j \leq n/2-1.$$

The storage of the complex conjugate-symmetric matrix z for C-interface is shown in [Table 3-8](#).

Table 3-8 C-interface Data Storage of CCS Format for the Real-to-Complex and Complex-to-Real Two-dimensional FFT

$z(0,0)$	$REz(0,1)$...	$REz(0, n/2-1)$	$z(0,n/2)$	0	$IMz(0,1)$...	$IMz(0, n/2-1)$	0
$REz(1,0)$	$REz(1,1)$...	$REz(1, n/2-1)$	$REz(1,n/2)$	0	$IMz(1,1)$...	$IMz(1, n/2-1)$	0
...	0	0
$REz(m/2-1, 0)$	$REz(m/2-1, 1)$...	$REz(m/2-1, n/2-1)$	$REz(m/2-1, n/2)$	0	$IMz(m/2-1, 1)$...	$IMz(m/2-1, n/2-1)$	0
$z(m/2,0)$	$REz(m/2,1)$...	$REz(m/2, n/2-1)$	$z(m/2,n/2)$	0	$IMz(m/2,1)$...	$IMz(m/2, n/2-1)$	0
0	$REz(m/2+1, 1)$...	$REz(m/2+1, n/2-1)$	0	0	$IMz(m/2+1, 1)$...	$IMz(m/2+1, n/2-1)$	0
$IMz(1,0)$	$REz(m/2+2, 1)$...	$REz(m/2+2, n/2-1)$	$IMz(1,n/2)$	0	$IMz(m/2+2, 1)$...	$IMz(m/2+2, n/2-1)$	0
...	0	0
$IMz(m/2-2, 0)$	$REz(m-1,1)$...	$REz(m-1, n/2-1)$	$IMz(m/2-2, n/2)$	0	$IMz(m-1,1)$...	$IMz(m-1, n/2-1)$	0
$IMz(m/2-1, 0)$	n/u	...	n/u	$IMz(m/2-1, n/2)$	n/u	n/u	...	n/u	n/u
0	n/u	...	n/u	0	n/u	n/u	...	n/u	n/u

Complex-to-Real Two-dimensional FFTs

Each of the complex-to-real routines computes a two-dimensional inverse FFT according to the mathematical equation:

$$t_{i,j} = \frac{1}{m*n} \sum_{k=0}^{m-1} \sum_{l=0}^{n-1} z_{k,l} * w_m^{i*k} * w_n^{j*l}, \quad 0 \leq i \leq m-1, \quad 0 \leq j \leq n-1$$

The mathematical input $z_{i,j}$, $0 \leq i \leq m-1$, $0 \leq j \leq n-1$, is a complex matrix of size (m,n) . Each column is the complex conjugate-symmetric vector as follows:

for $0 \leq j \leq n-1$,

$z(m/2+i, j) = \text{conjg}(z(m/2-i, j))$, $1 \leq i \leq m/2-1$.

Moreover, $z(0, j)$ and $z(m/2, j)$ are real values for $j=0$ and $j=n/2$.

This mathematical input can be stored in the complex two-dimensional array of size $(m/2+1, n/2+1)$ or in the real two-dimensional array of size $(m+2, n+2)$. For the details of data storage of CCS format see [Real-to-Complex One-dimensional FFTs](#) above.

The mathematical result of the transform is $t_{k,1} = \text{cplx}(r_{k,1}, 0)$, where $r_{k,1}$ is the real matrix, $0 \leq k \leq m-1$, $0 \leq l \leq n-1$.

csfft2d/zdffft2d

Fortran-interface routine.

Compute inverse FFT of a complex conjugate-symmetric matrix packed in CCS format (in-place).

```
call csfft2d ( r, m, n )
call zdffft2d ( r, m, n )
```

Discussion

See the equations of the operations for the [Complex-to-Real Two-dimensional FFTs](#) above. These routines are complementary to the real-to-complex transform routines [scfft2d/dzfft2d](#).

Input Parameters

r SINGLE PRECISION REAL*4 for csfft2d
 DOUBLE PRECISION REAL*8 for zdffft2d

Array, DIMENSION at least $(m+2, n+2)$, with its leading dimension equal to $(m+2)$. This array contains the complex conjugate-symmetric matrix in CCS format to be transformed. The input data layout is given in [Table 3-6](#).

- m* **INTEGER**. Column transform length (number of rows); *m* must be a power of 2.
- n* **INTEGER**. Row transform length (number of columns); *n* must be a power of 2.

Output Parameters

- r* Contains the real result returned by the transform. For the output data layout, see [Table 3-5](#).

csfft2dc/zdfft2dc

C-interface routines.

Compute inverse FFT of a complex conjugate-symmetric matrix packed in CCS format (in-place).

```
void csfft2dc ( float* r, int m, int n );
void zdfft2dc ( double* r, int m, int n );
```

Discussion

See the equations of the operations for the [Complex-to-Real Two-dimensional FFTs](#) above. These routines are complementary to the real-to-complex transform routines [scfft2dc/dzfft2dc](#).

Input Parameters

- r* **float*** for `csfft2dc`
double* for `zdfft2dc`
- Pointer to an array of size at least $(m+2, n+2)$, with its leading dimension equal to $(n+2)$. This array contains the complex conjugate-symmetric matrix in CCS format to be transformed. The input data layout is given in [Table 3-8](#).
- m* **int**. Column transform length; *m* must be a power of 2.

n `int`. Row transform length; *n* must be a power of 2.

Output Parameters

r Contains the real result returned by the transform. The output data layout is the same as that for the input data of `scfft2dc/dzfft2dc`. See [Table 3-7](#) for the details.

LAPACK Routines: Linear Equations

4

This chapter describes the Intel[®] Math Kernel Library implementation of routines from the LAPACK package that are used for solving systems of linear equations and performing a number of related computational tasks. The library includes LAPACK routines for both real and complex data.

Routines are supported for systems of equations with the following types of matrices:

- general
- banded
- symmetric or Hermitian positive-definite (both full and packed storage)
- symmetric or Hermitian positive-definite banded
- symmetric or Hermitian indefinite (both full and packed storage)
- symmetric or Hermitian indefinite banded
- triangular (both full and packed storage)
- triangular banded
- tridiagonal.

For each of the above matrix types, the library includes routines for performing the following computations: *factoring* the matrix (except for triangular matrices); *equilibrating* the matrix; *solving* a system of linear equations; *estimating the condition number* of a matrix; *refining* the solution of linear equations and computing its error bounds; *inverting* the matrix.

To solve a particular problem, you can either call two or more [computational routines](#) or call a corresponding [driver routine](#) that combines several tasks in one call, such as `?gesv` for factoring and solving. Thus, to solve a system of linear equations with a general matrix, you can first call `?getrf` (*LU* factorization) and then `?getrs` (computing the solution). Then, you might wish to call `?gerfs` to refine the solution and get the error bounds. Alternatively, you can just use the driver routine `?gesvx` which performs all these tasks in one call.

Routine Naming Conventions

For each routine introduced in this chapter, you can use the LAPACK name.

LAPACK names are listed in Tables 4-1 and 4-2, and have the structure `xyyzzz` or `xyyzz`, which is described below.

The initial letter `x` indicates the data type:

<code>s</code>	real, single precision	<code>c</code>	complex, single precision
<code>d</code>	real, double precision	<code>z</code>	complex, double precision

The second and third letters `yy` indicate the matrix type and storage scheme:

<code>ge</code>	general
<code>gb</code>	general band
<code>gt</code>	general tridiagonal
<code>po</code>	symmetric or Hermitian positive-definite
<code>pp</code>	symmetric or Hermitian positive-definite (packed storage)
<code>pb</code>	symmetric or Hermitian positive-definite band
<code>pt</code>	symmetric or Hermitian positive-definite tridiagonal
<code>sy</code>	symmetric indefinite
<code>sp</code>	symmetric indefinite (packed storage)
<code>he</code>	Hermitian indefinite
<code>hp</code>	Hermitian indefinite (packed storage)
<code>tr</code>	triangular
<code>tp</code>	triangular (packed storage)
<code>tb</code>	triangular band

For computational routines, the last three letters `zzz` indicate the computation performed:

<code>trf</code>	form a triangular matrix factorization
<code>trs</code>	solve the linear system with a factored matrix
<code>con</code>	estimate the matrix condition number
<code>rfs</code>	refine the solution and compute error bounds
<code>tri</code>	compute the inverse matrix using the factorization
<code>equ</code>	equilibrate a matrix.

For example, the routine `sgetrf` performs the triangular factorization of general real matrices in single precision; the corresponding routine for complex matrices is `cgetrf`.

For driver routines, the names can end either with `-sv` (meaning a *simple* driver), or with `-svx` (meaning an *expert* driver).

Matrix Storage Schemes

LAPACK routines use the following matrix storage schemes:

- *Full storage*: a matrix A is stored in a two-dimensional array a , with the matrix element a_{ij} stored in the array element $a(i, j)$.
- *Packed storage* scheme allows you to store symmetric, Hermitian, or triangular matrices more compactly: the upper or lower triangle of the matrix is packed by columns in a one-dimensional array.
- *Band storage*: an m by n band matrix with kl sub-diagonals and ku super-diagonals is stored compactly in a two-dimensional array ab with $kl+ku+1$ rows and n columns. Columns of the matrix are stored in the corresponding columns of the array, and *diagonals* of the matrix are stored in rows of the array.

In Chapters 4 and 5, arrays that hold matrices in packed storage have names ending in p ; arrays with matrices in band storage have names ending in b .

For more information on matrix storage schemes, see [Matrix Arguments](#) in Appendix A.

Mathematical Notation

Descriptions of LAPACK routines use the following notation:

$Ax = b$	A system of linear equations with an n by n matrix $A = \{a_{ij}\}$, a right-hand side vector $b = \{b_i\}$, and an unknown vector $x = \{x_i\}$.
$AX = B$	A set of systems with a common matrix A and multiple right-hand sides. The columns of B are individual right-hand sides, and the columns of X are the corresponding solutions.
$ x $	the vector with elements $ x_i $ (absolute values of x_i).
$ A $	the matrix with elements $ a_{ij} $ (absolute values of a_{ij}).
$\ x\ _\infty = \max_i x_i $	The <i>infinity-norm</i> of the vector x .
$\ A\ _\infty = \max_i \sum_j a_{ij} $	The <i>infinity-norm</i> of the matrix A .
$\ A\ _1 = \max_j \sum_i a_{ij} $	The <i>one-norm</i> of the matrix A . $\ A\ _1 = \ A^T\ _\infty = \ A^H\ _\infty$
$\kappa(A) = \ A\ \ A^{-1}\ $	The <i>condition number</i> of the matrix A .

Error Analysis

In practice, most computations are performed with rounding errors. Besides, you often need to solve a system $Ax = b$ where the data (the elements of A and b) are not known exactly. Therefore, it's important to understand how the data errors and rounding errors can affect the solution x .

Data perturbations. If x is the exact solution of $Ax = b$, and $x + \delta x$ is the exact solution of a perturbed problem $(A + \delta A)x = (b + \delta b)$, then

$$\frac{\|\delta x\|}{\|x\|} \leq \kappa(A) \left(\frac{\|\delta A\|}{\|A\|} + \frac{\|\delta b\|}{\|b\|} \right), \text{ where } \kappa(A) = \|A\| \|A^{-1}\|.$$

In other words, relative errors in A or b may be amplified in the solution vector x by a factor $\kappa(A) = \|A\| \|A^{-1}\|$ called the *condition number* of A .

Rounding errors have the same effect as relative perturbations $c(n)\epsilon$ in the original data. Here ϵ is the *machine precision*, and $c(n)$ is a modest function of the matrix order n . The corresponding solution error is $\|\delta x\| / \|x\| \leq c(n)\kappa(A)\epsilon$. (The value of $c(n)$ is seldom greater than $10n$.)

Thus, if your matrix A is *ill-conditioned* (that is, its condition number $\kappa(A)$ is very large), then the error in the solution x is also large; you may even encounter a complete loss of precision. LAPACK provides routines that allow you to estimate $\kappa(A)$ (see [Routines for Estimating the Condition Number](#)) and also give you a more precise estimate for the actual solution error (see [Refining the Solution and Estimating Its Error](#)).

Computational Routines

[Table 4-1](#) lists the LAPACK computational routines for factorizing, equilibrating, and inverting *real* matrices, estimating their condition numbers, solving systems of equations with real matrices, refining the solution, and estimating its error.

[Table 4-2](#) lists similar routines for *complex* matrices.

Table 4-1 Computational Routines for Systems of Equations with Real Matrices

Matrix type, storage scheme	Factorize matrix	Equilibrate matrix	Solve system	Condition number	Estimate error	Invert matrix
general	?getrf	?geequ	?getrs	?gecon	?gerfs	?getri
general band	?gbtrf	?gbequ	?gbtrs	?gbcon	?gbrfs	
general tridiagonal	?gttrf		?gttrs	?gtcon	?gtrfs	
symmetric positive-definite	?potrf	?poequ	?potrs	?pocon	?porfs	?potri
symmetric positive-definite, packed storage	?pptrf	?ppequ	?pptrs	?ppcon	?pprfs	?pptri
symmetric positive-definite, band	?pbtrf	?pbequ	?pbtrs	?pbcon	?pbrfs	
symmetric positive-definite, tridiagonal	?pttrf		?pttrs	?ptcon	?ptrfs	
symmetric indefinite	?sytrf		?sytrs	?sycon	?syrfs	?sytri
symmetric indefinite, packed storage	?sptrf		?sptrs	?spcon	?sprfs	?sptri
triangular			?trtrs	?trcon	?trrfs	?trtri
triangular, packed storage			?tptrs	?tpcon	?tprfs	?tptri
triangular band			?tbtrs	?tbcon	?tbrfs	

In this table ? denotes **s** (single precision) or **d** (double precision).

Table 4-2 Computational Routines for Systems of Equations with Complex Matrices

Matrix type, storage scheme	Factorize matrix	Equilibrate matrix	Solve system	Condition number	Estimate error	Invert matrix
general	<u>?getrf</u>	<u>?geequ</u>	<u>?getrs</u>	<u>?gecon</u>	<u>?gerfs</u>	<u>?getri</u>
general band	<u>?gbtrf</u>	<u>?gbequ</u>	<u>?gbtrs</u>	<u>?gbcon</u>	<u>?gbrfs</u>	
general tridiagonal	<u>?gttrf</u>		<u>?gttrs</u>	<u>?gtcon</u>	<u>?gtrfs</u>	
Hermitian positive-definite	<u>?potrf</u>	<u>?poequ</u>	<u>?potrs</u>	<u>?pocon</u>	<u>?porfs</u>	<u>?potri</u>
Hermitian positive-definite, packed storage	<u>?pptrf</u>	<u>?ppequ</u>	<u>?pptrs</u>	<u>?ppcon</u>	<u>?pprfs</u>	<u>?pptri</u>
Hermitian positive-definite, band	<u>?pbtrf</u>	<u>?pbequ</u>	<u>?pbtrs</u>	<u>?pbcon</u>	<u>?pbrfs</u>	
Hermitian positive-definite, tridiagonal	<u>?pttrf</u>		<u>?pttrs</u>	<u>?ptcon</u>	<u>?ptrfs</u>	
Hermitian indefinite	<u>?hetrf</u>		<u>?hetrs</u>	<u>?hecon</u>	<u>?herfs</u>	<u>?hetri</u>
symmetric indefinite	<u>?sytrf</u>		<u>?sytrs</u>	<u>?sycon</u>	<u>?syrf</u>	<u>?sytri</u>
Hermitian indefinite, packed storage	<u>?hptrf</u>		<u>?hptrs</u>	<u>?hpcon</u>	<u>?hprfs</u>	<u>?hptri</u>
symmetric indefinite, packed storage	<u>?sptrf</u>		<u>?sptrs</u>	<u>?spcon</u>	<u>?sprfs</u>	<u>?sptri</u>
triangular			<u>?trtrs</u>	<u>?trcon</u>	<u>?trrfs</u>	<u>?trtri</u>
triangular, packed storage			<u>?tptrs</u>	<u>?tpcon</u>	<u>?tprfs</u>	<u>?tptri</u>
triangular band			<u>?tbtrs</u>	<u>?tbcon</u>	<u>?tbrfs</u>	

In this table ? stands for **c** (single precision complex) or **z** (double precision complex).

Routines for Matrix Factorization

This section describes the LAPACK routines for matrix factorization. The following factorizations are supported:

- LU factorization
- Cholesky factorization of real symmetric positive-definite matrices
- Cholesky factorization of Hermitian positive-definite matrices
- Bunch-Kaufman factorization of real and complex symmetric matrices
- Bunch-Kaufman factorization of Hermitian matrices.

You can compute the LU factorization using full and band storage of matrices; the Cholesky factorization using full, packed, and band storage; and the Bunch-Kaufman factorization using full and packed storage.

?getrf

Computes the LU factorization of a general m by n matrix.

```
call sgetrf ( m, n, a, lda, ipiv, info )
call dgetrf ( m, n, a, lda, ipiv, info )
call cgetrf ( m, n, a, lda, ipiv, info )
call zgetrf ( m, n, a, lda, ipiv, info )
```

Discussion

The routine forms the LU factorization of a general m by n matrix A as

$$A = PLU$$

where P is a permutation matrix, L is lower triangular with unit diagonal elements (lower trapezoidal if $m > n$) and U is upper triangular (upper trapezoidal if $m < n$). Usually A is square ($m = n$), and both L and U are triangular. The routine uses partial pivoting, with row interchanges.

Input Parameters

<i>m</i>	INTEGER. The number of rows in the matrix <i>A</i> ($m \geq 0$).
<i>n</i>	INTEGER. The number of columns in <i>A</i> ($n \geq 0$).
<i>a</i>	REAL for <code>sgetrf</code> DOUBLE PRECISION for <code>dgetrf</code> COMPLEX for <code>cgetrf</code> DOUBLE COMPLEX for <code>zgetrf</code> . Array, DIMENSION (<i>lda</i> , *). Contains the matrix <i>A</i> . The second dimension of <i>a</i> must be at least $\max(1, n)$.
<i>lda</i>	INTEGER. The first dimension of <i>a</i> .

Output Parameters

<i>a</i>	Overwritten by <i>L</i> and <i>U</i> . The unit diagonal elements of <i>L</i> are not stored.
<i>ipiv</i>	INTEGER. Array, DIMENSION at least $\max(1, \min(m, n))$. The pivot indices: row <i>i</i> was interchanged with row <i>ipiv</i> (<i>i</i>).
<i>info</i>	INTEGER. If <i>info</i> =0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = <i>i</i> , <i>u</i> _{<i>ii</i>} is 0. The factorization has been completed, but <i>U</i> is exactly singular. Division by 0 will occur if you use the factor <i>U</i> for solving a system of linear equations.

Application Notes

The computed *L* and *U* are the exact factors of a perturbed matrix $A + E$, where

$$|E| \leq c(\min(m, n)) \varepsilon P|L||U|$$

$c(n)$ is a modest linear function of n , and ε is the machine precision.

The approximate number of floating-point operations for real flavors is

$$\begin{aligned} (2/3)n^3 & \quad \text{if } m = n, \\ (1/3)n^2(3m-n) & \quad \text{if } m > n, \end{aligned}$$

$$(1/3)m^2(3n-m) \quad \text{if } m < n.$$

The number of operations for complex flavors is 4 times greater.

After calling this routine with $m = n$, you can call the following:

- [?getrs](#) to solve $AX = B$ or $A^T X = B$ or $A^H X = B$;
- [?gecon](#) to estimate the condition number of A ;
- [?getri](#) to compute the inverse of A .

?gbtrf

Computes the *LU* factorization of a general *m* by *n* band matrix.

```
call sgbtrf ( m, n, kl, ku, ab, ldab, ipiv, info )
call dgbtrf ( m, n, kl, ku, ab, ldab, ipiv, info )
call cgbtrf ( m, n, kl, ku, ab, ldab, ipiv, info )
call zgbtrf ( m, n, kl, ku, ab, ldab, ipiv, info )
```

Discussion

The routine forms the *LU* factorization of a general *m* by *n* band matrix *A* with *kl* non-zero sub-diagonals and *ku* non-zero super-diagonals. Usually *A* is square (*m* = *n*), and then

$$A = PLU$$

where *P* is a permutation matrix; *L* is lower triangular with unit diagonal elements and at most *kl* non-zero elements in each column; *U* is an upper triangular band matrix with *kl* + *ku* super-diagonals. The routine uses partial pivoting, with row interchanges (which creates the additional *kl* super-diagonals in *U*).

Input Parameters

<i>m</i>	INTEGER. The number of rows in the matrix <i>A</i> (<i>m</i> ≥ 0).
<i>n</i>	INTEGER. The number of columns in <i>A</i> (<i>n</i> ≥ 0).
<i>kl</i>	INTEGER. The number of sub-diagonals within the band of <i>A</i> (<i>kl</i> ≥ 0).
<i>ku</i>	INTEGER. The number of super-diagonals within the band of <i>A</i> (<i>ku</i> ≥ 0).
<i>ab</i>	REAL for sgbtrf DOUBLE PRECISION for dgbtrf COMPLEX for cgbtrf DOUBLE COMPLEX for zgbtrf. Array, DIMENSION (<i>ldab</i> , *).

The array *ab* contains the matrix *A* in band storage (see [Matrix Storage Schemes](#)).

The second dimension of *ab* must be at least $\max(1, n)$.

ldab **INTEGER**. The first dimension of the array *ab*.
($ldab \geq 2kl + ku + 1$)

Output Parameters

ab Overwritten by *L* and *U*. The diagonal and $kl + ku$ super-diagonals of *U* are stored in the first $1 + kl + ku$ rows of *ab*. The multipliers used to form *L* are stored in the next *kl* rows.

ipiv **INTEGER**.
Array, **DIMENSION** at least $\max(1, \min(m, n))$.
The pivot indices: row *i* was interchanged with row *ipiv(i)*.

info **INTEGER**. If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i*, u_{ii} is 0. The factorization has been completed, but *U* is exactly singular. Division by 0 will occur if you use the factor *U* for solving a system of linear equations.

Application Notes

The computed *L* and *U* are the exact factors of a perturbed matrix $A + E$, where

$$|E| \leq c(kl + ku + 1)\epsilon P|L||U|$$

$c(k)$ is a modest linear function of *k*, and ϵ is the machine precision.

The total number of floating-point operations for real flavors varies between approximately $2n(ku+1)kl$ and $2n(kl+ku+1)kl$. The number of operations for complex flavors is 4 times greater. All these estimates assume that *kl* and *ku* are much less than $\min(m, n)$.

After calling this routine with $m = n$, you can call the following:

[?gbtrs](#) to solve $AX = B$ or $A^T X = B$ or $A^H X = B$;

[?gbcon](#) to estimate the condition number of *A*.

?gttrf

Computes the *LU* factorization of a tridiagonal matrix.

```
call sgttrf ( n, dl, d, du, du2, ipiv, info )
call dgttrf ( n, dl, d, du, du2, ipiv, info )
call cgttrf ( n, dl, d, du, du2, ipiv, info )
call zgttrf ( n, dl, d, du, du2, ipiv, info )
```

Discussion

The routine computes the *LU* factorization of a real or complex tridiagonal matrix *A* in the form

$$A = PLU$$

where *P* is a permutation matrix; *L* is lower bidiagonal with unit diagonal elements; and *U* is an upper triangular matrix with nonzeros in only the main diagonal and first two superdiagonals. The routine uses elimination with partial pivoting and row interchanges .

Input Parameters

n **INTEGER**. The order of the matrix *A* ($n \geq 0$).

dl, d, du **REAL** for *sgttrf*
DOUBLE PRECISION for *dgttrf*
COMPLEX for *cgttrf*
DOUBLE COMPLEX for *zgttrf*.

Arrays containing elements of *A*.
The array *dl* of dimension ($n - 1$) contains the sub-diagonal elements of *A*.
The array *d* of dimension *n* contains the diagonal elements of *A*.
The array *du* of dimension ($n - 1$) contains the super-diagonal elements of *A*.

Output Parameters

- dl* Overwritten by the $(n-1)$ multipliers that define the matrix L from the LU factorization of A .
- d* Overwritten by the n diagonal elements of the upper triangular matrix U from the LU factorization of A .
- du* Overwritten by the $(n-1)$ elements of the first super-diagonal of U .
- du2* REAL for `sgttrf`
DOUBLE PRECISION for `dgttrf`
COMPLEX for `cgttrf`
DOUBLE COMPLEX for `zgttrf`.
Array, dimension $(n-2)$. On exit, *du2* contains $(n-2)$ elements of the second super-diagonal of U .
- ipiv* INTEGER.
Array, dimension (n) .
The pivot indices: row i was interchanged with row *ipiv*(i).
- info* INTEGER. If *info* = 0, the execution is successful.
If *info* = $-i$, the i th parameter had an illegal value.
If *info* = i , u_{ii} is 0. The factorization has been completed, but U is exactly singular. Division by zero will occur if you use the factor U for solving a system of linear equations.

Application Notes

- [?gbtrs](#) to solve $AX = B$ or $A^T X = B$ or $A^H X = B$;
- [?gbcon](#) to estimate the condition number of A .

?potrf

Computes the Cholesky factorization of a symmetric (Hermitian) positive-definite matrix.

```
call spotrf ( uplo, n, a, lda, info )
call dpotrf ( uplo, n, a, lda, info )
call cpotrf ( uplo, n, a, lda, info )
call zpotrf ( uplo, n, a, lda, info )
```

Discussion

This routine forms the Cholesky factorization of a symmetric positive-definite or, for complex data, Hermitian positive-definite matrix A :

$$A = U^H U \quad \text{if } uplo = 'U'$$

$$A = LL^H \quad \text{if } uplo = 'L'$$

where L is a lower triangular matrix and U is upper triangular.

Input Parameters

<code>uplo</code>	CHARACTER*1. Must be 'U' or 'L'. Indicates whether the upper or lower triangular part of A is stored and how A is factored: If <code>uplo = 'U'</code> , the array <code>a</code> stores the upper triangular part of the matrix A , and A is factored as $U^H U$. If <code>uplo = 'L'</code> , the array <code>a</code> stores the lower triangular part of the matrix A ; A is factored as LL^H .
<code>n</code>	INTEGER. The order of matrix A ($n \geq 0$).
<code>a</code>	REAL for <code>spotrf</code> DOUBLE PRECISION for <code>dpotrf</code> COMPLEX for <code>cpotrf</code> DOUBLE COMPLEX for <code>zpotrf</code> . Array, DIMENSION (<code>lda, *</code>).

The array *a* contains either the upper or the lower triangular part of the matrix *A* (see *uplo*).
The second dimension of *a* must be at least $\max(1, n)$.

lda **INTEGER**. The first dimension of *a*.

Output Parameters

a The upper or lower triangular part of *a* is overwritten by the Cholesky factor *U* or *L*, as specified by *uplo*.

info **INTEGER**. If *info*=0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i*, the leading minor of order *i* (and hence the matrix *A* itself) is not positive-definite, and the factorization could not be completed. This may indicate an error in forming the matrix *A*.

Application Notes

If *uplo* = 'U', the computed factor *U* is the exact factor of a perturbed matrix $A + E$, where

$$|E| \leq c(n)\epsilon |U^H| |U|, \quad |e_{ij}| \leq c(n)\epsilon \sqrt{a_{ii}a_{jj}}$$

$c(n)$ is a modest linear function of n , and ϵ is the machine precision.

A similar estimate holds for *uplo* = 'L'.

The total number of floating-point operations is approximately $(1/3)n^3$ for real flavors or $(4/3)n^3$ for complex flavors.

After calling this routine, you can call the following:

[?potrs](#) to solve $AX = B$;
[?pocon](#) to estimate the condition number of *A*;
[?potri](#) to compute the inverse of *A*.

?pptrf

Computes the Cholesky factorization of a symmetric (Hermitian) positive-definite matrix using packed storage.

```
call spptrf ( uplo, n, ap, info )
call dpptrf ( uplo, n, ap, info )
call cpptrf ( uplo, n, ap, info )
call zpptrf ( uplo, n, ap, info )
```

Discussion

This routine forms the Cholesky factorization of a symmetric positive-definite or, for complex data, Hermitian positive-definite packed matrix A :

$$A = U^H U \quad \text{if } uplo = 'U'$$

$$A = LL^H \quad \text{if } uplo = 'L'$$

where L is a lower triangular matrix and U is upper triangular.

Input Parameters

<i>uplo</i>	CHARACTER*1 . Must be 'U' or 'L'. Indicates whether the upper or lower triangular part of A is packed in the array <i>ap</i> , and how A is factored: If <i>uplo</i> = 'U', the array <i>ap</i> stores the upper triangular part of the matrix A , and A is factored as $U^H U$. If <i>uplo</i> = 'L', the array <i>ap</i> stores the lower triangular part of the matrix A ; A is factored as LL^H .
<i>n</i>	INTEGER . The order of matrix A ($n \geq 0$).
<i>ap</i>	REAL for <i>spptrf</i> DOUBLE PRECISION for <i>dpptrf</i> COMPLEX for <i>cpptrf</i> DOUBLE COMPLEX for <i>zpptrf</i> . Array, DIMENSION at least $\max(1, n(n+1)/2)$.

The array *ap* contains either the upper or the lower triangular part of the matrix *A* (as specified by *uplo*) in *packed storage* (see [Matrix Storage Schemes](#)).

Output Parameters

- ap* The upper or lower triangular part of *A* in packed storage is overwritten by the Cholesky factor *U* or *L*, as specified by *uplo*.
- info* **INTEGER.** If *info*=0, the execution is successful. If *info* = -*i*, the *i*th parameter had an illegal value. If *info* = *i*, the leading minor of order *i* (and hence the matrix *A* itself) is not positive-definite, and the factorization could not be completed. This may indicate an error in forming the matrix *A*.

Application Notes

If *uplo* = 'U', the computed factor *U* is the exact factor of a perturbed matrix *A* + *E*, where

$$|E| \leq c(n)\varepsilon |U^H||U|, \quad |e_{ij}| \leq c(n)\varepsilon \sqrt{a_{ii}a_{jj}}$$

c(n) is a modest linear function of *n*, and ε is the machine precision.

A similar estimate holds for *uplo* = 'L'.

The total number of floating-point operations is approximately $(1/3)n^3$ for real flavors and $(4/3)n^3$ for complex flavors.

After calling this routine, you can call the following:

- [?pptrs](#) to solve $AX = B$;
- [?ppcon](#) to estimate the condition number of *A*;
- [?pptri](#) to compute the inverse of *A*.

?pbtrf

Computes the Cholesky factorization of a symmetric (Hermitian) positive-definite band matrix.

```
call spbtrf ( uplo, n, kd, ab, ldab, info )
call dpbtrf ( uplo, n, kd, ab, ldab, info )
call cpbtrf ( uplo, n, kd, ab, ldab, info )
call zpbtrf ( uplo, n, kd, ab, ldab, info )
```

Discussion

This routine forms the Cholesky factorization of a symmetric positive-definite or, for complex data, Hermitian positive-definite band matrix A :

$$A = U^H U \quad \text{if } uplo = 'U'$$

$$A = LL^H \quad \text{if } uplo = 'L'$$

where L is a lower triangular matrix and U is upper triangular.

Input Parameters

<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. Indicates whether the upper or lower triangular part of A is stored in the array <i>ab</i> , and how A is factored: If <i>uplo</i> = 'U', the array <i>ab</i> stores the upper triangular part of the matrix A , and A is factored as $U^H U$. If <i>uplo</i> = 'L', the array <i>ab</i> stores the lower triangular part of the matrix A ; A is factored as LL^H .
<i>n</i>	INTEGER. The order of matrix A ($n \geq 0$).
<i>kd</i>	INTEGER. The number of super-diagonals or sub-diagonals in the matrix A ($kd \geq 0$).
<i>ab</i>	REAL for spbtrf DOUBLE PRECISION for dpbtrf COMPLEX for cpbtrf DOUBLE COMPLEX for zpbtrf. Array, DIMENSION (<i>ldab</i> ,*).

The array *ap* contains either the upper or the lower triangular part of the matrix *A* (as specified by *uplo*) in *band storage* (see [Matrix Storage Schemes](#)).

The second dimension of *ab* must be at least $\max(1, n)$.

ldab **INTEGER**. The first dimension of the array *ab*.
(*ldab* $\geq kd + 1$)

Output Parameters

ap The upper or lower triangular part of *A* (in band storage) is overwritten by the Cholesky factor *U* or *L*, as specified by *uplo*.

info **INTEGER**. If *info*=0, the execution is successful.
If *info* = *-i*, the *i*th parameter had an illegal value.
If *info* = *i*, the leading minor of order *i* (and hence the matrix *A* itself) is not positive-definite, and the factorization could not be completed. This may indicate an error in forming the matrix *A*.

Application Notes

If *uplo* = 'U', the computed factor *U* is the exact factor of a perturbed matrix $A + E$, where

$$|E| \leq c(kd + 1)\epsilon \|U^H\| \|U\|, \quad |e_{ij}| \leq c(kd + 1)\epsilon \sqrt{a_{ii}a_{jj}}$$

$c(n)$ is a modest linear function of *n*, and ϵ is the machine precision.

A similar estimate holds for *uplo* = 'L'.

The total number of floating-point operations for real flavors is approximately $n(kd+1)^2$. The number of operations for complex flavors is 4 times greater. All these estimates assume that *kd* is much less than *n*.

After calling this routine, you can call the following:

[?pbtrs](#) to solve $AX = B$;
[?pbcon](#) to estimate the condition number of *A*;

?pttrf

Computes the factorization of a symmetric (Hermitian) positive-definite tridiagonal matrix.

```
call sptrrf ( n, d, e, info )
call dpttrf ( n, d, e, info )
call cpttrf ( n, d, e, info )
call zpttrf ( n, d, e, info )
```

Discussion

This routine forms the factorization of a symmetric positive-definite or, for complex data, Hermitian positive-definite tridiagonal matrix A :

$A = LDL^H$, where D is diagonal and L is unit lower bidiagonal. The factorization may also be regarded as having the form $A = U^H DU$, where D is unit upper bidiagonal.

Input Parameters

n **INTEGER**. The order of the matrix A ($n \geq 0$).

d **REAL** for `sptrrf`, `cpttrf`
DOUBLE PRECISION for `dpttrf`, `zpttrf`.
 Array, dimension (n). Contains the diagonal elements of A .

e **REAL** for `sptrrf`
DOUBLE PRECISION for `dpttrf`
COMPLEX for `cpttrf`
DOUBLE COMPLEX for `zpttrf`.
 Array, dimension ($n - 1$). Contains the sub-diagonal elements of A .

Output Parameters

d Overwritten by the n diagonal elements of the diagonal matrix D from the LDL^H factorization of A .

e Overwritten by the $(n - 1)$ off-diagonal elements of the unit bidiagonal factor L or U from the factorization of A .

info **INTEGER.** If *info*=0, the execution is successful.
If *info* = $-i$, the i th parameter had an illegal value.
If *info* = i , the leading minor of order i (and hence the matrix A itself) is not positive-definite; if $i < n$, the factorization could not be completed, while if $i = n$, the factorization was completed, but $d(n) = 0$.

?sytrf

Computes the Bunch-Kaufman factorization of a symmetric matrix.

```
call ssytrf ( uplo, n, a, lda, ipiv, work, lwork, info )
call dsytrf ( uplo, n, a, lda, ipiv, work, lwork, info )
call csytrf ( uplo, n, a, lda, ipiv, work, lwork, info )
call zsytrf ( uplo, n, a, lda, ipiv, work, lwork, info )
```

Discussion

This routine forms the Bunch-Kaufman factorization of a symmetric matrix:

if `uplo='U'`, $A = PUDU^T P^T$
 if `uplo='L'`, $A = PLDL^T P^T$

where A is the input matrix, P is a permutation matrix, U and L are upper and lower triangular matrices with unit diagonal, and D is a symmetric block-diagonal matrix with 1-by-1 and 2-by-2 diagonal blocks. U and L have 2-by-2 unit diagonal blocks corresponding to the 2-by-2 blocks of D .

Input Parameters

`uplo` CHARACTER*1. Must be 'U' or 'L'.
 Indicates whether the upper or lower triangular part of A is stored and how A is factored:
 If `uplo = 'U'`, the array `a` stores the upper triangular part of the matrix A , and A is factored as $PUDU^T P^T$.
 If `uplo = 'L'`, the array `a` stores the lower triangular part of the matrix A ; A is factored as $PLDL^T P^T$.

`n` INTEGER. The order of matrix A ($n \geq 0$).

`a` REAL for `ssytrf`
 DOUBLE PRECISION for `dsytrf`
 COMPLEX for `csytrf`
 DOUBLE COMPLEX for `zsytrf`.
 Array, DIMENSION (`lda`, *).

The array *a* contains either the upper or the lower triangular part of the matrix *A* (see *uplo*).

The second dimension of *a* must be at least $\max(1, n)$.

lda INTEGER. The first dimension of *a*; at least $\max(1, n)$.

work Same type as *a*. Workspace array of dimension *lwork*

lwork INTEGER. The size of the *work* array ($lwork \geq n$)

See [Application notes](#) for the suggested value of *lwork*.

Output Parameters

a The upper or lower triangular part of *a* is overwritten by details of the block-diagonal matrix *D* and the multipliers used to obtain the factor *U* (or *L*).

work(1) If *info*=0, on exit *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

ipiv INTEGER.

Array, DIMENSION at least $\max(1, n)$.

Contains details of the interchanges and the block structure of *D*.

If $ipiv(i) = k > 0$, then d_{ii} is a 1-by-1 block, and the *i*th row and column of *A* was interchanged with the *k*th row and column.

If $uplo = 'U'$ and $ipiv(i) = ipiv(i-1) = -m < 0$, then *D* has a 2-by-2 block in rows/columns *i* and *i-1*, and (*i-1*)th row and column of *A* was interchanged with the *m*th row and column.

If $uplo = 'L'$ and $ipiv(i) = ipiv(i+1) = -m < 0$, then *D* has a 2-by-2 block in rows/columns *i* and *i+1*, and (*i+1*)th row and column of *A* was interchanged with the *m*th row and column.

info INTEGER. If *info*=0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

If *info* = *i*, d_{ii} is 0. The factorization has been completed, but *D* is exactly singular. Division by 0 will occur if you use *D* for solving a system of linear equations.

Application Notes

For better performance, try using $lwork = n * blocksize$, where $blocksize$ is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*.

If you are in doubt how much workspace to supply, use a generous value of $lwork$ for the first run. On exit, examine $work(1)$ and use this value for subsequent runs.

The 2-by-2 unit diagonal blocks and the unit diagonal elements of U and L are not stored. The remaining elements of U and L are stored in the corresponding columns of the array a , but additional row interchanges are required to recover U or L explicitly (which is seldom necessary).

If $ipiv(i) = i$ for all $i = 1 . . . n$, then all off-diagonal elements of U (L) are stored explicitly in the corresponding elements of the array a .

If $uplo = 'U'$, the computed factors U and D are the exact factors of a perturbed matrix $A + E$, where

$$|E| \leq c(n) \varepsilon P |U| |D| |U^T| P^T$$

$c(n)$ is a modest linear function of n , and ε is the machine precision.

A similar estimate holds for the computed L and D when $uplo = 'L'$.

The total number of floating-point operations is approximately $(1/3)n^3$ for real flavors or $(4/3)n^3$ for complex flavors.

After calling this routine, you can call the following:

?sytrs	to solve $AX = B$;
?sycon	to estimate the condition number of A ;
?sytri	to compute the inverse of A .

?hetrf

Computes the Bunch-Kaufman factorization of a complex Hermitian matrix.

```
call chetrf ( uplo, n, a, lda, ipiv, work, lwork, info )
call zhetrf ( uplo, n, a, lda, ipiv, work, lwork, info )
```

Discussion

This routine forms the Bunch-Kaufman factorization of a Hermitian matrix:

if `uplo='U'`, $A = PUDU^H P^T$
 if `uplo='L'`, $A = PLDL^H P^T$

where A is the input matrix, P is a permutation matrix, U and L are upper and lower triangular matrices with unit diagonal, and D is a Hermitian block-diagonal matrix with 1-by-1 and 2-by-2 diagonal blocks. U and L have 2-by-2 unit diagonal blocks corresponding to the 2-by-2 blocks of D .

Input Parameters

`uplo` CHARACTER*1. Must be 'U' or 'L'.
 Indicates whether the upper or lower triangular part of A is stored and how A is factored:
 If `uplo = 'U'`, the array `a` stores the upper triangular part of the matrix A , and A is factored as $PUDU^H P^T$.
 If `uplo = 'L'`, the array `a` stores the lower triangular part of the matrix A ; A is factored as $PLDL^H P^T$.

`n` INTEGER. The order of matrix A ($n \geq 0$).

`a` COMPLEX for `chetrf`
 DOUBLE COMPLEX for `zhetrf`.
 Array, DIMENSION (`lda`, *).
 The array `a` contains either the upper or the lower triangular part of the matrix A (see `uplo`).
 The second dimension of `a` must be at least $\max(1, n)$.

<i>lda</i>	INTEGER. The first dimension of <i>a</i> ; at least $\max(1, n)$.
<i>work</i>	Same type as <i>a</i> . Workspace array of dimension <i>lwork</i>
<i>lwork</i>	INTEGER. The size of the <i>work</i> array ($lwork \geq n$) See Application notes for the suggested value of <i>lwork</i> .

Output Parameters

<i>a</i>	The upper or lower triangular part of <i>a</i> is overwritten by details of the block-diagonal matrix <i>D</i> and the multipliers used to obtain the factor <i>U</i> (or <i>L</i>).
<i>work(1)</i>	If <i>info</i> =0, on exit <i>work(1)</i> contains the minimum value of <i>lwork</i> required for optimum performance. Use this <i>lwork</i> for subsequent runs.
<i>ipiv</i>	INTEGER. Array, DIMENSION at least $\max(1, n)$. Contains details of the interchanges and the block structure of <i>D</i> . If $ipiv(i) = k > 0$, then d_{ii} is a 1-by-1 block, and the <i>i</i> th row and column of <i>A</i> was interchanged with the <i>k</i> th row and column. If $uplo = 'U'$ and $ipiv(i) = ipiv(i-1) = -m < 0$, then <i>D</i> has a 2-by-2 block in rows/columns <i>i</i> and <i>i-1</i> , and (<i>i-1</i>)th row and column of <i>A</i> was interchanged with the <i>m</i> th row and column. If $uplo = 'L'$ and $ipiv(i) = ipiv(i+1) = -m < 0$, then <i>D</i> has a 2-by-2 block in rows/columns <i>i</i> and <i>i+1</i> , and (<i>i+1</i>)th row and column of <i>A</i> was interchanged with the <i>m</i> th row and column.
<i>info</i>	INTEGER. If <i>info</i> =0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = <i>i</i> , d_{ii} is 0. The factorization has been completed, but <i>D</i> is exactly singular. Division by 0 will occur if you use <i>D</i> for solving a system of linear equations.

Application Notes

This routine is suitable for Hermitian matrices that are not known to be positive-definite. If A is in fact positive-definite, the routine does not perform interchanges, and no 2-by-2 diagonal blocks occur in D .

For better performance, try using $lwork = n * blocksize$, where $blocksize$ is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*.

If you are in doubt how much workspace to supply, use a generous value of $lwork$ for the first run. On exit, examine $work(1)$ and use this value for subsequent runs.

The 2-by-2 unit diagonal blocks and the unit diagonal elements of U and L are not stored. The remaining elements of U and L are stored in the corresponding columns of the array a , but additional row interchanges are required to recover U or L explicitly (which is seldom necessary).

If $ipiv(i) = i$ for all $i = 1 \dots n$, then all off-diagonal elements of U (L) are stored explicitly in the corresponding elements of the array a .

If $uplo = 'U'$, the computed factors U and D are the exact factors of a perturbed matrix $A + E$, where

$$|E| \leq c(n) \varepsilon P |U| |D| |U^T| P^T$$

$c(n)$ is a modest linear function of n , and ε is the machine precision.

A similar estimate holds for the computed L and D when $uplo = 'L'$.

The total number of floating-point operations is approximately $(4/3)n^3$.

After calling this routine, you can call the following:

- [?hetrs](#) to solve $AX = B$;
- [?hecon](#) to estimate the condition number of A ;
- [?hetri](#) to compute the inverse of A .

?sptf

Computes the Bunch-Kaufman factorization of a symmetric matrix using packed storage.

```
call ssptf ( uplo, n, ap, ipiv, info )
call dsptf ( uplo, n, ap, ipiv, info )
call csptf ( uplo, n, ap, ipiv, info )
call zsptf ( uplo, n, ap, ipiv, info )
```

Discussion

This routine forms the Bunch-Kaufman factorization of a symmetric matrix A using packed storage:

if $uplo='U'$, $A = PUDU^T P^T$
 if $uplo='L'$, $A = PLDL^T P^T$

where P is a permutation matrix, U and L are upper and lower triangular matrices with unit diagonal, and D is a symmetric block-diagonal matrix with 1-by-1 and 2-by-2 diagonal blocks. U and L have 2-by-2 unit diagonal blocks corresponding to the 2-by-2 blocks of D .

Input Parameters

$uplo$ CHARACTER*1. Must be 'U' or 'L'.

Indicates whether the upper or lower triangular part of A is packed in the array ap and how A is factored:

If $uplo = 'U'$, the array ap stores the upper triangular part of the matrix A , and A is factored as $PUDU^T P^T$.

If $uplo = 'L'$, the array ap stores the lower triangular part of the matrix A ; A is factored as $PLDL^T P^T$.

n INTEGER. The order of matrix A ($n \geq 0$).

ap REAL for *ssptf*
 DOUBLE PRECISION for *dsptf*
 COMPLEX for *csptf*
 DOUBLE COMPLEX for *zsptf*.
 Array, DIMENSION at least $\max(1, n(n+1)/2)$.
 The array *ap* contains either the upper or the lower triangular part of the matrix *A* (as specified by *uplo*) in packed storage (see [Matrix Storage Schemes](#)).

Output Parameters

ap The upper or lower triangle of *A* (as specified by *uplo*) is overwritten by details of the block-diagonal matrix *D* and the multipliers used to obtain the factor *U* (or *L*).

ipiv INTEGER.
 Array, DIMENSION at least $\max(1, n)$.
 Contains details of the interchanges and the block structure of *D*.
 If $ipiv(i) = k > 0$, then d_{ii} is a 1-by-1 block, and the *i*th row and column of *A* was interchanged with the *k*th row and column.
 If $uplo = 'U'$ and $ipiv(i) = ipiv(i-1) = -m < 0$, then *D* has a 2-by-2 block in rows/columns *i* and *i-1*, and (*i-1*)th row and column of *A* was interchanged with the *m*th row and column.
 If $uplo = 'L'$ and $ipiv(i) = ipiv(i+1) = -m < 0$, then *D* has a 2-by-2 block in rows/columns *i* and *i+1*, and (*i+1*)th row and column of *A* was interchanged with the *m*th row and column.

info INTEGER. If $info=0$, the execution is successful.
 If $info = -i$, the *i*th parameter had an illegal value.
 If $info = i$, d_{ii} is 0. The factorization has been completed, but *D* is exactly singular. Division by 0 will occur if you use *D* for solving a system of linear equations.

Application Notes

The 2-by-2 unit diagonal blocks and the unit diagonal elements of U and L are not stored. The remaining elements of U and L overwrite elements of the corresponding columns of the matrix A , but additional row interchanges are required to recover U or L explicitly (which is seldom necessary).

If $ipiv(i) = i$ for all $i = 1 \dots n$, then all off-diagonal elements of U (L) are stored explicitly in packed form.

If $uplo = 'U'$, the computed factors U and D are the exact factors of a perturbed matrix $A + E$, where

$$|E| \leq c(n)\epsilon P|U||D||U^T|P^T$$

$c(n)$ is a modest linear function of n , and ϵ is the machine precision.

A similar estimate holds for the computed L and D when $uplo = 'L'$.

The total number of floating-point operations is approximately $(1/3)n^3$ for real flavors or $(4/3)n^3$ for complex flavors.

After calling this routine, you can call the following:

- [?sptrs](#) to solve $AX = B$;
- [?spcon](#) to estimate the condition number of A ;
- [?sptri](#) to compute the inverse of A .

?hptrf

Computes the Bunch-Kaufman factorization of a complex Hermitian matrix using packed storage.

```
call chptrf ( uplo, n, ap, ipiv, info )
call zhptrf ( uplo, n, ap, ipiv, info )
```

Discussion

This routine forms the Bunch-Kaufman factorization of a Hermitian matrix using packed storage:

if `uplo='U'`, $A = PUDU^H P^T$
 if `uplo='L'`, $A = PLDL^H P^T$

where A is the input matrix, P is a permutation matrix, U and L are upper and lower triangular matrices with unit diagonal, and D is a Hermitian block-diagonal matrix with 1-by-1 and 2-by-2 diagonal blocks. U and L have 2-by-2 unit diagonal blocks corresponding to the 2-by-2 blocks of D .

Input Parameters

`uplo` CHARACTER*1. Must be 'U' or 'L'.
 Indicates whether the upper or lower triangular part of A is packed and how A is factored:
 If `uplo = 'U'`, the array `ap` stores the upper triangular part of the matrix A , and A is factored as $PUDU^H P^T$.
 If `uplo = 'L'`, the array `ap` stores the lower triangular part of the matrix A ; A is factored as $PLDL^H P^T$.

`n` INTEGER. The order of matrix A ($n \geq 0$).

`ap` COMPLEX for `chptrf`
 DOUBLE COMPLEX for `zhptrf`.
 Array, DIMENSION at least $\max(1, n(n+1)/2)$.

The array *ap* contains either the upper or the lower triangular part of the matrix *A* (as specified by *uplo*) in *packed storage* (see [Matrix Storage Schemes](#)).

Output Parameters

- ap* The upper or lower triangle of *A* (as specified by *uplo*) is overwritten by details of the block-diagonal matrix *D* and the multipliers used to obtain the factor *U* (or *L*).
- ipiv* **INTEGER.**
 Array, **DIMENSION** at least $\max(1,n)$.
 Contains details of the interchanges and the block structure of *D*.
 If $ipiv(i) = k > 0$, then d_{ii} is a 1-by-1 block, and the *i*th row and column of *A* was interchanged with the *k*th row and column.
 If $uplo = 'U'$ and $ipiv(i) = ipiv(i-1) = -m < 0$, then *D* has a 2-by-2 block in rows/columns *i* and *i-1*, and (*i-1*)th row and column of *A* was interchanged with the *m*th row and column.
 If $uplo = 'L'$ and $ipiv(i) = ipiv(i+1) = -m < 0$, then *D* has a 2-by-2 block in rows/columns *i* and *i+1*, and (*i+1*)th row and column of *A* was interchanged with the *m*th row and column.
- info* **INTEGER.** If $info=0$, the execution is successful.
 If $info = -i$, the *i*th parameter had an illegal value.
 If $info = i$, d_{ii} is 0. The factorization has been completed, but *D* is exactly singular. Division by 0 will occur if you use *D* for solving a system of linear equations.

Application Notes

The 2-by-2 unit diagonal blocks and the unit diagonal elements of U and L are not stored. The remaining elements of U and L are stored in the corresponding columns of the array a , but additional row interchanges are required to recover U or L explicitly (which is seldom necessary).

If $ipiv(i) = i$ for all $i = 1 \dots n$, then all off-diagonal elements of U (L) are stored explicitly in the corresponding elements of the array a .

If $uplo = 'U'$, the computed factors U and D are the exact factors of a perturbed matrix $A + E$, where

$$|E| \leq c(n) \varepsilon P|U||D||U^T|P^T$$

$c(n)$ is a modest linear function of n , and ε is the machine precision.

A similar estimate holds for the computed L and D when $uplo = 'L'$.

The total number of floating-point operations is approximately $(4/3)n^3$.

After calling this routine, you can call the following:

- [?hptrs](#) to solve $AX = B$;
- [?hpcon](#) to estimate the condition number of A ;
- [?hptri](#) to compute the inverse of A .

Routines for Solving Systems of Linear Equations

This section describes the LAPACK routines for solving systems of linear equations. Before calling most of these routines, you need to factorize the matrix of your system of equations (see [Routines for Matrix Factorization](#) in this chapter). However, the factorization is not necessary if your system of equations has a triangular matrix.

?getrs

Solves a system of linear equations with an LU-factored square matrix, with multiple right-hand sides.

```
call sgetrs (trans, n, nrhs, a, lda, ipiv, b, ldb, info)
call dgetrs (trans, n, nrhs, a, lda, ipiv, b, ldb, info)
call cgetrs (trans, n, nrhs, a, lda, ipiv, b, ldb, info)
call zgetrs (trans, n, nrhs, a, lda, ipiv, b, ldb, info)
```

Discussion

This routine solves for X the following systems of linear equations:

$AX = B$ if $trans = 'N'$,
 $A^T X = B$ if $trans = 'T'$,
 $A^H X = B$ if $trans = 'C'$ (for complex matrices only).

Before calling this routine, you must call [?getrf](#) to compute the LU factorization of A .

Input Parameters

trans CHARACTER*1. Must be 'N' or 'T' or 'C'.
 Indicates the form of the equations:
 If $trans = 'N'$, then $AX = B$ is solved for X .
 If $trans = 'T'$, then $A^T X = B$ is solved for X .
 If $trans = 'C'$, then $A^H X = B$ is solved for X .

n INTEGER. The order of A ; the number of rows in B ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides ($nrhs \geq 0$).

a, b REAL for *sgetrs*
 DOUBLE PRECISION for *dgetrs*
 COMPLEX for *cgetrs*
 DOUBLE COMPLEX for *zgetrs*.
 Arrays: $a(lda, *)$, $b(ldb, *)$.

The array *a* contains the matrix *A*.

The array *b* contains the matrix *B* whose columns are the right-hand sides for the systems of equations.

The second dimension of *a* must be at least $\max(1, n)$, the second dimension of *b* at least $\max(1, nrhs)$.

lda **INTEGER**. The first dimension of *a*; $lda \geq \max(1, n)$.
ldb **INTEGER**. The first dimension of *b*; $ldb \geq \max(1, n)$.
ipiv **INTEGER**.
 Array, **DIMENSION** at least $\max(1, n)$.
 The *ipiv* array, as returned by [?getrf](#).

Output Parameters

b Overwritten by the solution matrix *X*.
info **INTEGER**. If *info*=0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For each right-hand side *b*, the computed solution is the exact solution of a perturbed system of equations $(A + E)x = b$ where

$$|E| \leq c(n) \varepsilon P|L||U|$$

$c(n)$ is a modest linear function of *n*, and ε is the machine precision.

If x_0 is the true solution, the computed solution *x* satisfies this error bound:

$$\frac{\|x - x_0\|_\infty}{\|x\|_\infty} \leq c(n) \text{cond}(A, x) \varepsilon$$

where $\text{cond}(A, x) = \frac{\|A^{-1}\|_\infty \|A\|_\infty \|x\|_\infty}{\|x\|_\infty} \leq \|A^{-1}\|_\infty \|A\|_\infty = \kappa_\infty(A)$.

Note that $\text{cond}(A, x)$ can be much smaller than $\kappa_\infty(A)$; the condition number of A^T and A^H might or might not be equal to $\kappa_\infty(A)$.

The approximate number of floating-point operations for one right-hand side vector *b* is $2n^2$ for real flavors and $8n^2$ for complex flavors.

To estimate the condition number $\kappa_\infty(A)$, call [?gecon](#).

To refine the solution and estimate the error, call [?gerfs](#).

?gbtrs

Solves a system of linear equations with an LU-factored band matrix, with multiple right-hand sides.

```
call sgbtrs (trans, n, kl, ku, nrhs, ab, ldab, ipiv, b, ldb, info)
call dgbtrs (trans, n, kl, ku, nrhs, ab, ldab, ipiv, b, ldb, info)
call cgbtrs (trans, n, kl, ku, nrhs, ab, ldab, ipiv, b, ldb, info)
call zgbtrs (trans, n, kl, ku, nrhs, ab, ldab, ipiv, b, ldb, info)
```

Discussion

This routine solves for X the following systems of linear equations:

$AX = B$ if $trans = 'N'$,
 $A^T X = B$ if $trans = 'T'$,
 $A^H X = B$ if $trans = 'C'$ (for complex matrices only).

Here A is an LU-factored general band matrix of order n with kl non-zero sub-diagonals and ku non-zero super-diagonals. Before calling this routine, you must call [?gbtrf](#) to compute the LU factorization of A .

Input Parameters

$trans$ CHARACTER*1. Must be 'N' or 'T' or 'C'.
 n INTEGER. The order of A ; the number of rows in B ($n \geq 0$).
 kl INTEGER. The number of sub-diagonals within the band of A ($kl \geq 0$).
 ku INTEGER. The number of super-diagonals within the band of A ($ku \geq 0$).
 $nrhs$ INTEGER. The number of right-hand sides ($nrhs \geq 0$).
 ab, b REAL for sgbtrs
 DOUBLE PRECISION for dgbtrs
 COMPLEX for cgbtrs
 DOUBLE COMPLEX for zgbtrs.
 Arrays: $ab(ldab, *)$, $b(ldb, *)$.

The array *ab* contains the matrix *A* in *band storage* (see [Matrix Storage Schemes](#)).

The array *b* contains the matrix *B* whose columns are the right-hand sides for the systems of equations.

The second dimension of *ab* must be at least $\max(1, n)$, the second dimension of *b* at least $\max(1, nrhs)$.

ldab **INTEGER**. The first dimension of the array *ab*.
($ldab \geq 2kl + ku + 1$).

ldb **INTEGER**. The first dimension of *b*; $ldb \geq \max(1, n)$.

ipiv **INTEGER**. Array, **DIMENSION** at least $\max(1, n)$.
The *ipiv* array, as returned by [?gbtrf](#).

Output Parameters

b Overwritten by the solution matrix *X*.

info **INTEGER**. If *info*=0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For each right-hand side *b*, the computed solution is the exact solution of a perturbed system of equations $(A + E)x = b$, where

$$|E| \leq c(kl + ku + 1) \varepsilon P|L||U|$$

$c(k)$ is a modest linear function of *k*, and ε is the machine precision.

If x_0 is the true solution, the computed solution *x* satisfies this error bound:

$$\frac{\|x - x_0\|_\infty}{\|x\|_\infty} \leq c(kl + ku + 1) \text{cond}(A, x) \varepsilon$$

where $\text{cond}(A, x) = \frac{\|A^{-1}\|_\infty \|A\|_\infty \|x\|_\infty}{\|x\|_\infty} \leq \|A^{-1}\|_\infty \|A\|_\infty = \kappa_\infty(A)$.

Note that $\text{cond}(A, x)$ can be much smaller than $\kappa_\infty(A)$; the condition number of A^T and A^H might or might not be equal to $\kappa_\infty(A)$.

The approximate number of floating-point operations for one right-hand side vector is $2n(ku + 2kl)$ for real flavors. The number of operations for complex flavors is 4 times greater. All these estimates assume that *kl* and *ku* are much less than $\min(m, n)$.

To estimate the condition number $\kappa_\infty(A)$, call [?gbcon](#).

To refine the solution and estimate the error, call [?gbrfs](#).

?gttrs

Solves a system of linear equations with a tridiagonal matrix using the LU factorization computed by ?gttrf.

```
call sgttrs (trans, n, nrhs, dl, d, du, du2, ipiv, b, ldb, info)
call dgttrs (trans, n, nrhs, dl, d, du, du2, ipiv, b, ldb, info)
call cgttrs (trans, n, nrhs, dl, d, du, du2, ipiv, b, ldb, info)
call zgttrs (trans, n, nrhs, dl, d, du, du2, ipiv, b, ldb, info)
```

Discussion

This routine solves for X the following systems of linear equations with multiple right hand sides:

$$\begin{aligned}
 AX &= B && \text{if } trans = 'N', \\
 A^T X &= B && \text{if } trans = 'T', \\
 A^H X &= B && \text{if } trans = 'C' \text{ (for complex matrices only).}
 \end{aligned}$$

Before calling this routine, you must call ?gttrf to compute the LU factorization of A .

Input Parameters

trans CHARACTER*1. Must be 'N' or 'T' or 'C'.
Indicates the form of the equations:
If *trans* = 'N', then $AX = B$ is solved for X .
If *trans* = 'T', then $A^T X = B$ is solved for X .
If *trans* = 'C', then $A^H X = B$ is solved for X .

n INTEGER. The order of A ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides, i.e., the number of columns in B ($nrhs \geq 0$).

dl,d,du,du2,b REAL for sgttrs
DOUBLE PRECISION for dgttrs
COMPLEX for cgttrs
DOUBLE COMPLEX for zgttrf.

Arrays: $d1(n-1)$, $d(n)$, $du(n-1)$, $du2(n-2)$,
 $b(l\delta b, nrhs)$.

The array $d1$ contains the $(n-1)$ multipliers that define the matrix L from the LU factorization of A .

The array d contains the n diagonal elements of the upper triangular matrix U from the LU factorization of A .

The array du contains the $(n-1)$ elements of the first super-diagonal of U .

The array $du2$ contains the $(n-2)$ elements of the second super-diagonal of U .

The array b contains the matrix B whose columns are the right-hand sides for the systems of equations.

$l\delta b$ INTEGER. The leading dimension of b ; $l\delta b \geq \max(1, n)$.

$ipiv$ INTEGER.
 Array, DIMENSION (n) .
 The $ipiv$ array, as returned by [?gttrf](#).

Output Parameters

b Overwritten by the solution matrix X .

$info$ INTEGER. If $info=0$, the execution is successful.
 If $info = -i$, the i th parameter had an illegal value.

Application Notes

For each right-hand side b , the computed solution is the exact solution of a perturbed system of equations $(A + E)x = b$ where

$$|E| \leq c(n)\epsilon P|L||U|$$

$c(n)$ is a modest linear function of n , and ϵ is the machine precision.

If x_0 is the true solution, the computed solution x satisfies this error bound:

$$\frac{\|x - x_0\|_\infty}{\|x\|_\infty} \leq c(n) \text{cond}(A, x)\epsilon$$

where $\text{cond}(A, x) = \frac{\| |A^{-1}| |A| |x| \|_\infty}{\|x\|_\infty} \leq \| |A^{-1}| \|_\infty \| |A| \|_\infty = \kappa_\infty(A)$.

Note that $\text{cond}(A,x)$ can be much smaller than $\kappa_\infty(A)$; the condition number of A^T and A^H might or might not be equal to $\kappa_\infty(A)$.

The approximate number of floating-point operations for one right-hand side vector b is $2n^2$ for real flavors and $8n^2$ for complex flavors.

To estimate the condition number $\kappa_\infty(A)$, call [?gecon](#).

To refine the solution and estimate the error, call [?gerfs](#).

?potrs

Solves a system of linear equations with a Cholesky-factored symmetric (Hermitian) positive-definite matrix.

```
call spotrs ( uplo, n, nrhs, a, lda, b, ldb, info )
call dpotrs ( uplo, n, nrhs, a, lda, b, ldb, info )
call cpotrs ( uplo, n, nrhs, a, lda, b, ldb, info )
call zpotrs ( uplo, n, nrhs, a, lda, b, ldb, info )
```

Discussion

This routine solves for X the system of linear equations $AX = B$ with a symmetric positive-definite or, for complex data, Hermitian positive-definite matrix A , given the Cholesky factorization of A :

$$A = U^H U \quad \text{if } uplo = 'U'$$

$$A = LL^H \quad \text{if } uplo = 'L'$$

where L is a lower triangular matrix and U is upper triangular. The system is solved with multiple right-hand sides stored in the columns of the matrix B .

Before calling this routine, you must call [?potrf](#) to compute the Cholesky factorization of A .

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If $uplo = 'U'$, the array a stores the factor U of the Cholesky factorization $A = U^H U$.
If $uplo = 'L'$, the array a stores the factor L of the Cholesky factorization $A = LL^H$.

n INTEGER. The order of matrix A ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides ($nrhs \geq 0$).

a, b REAL for `spotrs`
 DOUBLE PRECISION for `dpotrs`
 COMPLEX for `cpotrs`
 DOUBLE COMPLEX for `zpotrs`.
 Arrays: `a(lda,*)`, `b(ldb,*)`.
 The array `a` contains the factor U or L (see `uplo`).
 The array `b` contains the matrix B whose columns are the right-hand sides for the systems of equations.
 The second dimension of `a` must be at least $\max(1, n)$, the second dimension of `b` at least $\max(1, nrhs)$.

lda INTEGER. The first dimension of `a`; $lda \geq \max(1, n)$.

ldb INTEGER. The first dimension of `b`; $ldb \geq \max(1, n)$.

Output Parameters

b Overwritten by the solution matrix X .

info INTEGER. If `info` = 0, the execution is successful.
 If `info` = $-i$, the i th parameter had an illegal value.

Application Notes

If `uplo` = 'U', the computed solution for each right-hand side b is the exact solution of a perturbed system of equations $(A + E)x = b$, where

$$|E| \leq c(n)\varepsilon |U^H| |U|$$

$c(n)$ is a modest linear function of n , and ε is the machine precision.

A similar estimate holds for `uplo` = 'L'.

If x_0 is the true solution, the computed solution x satisfies this error bound:

$$\frac{\|x - x_0\|_\infty}{\|x\|_\infty} \leq c(n) \text{cond}(A, x) \varepsilon$$

where $\text{cond}(A, x) = \| |A^{-1}| |A| |x| \|_\infty / \|x\|_\infty \leq \|A^{-1}\|_\infty \|A\|_\infty = \kappa_\infty(A)$.

Note that $\text{cond}(A, x)$ can be much smaller than $\kappa_\infty(A)$.

The approximate number of floating-point operations for one right-hand side vector b is $2n^2$ for real flavors and $8n^2$ for complex flavors.

To estimate the condition number $\kappa_\infty(A)$, call [?pocon](#).

To refine the solution and estimate the error, call [?porfs](#).

?pptrs

Solves a system of linear equations with a packed Cholesky-factored symmetric (Hermitian) positive-definite matrix.

```
call spptrs ( uplo, n, nrhs, ap, b, ldb, info )
call dpptrs ( uplo, n, nrhs, ap, b, ldb, info )
call cpptrs ( uplo, n, nrhs, ap, b, ldb, info )
call zpptrs ( uplo, n, nrhs, ap, b, ldb, info )
```

Discussion

This routine solves for X the system of linear equations $AX = B$ with a packed symmetric positive-definite or, for complex data, Hermitian positive-definite matrix A , given the Cholesky factorization of A :

$$A = U^H U \quad \text{if } uplo = 'U'$$

$$A = LL^H \quad \text{if } uplo = 'L'$$

where L is a lower triangular matrix and U is upper triangular. The system is solved with multiple right-hand sides stored in the columns of the matrix B .

Before calling this routine, you must call [?pptrf](#) to compute the Cholesky factorization of A .

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If $uplo = 'U'$, the array a stores the packed factor U of the Cholesky factorization $A = U^H U$.
If $uplo = 'L'$, the array a stores the packed factor L of the Cholesky factorization $A = LL^H$.

n INTEGER. The order of matrix A ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides ($nrhs \geq 0$).

ap, b REAL for *spptrs*
 DOUBLE PRECISION for *dpptrs*
 COMPLEX for *cpptrs*
 DOUBLE COMPLEX for *zpptrs*.
 Arrays: *ap(*), b(ldb,*)*
 The dimension of *ap* must be at least $\max(1, n(n+1)/2)$.
 The array *ap* contains the factor *U* or *L*, as specified by *uplo*, in *packed storage* (see [Matrix Storage Schemes](#)).
 The array *b* contains the matrix *B* whose columns are the right-hand sides for the systems of equations. The second dimension of *b* must be at least $\max(1, nrhs)$.

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

Output Parameters

b Overwritten by the solution matrix *X*.
info INTEGER. If *info* = 0, the execution is successful.
 If *info* = *-i*, the *i*th parameter had an illegal value.

Application Notes

If *uplo* = 'U', the computed solution for each right-hand side *b* is the exact solution of a perturbed system of equations $(A + E)x = b$, where

$$|E| \leq c(n) \varepsilon |U^H| |U|$$

c(n) is a modest linear function of *n*, and ε is the machine precision.

A similar estimate holds for *uplo* = 'L'.

If *x*₀ is the true solution, the computed solution *x* satisfies this error bound:

$$\frac{\|x - x_0\|_\infty}{\|x\|_\infty} \leq c(n) \text{cond}(A, x) \varepsilon$$

where $\text{cond}(A, x) = \frac{\|A^{-1}\|_\infty \|A\|_\infty \|x\|_\infty}{\|x\|_\infty} \leq \|A^{-1}\|_\infty \|A\|_\infty = \kappa_\infty(A)$.

Note that $\text{cond}(A, x)$ can be much smaller than $\kappa_\infty(A)$.

The approximate number of floating-point operations for one right-hand side vector *b* is $2n^2$ for real flavors and $8n^2$ for complex flavors.

To estimate the condition number $\kappa_\infty(A)$, call [?ppcon](#).

To refine the solution and estimate the error, call [?pprfs](#).

?pbtrs

Solves a system of linear equations with a Cholesky-factored symmetric (Hermitian) positive-definite band matrix.

```
call spbtrs (uplo, n, kd, nrhs, ab, ldab, b, ldb, info)
call dpbtrs (uplo, n, kd, nrhs, ab, ldab, b, ldb, info)
call cpbtrs (uplo, n, kd, nrhs, ab, ldab, b, ldb, info)
call zpbtrs (uplo, n, kd, nrhs, ab, ldab, b, ldb, info)
```

Discussion

This routine solves for X the system of linear equations $AX = B$ with a symmetric positive-definite or, for complex data, Hermitian positive-definite **band** matrix A , given the Cholesky factorization of A :

$$\begin{aligned} A &= U^H U && \text{if } uplo = 'U' \\ A &= L L^H && \text{if } uplo = 'L' \end{aligned}$$

where L is a lower triangular matrix and U is upper triangular. The system is solved with multiple right-hand sides stored in the columns of the matrix B .

Before calling this routine, you must call [?pbtrf](#) to compute the Cholesky factorization of A in the band storage form.

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If *uplo* = 'U', the array *a* stores the factor U of the factorization $A = U^H U$ in the band storage form.
If *uplo* = 'L', the array *a* stores the factor L of the factorization $A = L L^H$ in the band storage form.

n INTEGER. The order of matrix A ($n \geq 0$).

kd INTEGER. The number of super-diagonals or sub-diagonals in the matrix A ($kd \geq 0$).

nrhs INTEGER. The number of right-hand sides ($nrhs \geq 0$).

ab, *b* REAL for *spbtrs*
 DOUBLE PRECISION for *dpbtrs*
 COMPLEX for *cpbtrs*
 DOUBLE COMPLEX for *zpbtrs*.
 Arrays: *ab(ldab,*)*, *b(ldb,*)*.
 The array *ab* contains the Cholesky factor, as returned by the factorization routine, in *band storage* form.
 The array *b* contains the matrix *B* whose columns are the right-hand sides for the systems of equations.
 The second dimension of *ab* must be at least $\max(1, n)$, the second dimension of *b* at least $\max(1, nrhs)$.

ldab INTEGER. The first dimension of the array *ab*.
 ($ldab \geq kd + 1$).

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

Output Parameters

b Overwritten by the solution matrix *X*.
info INTEGER. If *info*=0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For each right-hand side *b*, the computed solution is the exact solution of a perturbed system of equations $(A + E)x = b$, where

$$|E| \leq c(kd + 1)\epsilon P|U^H||U| \quad \text{or} \quad |E| \leq c(kd + 1)\epsilon P|L^H||L|$$

$c(k)$ is a modest linear function of k , and ϵ is the machine precision.

If x_0 is the true solution, the computed solution x satisfies this error bound:

$$\frac{\|x - x_0\|_\infty}{\|x\|_\infty} \leq c(kd + 1) \text{cond}(A, x)\epsilon$$

where $\text{cond}(A, x) = \frac{\|A^{-1}\|_\infty \|A\|_\infty \|x\|_\infty}{\|x\|_\infty} \leq \|A^{-1}\|_\infty \|A\|_\infty = \kappa_\infty(A)$.

Note that $\text{cond}(A, x)$ can be much smaller than $\kappa_\infty(A)$.

The approximate number of floating-point operations for one right-hand side vector is $4n*kd$ for real flavors and $16n*kd$ for complex flavors.

To estimate the condition number $\kappa_\infty(A)$, call [?pbcon](#).

To refine the solution and estimate the error, call [?pbrfs](#).

?pttrs

Solves a system of linear equations with a symmetric (Hermitian) positive-definite tridiagonal matrix using the factorization computed by `?pttrf`.

```
call spttrs (n, nrhs, d, e, b, ldb, info)
call dpttrs (n, nrhs, d, e, b, ldb, info)
call cpttrs (uplo, n, nrhs, d, e, b, ldb, info)
call zpttrs (uplo, n, nrhs, d, e, b, ldb, info)
```

Discussion

This routine solves for X a system of linear equations $AX = B$ with a symmetric (Hermitian) positive-definite tridiagonal matrix A .

Before calling this routine, you must call `?pttrf` to compute the LDL^H or $U^H DU$ factorization of A .

Input Parameters

`uplo` CHARACTER*1. Used for `cpttrs/zpttrs` only. Must be 'U' or 'L'. Specifies whether the superdiagonal or the subdiagonal of the tridiagonal matrix A is stored and how A is factored:
If `uplo = 'U'`, the array `e` stores the superdiagonal of A , and A is factored as $U^H DU$;
If `uplo = 'L'`, the array `e` stores the subdiagonal of A , and A is factored as LDL^H .

`n` INTEGER. The order of A ($n \geq 0$).

`nrhs` INTEGER. The number of right-hand sides, i.e., the number of columns of the matrix B ($nrhs \geq 0$).

<i>d</i>	<p>REAL for <i>spttrs</i>, <i>cpttrs</i> DOUBLE PRECISION for <i>dpttrs</i>, <i>zpttrs</i>. Array, dimension (<i>n</i>). Contains the diagonal elements of the diagonal matrix <i>D</i> from the factorization computed by ?pttrf.</p>
<i>e</i> , <i>b</i>	<p>REAL for <i>spttrs</i> DOUBLE PRECISION for <i>dpttrs</i> COMPLEX for <i>cpttrs</i> DOUBLE COMPLEX for <i>zpttrs</i>. Arrays: <i>e</i>(<i>n</i> - 1), <i>b</i>(<i>ldb</i>, <i>nrhs</i>). The array <i>e</i> contains the (<i>n</i> - 1) off-diagonal elements of the unit bidiagonal factor <i>U</i> or <i>L</i> from the factorization computed by ?pttrf (see <i>uplo</i>). The array <i>b</i> contains the matrix <i>B</i> whose columns are the right-hand sides for the systems of equations.</p>
<i>ldb</i>	<p>INTEGER. The leading dimension of <i>b</i>; <i>ldb</i> ≥ max(1, <i>n</i>).</p>

Output Parameters

<i>b</i>	Overwritten by the solution matrix <i>X</i> .
<i>info</i>	<p>INTEGER. If <i>info</i>=0, the execution is successful. If <i>info</i> = -<i>i</i>, the <i>i</i>th parameter had an illegal value.</p>

?sytrs

Solves a system of linear equations with a UDU- or LDL-factored symmetric matrix.

```
call ssytrs (uplo, n, nrhs, a, lda, ipiv, b, ldb, info)
call dsytrs (uplo, n, nrhs, a, lda, ipiv, b, ldb, info)
call csytrs (uplo, n, nrhs, a, lda, ipiv, b, ldb, info)
call zsytrs (uplo, n, nrhs, a, lda, ipiv, b, ldb, info)
```

Discussion

This routine solves for X the system of linear equations $AX = B$ with a symmetric matrix A , given the Bunch-Kaufman factorization of A :

if $uplo='U'$, $A = PUDU^T P^T$
if $uplo='L'$, $A = PLDL^T P^T$

where P is a permutation matrix, U and L are upper and lower triangular matrices with unit diagonal, and D is a symmetric block-diagonal matrix. The system is solved with multiple right-hand sides stored in the columns of the matrix B . You must supply to this routine the factor U (or L) and the array $ipiv$ returned by the factorization routine [?sytrf](#).

Input Parameters

$uplo$ CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If $uplo = 'U'$, the array a stores the upper triangular factor U of the factorization $A = PUDU^T P^T$.
If $uplo = 'L'$, the array a stores the lower triangular factor L of the factorization $A = PLDL^T P^T$.

n INTEGER. The order of matrix A ($n \geq 0$).

$nrhs$ INTEGER. The number of right-hand sides ($nrhs \geq 0$).

$ipiv$ INTEGER. Array, DIMENSION at least $\max(1,n)$.
The $ipiv$ array, as returned by [?sytrf](#).

a, *b* REAL for `ssytrs`
 DOUBLE PRECISION for `dsytrs`
 COMPLEX for `csytrs`
 DOUBLE COMPLEX for `zsytrs`.
 Arrays: *a*(*lda*, *), *b*(*ldb*, *).
 The array *a* contains the factor *U* or *L* (see `uplo`).
 The array *b* contains the matrix *B* whose columns are the right-hand sides for the system of equations.
 The second dimension of *a* must be at least $\max(1, n)$,
 the second dimension of *b* at least $\max(1, nrhs)$.

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, n)$.
ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

Output Parameters

b Overwritten by the solution matrix *X*.
info INTEGER. If *info*=0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For each right-hand side *b*, the computed solution is the exact solution of a perturbed system of equations $(A + E)x = b$, where

$$|E| \leq c(n)\varepsilon P|U||D||U^T|P^T \quad \text{or} \quad |E| \leq c(n)\varepsilon P|L||D||L^T|P^T$$

$c(n)$ is a modest linear function of *n*, and ε is the machine precision.

If x_0 is the true solution, the computed solution *x* satisfies this error bound:

$$\frac{\|x - x_0\|_\infty}{\|x\|_\infty} \leq c(n) \text{cond}(A, x)\varepsilon$$

where $\text{cond}(A, x) = \frac{\|A^{-1}\|_\infty \|A\|_\infty \|x\|_\infty}{\|x\|_\infty} \leq \|A^{-1}\|_\infty \|A\|_\infty = \kappa_\infty(A)$.

Note that $\text{cond}(A, x)$ can be much smaller than $\kappa_\infty(A)$.

The total number of floating-point operations for one right-hand side vector is approximately $2n^2$ for real flavors or $8n^2$ for complex flavors.

To estimate the condition number $\kappa_\infty(A)$, call [?sycon](#).

To refine the solution and estimate the error, call [?sytrfs](#).

?hetrs

Solves a system of linear equations with a UDU- or LDL-factored Hermitian matrix.

```
call chetrs (uplo, n, nrhs, a, lda, ipiv, b, ldb, info)
call zhetrs (uplo, n, nrhs, a, lda, ipiv, b, ldb, info)
```

Discussion

This routine solves for X the system of linear equations $AX = B$ with a Hermitian matrix A , given the Bunch-Kaufman factorization of A :

$$\text{if } uplo = 'U', \quad A = PUDU^H P^T$$

$$\text{if } uplo = 'L', \quad A = PLDL^H P^T$$

where P is a permutation matrix, U and L are upper and lower triangular matrices with unit diagonal, and D is a symmetric block-diagonal matrix. The system is solved with multiple right-hand sides stored in the columns of the matrix B . You must supply to this routine the factor U (or L) and the array $ipiv$ returned by the factorization routine [?hetrf](#).

Input Parameters

<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. Indicates how the input matrix A has been factored: If $uplo = 'U'$, the array a stores the upper triangular factor U of the factorization $A = PUDU^H P^T$. If $uplo = 'L'$, the array a stores the lower triangular factor L of the factorization $A = PLDL^H P^T$.
<i>n</i>	INTEGER. The order of matrix A ($n \geq 0$).
<i>nrhs</i>	INTEGER. The number of right-hand sides ($nrhs \geq 0$).
<i>ipiv</i>	INTEGER. Array, DIMENSION at least $\max(1, n)$. The $ipiv$ array, as returned by ?hetrf .

a, *b* **COMPLEX** for **chetrs**.
DOUBLE COMPLEX for **zhetsr**.
 Arrays: *a*(*lda*, *), *b*(*ldb*, *).
 The array *a* contains the factor *U* or *L* (see **uplo**).
 The array *b* contains the matrix *B* whose columns are the right-hand sides for the system of equations.
 The second dimension of *a* must be at least $\max(1, n)$,
 the second dimension of *b* at least $\max(1, nrhs)$.

lda **INTEGER**. The first dimension of *a*; $lda \geq \max(1, n)$.

ldb **INTEGER**. The first dimension of *b*; $ldb \geq \max(1, n)$.

Output Parameters

b Overwritten by the solution matrix *X*.

info **INTEGER**. If *info*=0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For each right-hand side *b*, the computed solution is the exact solution of a perturbed system of equations $(A + E)x = b$, where

$$|E| \leq c(n)\varepsilon P|U||D||U^H|P^T \quad \text{or} \quad |E| \leq c(n)\varepsilon P|L||D||L^H|P^T$$

$c(n)$ is a modest linear function of *n*, and ε is the machine precision.

If x_0 is the true solution, the computed solution *x* satisfies this error bound:

$$\frac{\|x - x_0\|_\infty}{\|x\|_\infty} \leq c(n) \text{cond}(A, x)\varepsilon$$

where $\text{cond}(A, x) = \| |A^{-1}| |A| |x| \|_\infty / \|x\|_\infty \leq \|A^{-1}\|_\infty \|A\|_\infty = \kappa_\infty(A)$.

Note that $\text{cond}(A, x)$ can be much smaller than $\kappa_\infty(A)$.

The total number of floating-point operations for one right-hand side vector is approximately $8n^2$.

To estimate the condition number $\kappa_\infty(A)$, call [?hecon](#).

To refine the solution and estimate the error, call [?herfs](#).

?spttrs

Solves a system of linear equations with a UDU- or LDL-factored symmetric matrix using packed storage.

```
call sspttrs ( uplo, n, nrhs, ap, ipiv, b, ldb, info )
call dspttrs ( uplo, n, nrhs, ap, ipiv, b, ldb, info )
call cspttrs ( uplo, n, nrhs, ap, ipiv, b, ldb, info )
call zsptrs ( uplo, n, nrhs, ap, ipiv, b, ldb, info )
```

Discussion

This routine solves for X the system of linear equations $AX = B$ with a symmetric matrix A , given the Bunch-Kaufman factorization of A :

$$\begin{aligned} \text{if } \text{uplo} = \text{'U'}, & \quad A = PUDU^T P^T \\ \text{if } \text{uplo} = \text{'L'}, & \quad A = PLDL^T P^T \end{aligned}$$

where P is a permutation matrix, U and L are upper and lower **packed** triangular matrices with unit diagonal, and D is a symmetric block-diagonal matrix. The system is solved with multiple right-hand sides stored in the columns of the matrix B . You must supply the factor U (or L) and the array $ipiv$ returned by the factorization routine [?sptrf](#).

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If $uplo = \text{'U'}$, the array ap stores the packed factor U of the factorization $A = PUDU^T P^T$.
If $uplo = \text{'L'}$, the array ap stores the packed factor L of the factorization $A = PLDL^T P^T$.

n INTEGER. The order of matrix A ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides ($nrhs \geq 0$).

ipiv INTEGER. Array, DIMENSION at least $\max(1, n)$.
The $ipiv$ array, as returned by [?sptrf](#).

ap, *b* REAL for *ssptrs*
 DOUBLE PRECISION for *dsptrs*
 COMPLEX for *csptrs*
 DOUBLE COMPLEX for *zsptrs*.
 Arrays: *ap*(*), *b*(*ldb*,*)
 The dimension of *ap* must be at least $\max(1, n(n+1)/2)$.
 The array *ap* contains the factor *U* or *L*, as specified by *uplo*, in *packed storage* (see [Matrix Storage Schemes](#)).
 The array *b* contains the matrix *B* whose columns are the right-hand sides for the system of equations. The second dimension of *b* must be at least $\max(1, nrhs)$.

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

Output Parameters

b Overwritten by the solution matrix *X*.
info INTEGER. If *info*=0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For each right-hand side *b*, the computed solution is the exact solution of a perturbed system of equations $(A + E)x = b$, where

$$|E| \leq c(n)\varepsilon P|U||D||U^T|P^T \quad \text{or} \quad |E| \leq c(n)\varepsilon P|L||D||L^T|P^T$$

$c(n)$ is a modest linear function of *n*, and ε is the machine precision.

If x_0 is the true solution, the computed solution *x* satisfies this error bound:

$$\frac{\|x - x_0\|_\infty}{\|x\|_\infty} \leq c(n) \text{cond}(A, x)\varepsilon$$

where $\text{cond}(A, x) = \frac{\|A^{-1}\|_\infty \|A\|_\infty \|x\|_\infty}{\|x\|_\infty} \leq \|A^{-1}\|_\infty \|A\|_\infty = \kappa_\infty(A)$.

Note that $\text{cond}(A, x)$ can be much smaller than $\kappa_\infty(A)$.

The total number of floating-point operations for one right-hand side vector is approximately $2n^2$ for real flavors or $8n^2$ for complex flavors.

To estimate the condition number $\kappa_\infty(A)$, call [?spcon](#).

To refine the solution and estimate the error, call [?sprfs](#).

?hptrs

Solves a system of linear equations with a UDU- or LDL-factored Hermitian matrix using packed storage.

```
call chptrs ( uplo, n, nrhs, ap, ipiv, b, ldb, info )
call zhptrs ( uplo, n, nrhs, ap, ipiv, b, ldb, info )
```

Discussion

This routine solves for X the system of linear equations $AX = B$ with a Hermitian matrix A , given the Bunch-Kaufman factorization of A :

if $uplo='U'$, $A = PUDU^H P^T$
 if $uplo='L'$, $A = PLDL^H P^T$

where P is a permutation matrix, U and L are upper and lower *packed* triangular matrices with unit diagonal, and D is a symmetric block-diagonal matrix. The system is solved with multiple right-hand sides stored in the columns of the matrix B .

You must supply to this routine the arrays ap (containing U or L) and $ipiv$ in the form returned by the factorization routine [?hptrf](#).

Input Parameters

$uplo$ CHARACTER*1. Must be 'U' or 'L'.
 Indicates how the input matrix A has been factored:
 If $uplo = 'U'$, the array ap stores the packed factor U of the factorization $A = PUDU^H P^T$.
 If $uplo = 'L'$, the array ap stores the packed factor L of the factorization $A = PLDL^H P^T$.

n INTEGER. The order of matrix A ($n \geq 0$).

$nrhs$ INTEGER. The number of right-hand sides ($nrhs \geq 0$).

$ipiv$ INTEGER. Array, DIMENSION at least $\max(1,n)$.
 The $ipiv$ array, as returned by [?hptrf](#).

ap, *b* **COMPLEX** for **chptrs**.
DOUBLE COMPLEX for **zhptrs**.
 Arrays: *ap*(*), *b*(*ldb*, *)
 The dimension of *ap* must be at least $\max(1, n(n+1)/2)$.
 The array *ap* contains the factor *U* or *L*, as specified by *uplo*, in *packed storage* (see [Matrix Storage Schemes](#)).
 The array *b* contains the matrix *B* whose columns are the right-hand sides for the system of equations. The second dimension of *b* must be at least $\max(1, nrhs)$.

ldb **INTEGER**. The first dimension of *b*; $ldb \geq \max(1, n)$.

Output Parameters

b Overwritten by the solution matrix *X*.
info **INTEGER**. If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For each right-hand side *b*, the computed solution is the exact solution of a perturbed system of equations $(A + E)x = b$, where

$$|E| \leq c(n)\varepsilon P|U||D||U^H|P^T \quad \text{or} \quad |E| \leq c(n)\varepsilon P|L||D||L^H|P^T$$

$c(n)$ is a modest linear function of n , and ε is the machine precision.

If x_0 is the true solution, the computed solution x satisfies this error bound:

$$\frac{\|x - x_0\|_\infty}{\|x\|_\infty} \leq c(n) \text{cond}(A, x)\varepsilon$$

where $\text{cond}(A, x) = \frac{\|A^{-1}\|_\infty \|A\|_\infty \|x\|_\infty}{\|x\|_\infty} \leq \|A^{-1}\|_\infty \|A\|_\infty = \kappa_\infty(A)$.

Note that $\text{cond}(A, x)$ can be much smaller than $\kappa_\infty(A)$.

The total number of floating-point operations for one right-hand side vector is approximately $8n^2$ for complex flavors.

To estimate the condition number $\kappa_\infty(A)$, call [?hpcon](#).

To refine the solution and estimate the error, call [?hprfs](#).

?trtrs

Solves a system of linear equations with a triangular matrix, with multiple right-hand sides.

```
call strtrs (uplo,trans,diag,n,nrhs,a,lda,b,ldb,info)
call dtrtrs (uplo,trans,diag,n,nrhs,a,lda,b,ldb,info)
call ctrtrs (uplo,trans,diag,n,nrhs,a,lda,b,ldb,info)
call ztrtrs (uplo,trans,diag,n,nrhs,a,lda,b,ldb,info)
```

Discussion

This routine solves for X the following systems of linear equations with a triangular matrix A , with multiple right-hand sides stored in B :

$AX = B$ if $trans = 'N'$,
 $A^T X = B$ if $trans = 'T'$,
 $A^H X = B$ if $trans = 'C'$ (for complex matrices only).

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
 Indicates whether A is upper or lower triangular:
 If $uplo = 'U'$, then A is upper triangular.
 If $uplo = 'L'$, then A is lower triangular.

trans CHARACTER*1. Must be 'N' or 'T' or 'C'.
 If $trans = 'N'$, then $AX = B$ is solved for X .
 If $trans = 'T'$, then $A^T X = B$ is solved for X .
 If $trans = 'C'$, then $A^H X = B$ is solved for X .

diag CHARACTER*1. Must be 'N' or 'U'.
 If $diag = 'N'$, then A is not a unit triangular matrix.
 If $diag = 'U'$, then A is unit triangular: diagonal elements of A are assumed to be 1 and not referenced in the array a .

n INTEGER. The order of A ; the number of rows in B ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides ($nrhs \geq 0$).

a, b REAL for `strtrs`
 DOUBLE PRECISION for `dtrtrs`
 COMPLEX for `ctrtrs`
 DOUBLE COMPLEX for `ztrtrs`.
 Arrays: `a(lda,*)`, `b(ldb,*)`.
 The array *a* contains the matrix *A*.
 The array *b* contains the matrix *B* whose columns are the right-hand sides for the systems of equations.
 The second dimension of *a* must be at least $\max(1, n)$, the second dimension of *b* at least $\max(1, nrhs)$.

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, n)$.

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

Output Parameters

b Overwritten by the solution matrix *X*.

info INTEGER. If *info*=0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For each right-hand side *b*, the computed solution is the exact solution of a perturbed system of equations $(A + E)x = b$ where

$$|E| \leq c(n) \varepsilon |A|$$

$c(n)$ is a modest linear function of *n*, and ε is the machine precision.

If x_0 is the true solution, the computed solution *x* satisfies this error bound:

$$\frac{\|x - x_0\|_\infty}{\|x\|_\infty} \leq c(n) \text{cond}(A, x) \varepsilon, \text{ provided } c(n) \text{cond}(A, x) \varepsilon < 1$$

where $\text{cond}(A, x) = \frac{\|A^{-1}\|_\infty \|A\|_\infty \|x\|_\infty}{\|x\|_\infty} \leq \|A^{-1}\|_\infty \|A\|_\infty = \kappa_\infty(A)$.

Note that $\text{cond}(A, x)$ can be much smaller than $\kappa_\infty(A)$; the condition number of A^T and A^H might or might not be equal to $\kappa_\infty(A)$.

The approximate number of floating-point operations for one right-hand side vector *b* is n^2 for real flavors and $4n^2$ for complex flavors.

To estimate the condition number $\kappa_\infty(A)$, call [?trcon](#).

To estimate the error in the solution, call [?trrfs](#).

?tptrs

Solves a system of linear equations with a packed triangular matrix, with multiple right-hand sides.

```
call stptrs (uplo,trans,diag,n,nrhs,ap,b,ldb,info)
call dtptrs (uplo,trans,diag,n,nrhs,ap,b,ldb,info)
call ctptrs (uplo,trans,diag,n,nrhs,ap,b,ldb,info)
call ztptrs (uplo,trans,diag,n,nrhs,ap,b,ldb,info)
```

Discussion

This routine solves for X the following systems of linear equations with a packed triangular matrix A , with multiple right-hand sides stored in B :

$AX = B$ if $trans = 'N'$,
 $A^T X = B$ if $trans = 'T'$,
 $A^H X = B$ if $trans = 'C'$ (for complex matrices only).

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
 Indicates whether A is upper or lower triangular:
 If $uplo = 'U'$, then A is upper triangular.
 If $uplo = 'L'$, then A is lower triangular.

trans CHARACTER*1. Must be 'N' or 'T' or 'C'.
 If $trans = 'N'$, then $AX = B$ is solved for X .
 If $trans = 'T'$, then $A^T X = B$ is solved for X .
 If $trans = 'C'$, then $A^H X = B$ is solved for X .

diag CHARACTER*1. Must be 'N' or 'U'.
 If $diag = 'N'$, then A is not a unit triangular matrix.
 If $diag = 'U'$, then A is unit triangular: diagonal elements are assumed to be 1 and not referenced in the array ap .

n INTEGER. The order of A ; the number of rows in B ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides ($nrhs \geq 0$).

ap, b REAL for *stptrs*
 DOUBLE PRECISION for *dtptrs*
 COMPLEX for *ctptrs*
 DOUBLE COMPLEX for *ztptrs*.
 Arrays: *ap(*), b(ldb,*)*
 The dimension of *ap* must be at least $\max(1, n(n+1)/2)$.
 The array *ap* contains the matrix *A* in *packed storage* (see [Matrix Storage Schemes](#)).
 The array *b* contains the matrix *B* whose columns are the right-hand sides for the system of equations. The second dimension of *b* must be at least $\max(1, nrhs)$.

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

Output Parameters

b Overwritten by the solution matrix *X*.

info INTEGER. If *info*=0, the execution is successful.
 If *info* = *-i*, the *i*th parameter had an illegal value.

Application Notes

For each right-hand side *b*, the computed solution is the exact solution of a perturbed system of equations $(A + E)x = b$ where

$$|E| \leq c(n) \varepsilon |A|$$

$c(n)$ is a modest linear function of *n*, and ε is the machine precision.

If x_0 is the true solution, the computed solution *x* satisfies this error bound:

$$\frac{\|x - x_0\|_\infty}{\|x\|_\infty} \leq c(n) \text{cond}(A, x) \varepsilon, \text{ provided } c(n) \text{cond}(A, x) \varepsilon < 1$$

where $\text{cond}(A, x) = \frac{\|A^{-1}\|_\infty \|A\|_\infty \|x\|_\infty}{\|x\|_\infty} \leq \|A^{-1}\|_\infty \|A\|_\infty = \kappa_\infty(A)$.

Note that $\text{cond}(A, x)$ can be much smaller than $\kappa_\infty(A)$; the condition number of A^T and A^H might or might not be equal to $\kappa_\infty(A)$.

The approximate number of floating-point operations for one right-hand side vector *b* is n^2 for real flavors and $4n^2$ for complex flavors.

To estimate the condition number $\kappa_\infty(A)$, call [?tpcon](#).

To estimate the error in the solution, call [?tprfs](#).

?tbtrs

Solves a system of linear equations with a band triangular matrix, with multiple right-hand sides.

```
call stbtrs (uplo, trans, diag, n, kd, nrhs, ab, ldab, b, ldb, info)
call dtbtrs (uplo, trans, diag, n, kd, nrhs, ab, ldab, b, ldb, info)
call ctbtrs (uplo, trans, diag, n, kd, nrhs, ab, ldab, b, ldb, info)
call ztbtrs (uplo, trans, diag, n, kd, nrhs, ab, ldab, b, ldb, info)
```

Discussion

This routine solves for X the following systems of linear equations with a band triangular matrix A , with multiple right-hand sides stored in B :

$AX = B$ if $trans = 'N'$,
 $A^T X = B$ if $trans = 'T'$,
 $A^H X = B$ if $trans = 'C'$ (for complex matrices only).

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
 Indicates whether A is upper or lower triangular:
 If $uplo = 'U'$, then A is upper triangular.
 If $uplo = 'L'$, then A is lower triangular.

trans CHARACTER*1. Must be 'N' or 'T' or 'C'.
 If $trans = 'N'$, then $AX = B$ is solved for X .
 If $trans = 'T'$, then $A^T X = B$ is solved for X .
 If $trans = 'C'$, then $A^H X = B$ is solved for X .

diag CHARACTER*1. Must be 'N' or 'U'.
 If $diag = 'N'$, then A is not a unit triangular matrix.
 If $diag = 'U'$, then A is unit triangular: diagonal elements are assumed to be 1 and not referenced in the array ab .

n INTEGER. The order of A ; the number of rows in B ($n \geq 0$).

kd INTEGER. The number of super-diagonals or sub-diagonals in the matrix A ($kd \geq 0$).

nrhs INTEGER. The number of right-hand sides ($nrhs \geq 0$).

ab, *b* REAL for `stbtrs`
 DOUBLE PRECISION for `dtbtrs`
 COMPLEX for `ctbtrs`
 DOUBLE COMPLEX for `ztbtrs`.
 Arrays: *ab*(*ldab*,*), *b*(*ldb*,*).
 The array *ab* contains the matrix *A* in *band storage* form.
 The array *b* contains the matrix *B* whose columns are the
 right-hand sides for the systems of equations.
 The second dimension of *ab* must be at least $\max(1, n)$,
 the second dimension of *b* at least $\max(1, nrhs)$.

ldab INTEGER. The first dimension of *ab*; $ldab \geq kd + 1$.
ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

Output Parameters

b Overwritten by the solution matrix *X*.
info INTEGER. If *info*=0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For each right-hand side *b*, the computed solution is the exact solution of a perturbed system of equations $(A + E)x = b$ where

$$|E| \leq c(n) \varepsilon |A|$$

$c(n)$ is a modest linear function of n , and ε is the machine precision.

If x_0 is the true solution, the computed solution x satisfies this error bound:

$$\frac{\|x - x_0\|_\infty}{\|x\|_\infty} \leq c(n) \text{cond}(A, x) \varepsilon, \text{ provided } c(n) \text{cond}(A, x) \varepsilon < 1$$

where $\text{cond}(A, x) = \| |A^{-1}| |A| |x| \|_\infty / \|x\|_\infty \leq \|A^{-1}\|_\infty \|A\|_\infty = \kappa_\infty(A)$.

Note that $\text{cond}(A, x)$ can be much smaller than $\kappa_\infty(A)$; the condition number of A^T and A^H might or might not be equal to $\kappa_\infty(A)$.

The approximate number of floating-point operations for one right-hand side vector *b* is $2n * kd$ for real flavors and $8n * kd$ for complex flavors.

To estimate the condition number $\kappa_\infty(A)$, call [?tbcon](#).

To estimate the error in the solution, call [?tbrfs](#).

Routines for Estimating the Condition Number

This section describes the LAPACK routines for estimating the *condition number* of a matrix. The condition number is used for analyzing the errors in the solution of a system of linear equations (see [Error Analysis](#)). Since the condition number may be arbitrarily large when the matrix is nearly singular, the routines actually compute the *reciprocal* condition number.

?gecon

Estimates the reciprocal of the condition number of a general matrix in either the 1-norm or the infinity-norm.

```
call sgecon ( norm,n,a,lda,anorm,rcond,work,iwork,info )
call dgecon ( norm,n,a,lda,anorm,rcond,work,iwork,info )
call cgecon ( norm,n,a,lda,anorm,rcond,work,rwork,info )
call zgecon ( norm,n,a,lda,anorm,rcond,work,rwork,info )
```

Discussion

This routine estimates the reciprocal of the condition number of a general matrix A in either the 1-norm or infinity-norm:

$$\begin{aligned}\kappa_1(A) &= \|A\|_1 \|A^{-1}\|_1 = \kappa_\infty(A^T) = \kappa_\infty(A^H) \\ \kappa_\infty(A) &= \|A\|_\infty \|A^{-1}\|_\infty = \kappa_1(A^T) = \kappa_1(A^H).\end{aligned}$$

Before calling this routine:

- compute *anorm* (either $\|A\|_1 = \max_j \sum_i |a_{ij}|$ or $\|A\|_\infty = \max_i \sum_j |a_{ij}|$)
- call [?getrf](#) to compute the *LU* factorization of A .

Input Parameters

norm CHARACTER*1. Must be '1' or 'O' or 'I'.
If *norm* = '1' or 'O', then the routine estimates $\kappa_1(A)$.
If *norm* = 'I', then the routine estimates $\kappa_\infty(A)$.

n INTEGER. The order of the matrix A ($n \geq 0$).

<i>a</i> , <i>work</i>	<p>REAL for <code>sgecon</code> DOUBLE PRECISION for <code>dgecon</code> COMPLEX for <code>cgecon</code> DOUBLE COMPLEX for <code>zgecon</code>.</p> <p>Arrays: <i>a</i>(<i>lda</i>, *), <i>work</i>(*).</p> <p>The array <i>a</i> contains the <i>LU</i>-factored matrix <i>A</i>, as returned by <code>?getrf</code>.</p> <p>The second dimension of <i>a</i> must be at least $\max(1, n)$.</p> <p>The array <i>work</i> is a workspace for the routine.</p> <p>The dimension of <i>work</i> must be at least $\max(1, 4 * n)$ for real flavors and $\max(1, 2 * n)$ for complex flavors.</p>
<i>anorm</i>	<p>REAL for single precision flavors. DOUBLE PRECISION for double precision flavors.</p> <p>The norm of the <i>original</i> matrix <i>A</i> (see Discussion).</p>
<i>lda</i>	INTEGER. The first dimension of <i>a</i> ; $lda \geq \max(1, n)$.
<i>iwork</i>	INTEGER. Workspace array, DIMENSION at least $\max(1, n)$.
<i>rwork</i>	<p>REAL for <code>cgecon</code> DOUBLE PRECISION for <code>zgecon</code></p> <p>Workspace array, DIMENSION at least $\max(1, 2 * n)$.</p>

Output Parameters

<i>rcond</i>	<p>REAL for single precision flavors. DOUBLE PRECISION for double precision flavors.</p> <p>An estimate of the reciprocal of the condition number. The routine sets <i>rcond</i>=0 if the estimate underflows; in this case the matrix is singular (to working precision). However, anytime <i>rcond</i> is small compared to 1.0, for the working precision, the matrix may be poorly conditioned or even singular.</p>
<i>info</i>	<p>INTEGER. If <i>info</i>=0, the execution is successful. If <i>info</i> = -<i>i</i>, the <i>i</i>th parameter had an illegal value.</p>

Application Notes

The computed *rcond* is never less than ρ (the reciprocal of the true condition number) and in practice is nearly always less than 10ρ . A call to this routine involves solving a number of systems of linear equations $Ax = b$ or $A^Hx = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $2n^2$ floating-point operations for real flavors and $8n^2$ for complex flavors.

?gbcon

Estimates the reciprocal of the condition number of a band matrix in either the 1-norm or the infinity-norm.

```
call sgbcon (norm,n,kl,ku,ab,ldab,ipiv,anorm,rcond,work,iwork,info)
call dgbcon (norm,n,kl,ku,ab,ldab,ipiv,anorm,rcond,work,iwork,info)
call cgbcon (norm,n,kl,ku,ab,ldab,ipiv,anorm,rcond,work,rwork,info)
call zgbcon (norm,n,kl,ku,ab,ldab,ipiv,anorm,rcond,work,rwork,info)
```

Discussion

This routine estimates the reciprocal of the condition number of a general band matrix A in either the 1-norm or infinity-norm:

$$\begin{aligned}\kappa_1(A) &= \|A\|_1 \|A^{-1}\|_1 = \kappa_\infty(A^T) = \kappa_\infty(A^H) \\ \kappa_\infty(A) &= \|A\|_\infty \|A^{-1}\|_\infty = \kappa_1(A^T) = \kappa_1(A^H).\end{aligned}$$

Before calling this routine:

- compute *anorm* (either $\|A\|_1 = \max_j \sum_i |a_{ij}|$ or $\|A\|_\infty = \max_i \sum_j |a_{ij}|$)
- call [?gbtrf](#) to compute the *LU* factorization of A .

Input Parameters

norm CHARACTER*1. Must be '1' or 'O' or 'I'.
If *norm* = '1' or 'O', then the routine estimates $\kappa_1(A)$.
If *norm* = 'I', then the routine estimates $\kappa_\infty(A)$.

n INTEGER. The order of the matrix A ($n \geq 0$).

kl INTEGER. The number of sub-diagonals within the band of A ($kl \geq 0$).

ku INTEGER. The number of super-diagonals within the band of A ($ku \geq 0$).

ldab INTEGER. The first dimension of the array *ab*.
($ldab \geq 2kl + ku + 1$).

ipiv INTEGER. Array, DIMENSION at least $\max(1,n)$.
The *ipiv* array, as returned by [?gbtrf](#).

<i>ab, work</i>	<p>REAL for <code>sgbcon</code> DOUBLE PRECISION for <code>dgbcon</code> COMPLEX for <code>cgbcon</code> DOUBLE COMPLEX for <code>zgbcon</code>. Arrays: <code>ab(ldab,*)</code>, <code>work(*)</code>.</p> <p>The array <code>ab</code> contains the factored band matrix <code>A</code>, as returned by <code>?gbtrf</code>.</p> <p>The second dimension of <code>ab</code> must be at least $\max(1, n)$. The array <code>work</code> is a workspace for the routine.</p> <p>The dimension of <code>work</code> must be at least $\max(1, 3 * n)$ for real flavors and $\max(1, 2 * n)$ for complex flavors.</p>
<i>anorm</i>	<p>REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. The norm of the <i>original</i> matrix <code>A</code> (see Discussion).</p>
<i>iwork</i>	<p>INTEGER. Workspace array, DIMENSION at least $\max(1, n)$.</p>
<i>rwork</i>	<p>REAL for <code>cgbcon</code> DOUBLE PRECISION for <code>zgbcon</code> Workspace array, DIMENSION at least $\max(1, 2 * n)$.</p>

Output Parameters

<i>rcond</i>	<p>REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. An estimate of the reciprocal of the condition number. The routine sets <code>rcond</code>=0 if the estimate underflows; in this case the matrix is singular (to working precision). However, anytime <code>rcond</code> is small compared to 1.0, for the working precision, the matrix may be poorly conditioned or even singular.</p>
<i>info</i>	<p>INTEGER. If <code>info</code>=0, the execution is successful. If <code>info</code> = <code>-i</code>, the <code>i</code>th parameter had an illegal value.</p>

Application Notes

The computed *rcond* is never less than ρ (the reciprocal of the true condition number) and in practice is nearly always less than 10ρ . A call to this routine involves solving a number of systems of linear equations $Ax = b$ or $A^Hx = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $2n(ku + 2kl)$ floating-point operations for real flavors and $8n(ku + 2kl)$ for complex flavors.

?gtcon

Estimates the reciprocal of the condition number of a tridiagonal matrix using the factorization computed by ?gttrf.

```
call sgtcon ( norm,n,d1,d,du,du2,ipiv,anorm,rcond,work,iwork,info )
call dgtcon ( norm,n,d1,d,du,du2,ipiv,anorm,rcond,work,iwork,info )
call cgtcon ( norm,n,d1,d,du,du2,ipiv,anorm,rcond,work,info )
call zgtcon ( norm,n,d1,d,du,du2,ipiv,anorm,rcond,work,info )
```

Discussion

This routine estimates the reciprocal of the condition number of a real or complex tridiagonal matrix A in either the 1-norm or infinity-norm:

$$\kappa_1(A) = \| |A| \|_1 \| |A^{-1}| \|_1$$

$$\kappa_\infty(A) = \| |A| \|_\infty \| |A^{-1}| \|_\infty$$

An estimate is obtained for $\| |A^{-1}| \|$, and the reciprocal of the condition number is computed as $rcond = 1 / (\| |A| \| \| |A^{-1}| \|)$.

Before calling this routine:

- compute *anorm* (either $\| |A| \|_1 = \max_j \sum_i |a_{ij}|$ or $\| |A| \|_\infty = \max_i \sum_j |a_{ij}|$)
- call [?gttrf](#) to compute the *LU* factorization of A .

Input Parameters

<i>norm</i>	<p>CHARACTER*1. Must be '1' or 'O' or 'I'.</p> <p>If <i>norm</i> = '1' or 'O', then the routine estimates $\kappa_1(A)$.</p> <p>If <i>norm</i> = 'I', then the routine estimates $\kappa_\infty(A)$.</p>
<i>n</i>	<p>INTEGER. The order of the matrix <i>A</i> ($n \geq 0$).</p>
<i>dl, d, du, du2</i>	<p>REAL for <i>sgtcon</i></p> <p>DOUBLE PRECISION for <i>dgtcon</i></p> <p>COMPLEX for <i>cgtcon</i></p> <p>DOUBLE COMPLEX for <i>zgtcon</i>.</p> <p>Arrays: <i>dl</i>($n - 1$), <i>d</i>(n), <i>du</i>($n - 1$), <i>du2</i>($n - 2$).</p> <p>The array <i>dl</i> contains the ($n - 1$) multipliers that define the matrix <i>L</i> from the <i>LU</i> factorization of <i>A</i> as computed by ?gttrf.</p> <p>The array <i>d</i> contains the <i>n</i> diagonal elements of the upper triangular matrix <i>U</i> from the <i>LU</i> factorization of <i>A</i>.</p> <p>The array <i>du</i> contains the ($n - 1$) elements of the first super-diagonal of <i>U</i>.</p> <p>The array <i>du2</i> contains the ($n - 2$) elements of the second super-diagonal of <i>U</i>.</p>
<i>ipiv</i>	<p>INTEGER.</p> <p>Array, DIMENSION (<i>n</i>).</p> <p>The array of pivot indices, as returned by ?gttrf.</p>
<i>anorm</i>	<p>REAL for single precision flavors.</p> <p>DOUBLE PRECISION for double precision flavors.</p> <p>The norm of the <i>original</i> matrix <i>A</i> (see <i>Discussion</i>).</p>
<i>work</i>	<p>REAL for <i>sgtcon</i></p> <p>DOUBLE PRECISION for <i>dgtcon</i></p> <p>COMPLEX for <i>cgtcon</i></p> <p>DOUBLE COMPLEX for <i>zgtcon</i>.</p> <p>Workspace array, DIMENSION ($2 * n$).</p>
<i>iwork</i>	<p>INTEGER.</p> <p>Workspace array, DIMENSION (<i>n</i>).</p> <p>Used for real flavors only.</p>

Output Parameters

- rcond* **REAL** for single precision flavors.
 DOUBLE PRECISION for double precision flavors.
 An estimate of the reciprocal of the condition number.
 The routine sets *rcond*=0 if the estimate underflows; in
 this case the matrix is singular (to working precision).
 However, anytime *rcond* is small compared to 1.0,
 for the working precision, the matrix may be poorly
 conditioned or even singular.
- info* **INTEGER**.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The computed *rcond* is never less than ρ (the reciprocal of the true condition number) and in practice is nearly always less than 10ρ . A call to this routine involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $2n^2$ floating-point operations for real flavors and $8n^2$ for complex flavors.

?pocon

Estimates the reciprocal of the condition number of a symmetric (Hermitian) positive-definite matrix.

```
call spocon ( uplo,n,a,lda,anorm,rcond,work,iwork,info )
call dpocon ( uplo,n,a,lda,anorm,rcond,work,iwork,info )
call cpocon ( uplo,n,a,lda,anorm,rcond,work,rwork,info )
call zpocon ( uplo,n,a,lda,anorm,rcond,work,rwork,info )
```

Discussion

This routine estimates the reciprocal of the condition number of a symmetric (Hermitian) positive-definite matrix A :

$$\kappa_1(A) = \|A\|_1 \|A^{-1}\|_1 \quad (\text{since } A \text{ is symmetric or Hermitian, } \kappa_\infty(A) = \kappa_1(A)).$$

Before calling this routine:

- compute *anorm* (either $\|A\|_1 = \max_j \sum_i |a_{ij}|$ or $\|A\|_\infty = \max_i \sum_j |a_{ij}|$)
- call [?potrf](#) to compute the Cholesky factorization of A .

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If *uplo* = 'U', the array *a* stores the upper triangular factor U of the factorization $A = U^H U$.
If *uplo* = 'L', the array *a* stores the lower triangular factor L of the factorization $A = L L^H$.

n INTEGER. The order of the matrix A ($n \geq 0$).

a, *work* REAL for *spocon*
DOUBLE PRECISION for *dpocon*
COMPLEX for *cpocon*
DOUBLE COMPLEX for *zpocon*.
Arrays: *a*(*lda*,*), *work*(*).

The array *a* contains the factored matrix *A*, as returned by `?potrf`.

The second dimension of *a* must be at least $\max(1, n)$.

The array *work* is a workspace for the routine.

The dimension of *work* must be at least $\max(1, 3 * n)$ for real flavors and $\max(1, 2 * n)$ for complex flavors.

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, n)$.

anorm REAL for single precision flavors.

DOUBLE PRECISION for double precision flavors.

The norm of the *original* matrix *A* (see *Discussion*).

iwork INTEGER.

Workspace array, DIMENSION at least $\max(1, n)$.

rwork REAL for `cpocon`

DOUBLE PRECISION for `zpocon`

Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

rcond REAL for single precision flavors.

DOUBLE PRECISION for double precision flavors.

An estimate of the reciprocal of the condition number.

The routine sets *rcond*=0 if the estimate underflows; in this case the matrix is singular (to working precision).

However, anytime *rcond* is small compared to 1.0,

for the working precision, the matrix may be poorly conditioned or even singular.

info INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The computed *rcond* is never less than ρ (the reciprocal of the true condition number) and in practice is nearly always less than 10ρ . A call to this routine involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $2n^2$ floating-point operations for real flavors and $8n^2$ for complex flavors.

?ppcon

Estimates the reciprocal of the condition number of a packed symmetric (Hermitian) positive-definite matrix.

```
call sppcon ( uplo, n, ap, anorm, rcond, work, iwork, info )
call dppcon ( uplo, n, ap, anorm, rcond, work, iwork, info )
call cppcon ( uplo, n, ap, anorm, rcond, work, rwork, info )
call zppcon ( uplo, n, ap, anorm, rcond, work, rwork, info )
```

Discussion

This routine estimates the reciprocal of the condition number of a packed symmetric (Hermitian) positive-definite matrix A :

$$\kappa_1(A) = \|A\|_1 \|A^{-1}\|_1 \quad (\text{since } A \text{ is symmetric or Hermitian, } \kappa_\infty(A) = \kappa_1(A)).$$

Before calling this routine:

- compute *anorm* (either $\|A\|_1 = \max_j \sum_i |a_{ij}|$ or $\|A\|_\infty = \max_i \sum_j |a_{ij}|$)
- call [?pptrf](#) to compute the Cholesky factorization of A .

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If *uplo* = 'U', the array *ap* stores the upper triangular factor U of the factorization $A = U^H U$.
If *uplo* = 'L', the array *ap* stores the lower triangular factor L of the factorization $A = L L^H$.

n INTEGER. The order of the matrix A ($n \geq 0$).

ap, work REAL for *sppcon*
DOUBLE PRECISION for *dppcon*
COMPLEX for *cppcon*
DOUBLE COMPLEX for *zppcon*.
Arrays: *ap*(*), *work*(*).

The array *ap* contains the packed factored matrix *A*, as returned by [?pptrf](#).

The dimension of *ap* must be at least $\max(1, n(n+1)/2)$.

The array *work* is a workspace for the routine.

The dimension of *work* must be at least $\max(1, 3 * n)$ for real flavors and $\max(1, 2 * n)$ for complex flavors.

anorm REAL for single precision flavors.

DOUBLE PRECISION for double precision flavors.

The norm of the *original* matrix *A* (see *Discussion*).

iwork INTEGER.

Workspace array, DIMENSION at least $\max(1, n)$.

rwork REAL for *cppcon*

DOUBLE PRECISION for *zppcon*

Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

rcond REAL for single precision flavors.

DOUBLE PRECISION for double precision flavors.

An estimate of the reciprocal of the condition number.

The routine sets *rcond*=0 if the estimate underflows; in this case the matrix is singular (to working precision).

However, anytime *rcond* is small compared to 1.0, for the working precision, the matrix may be poorly conditioned or even singular.

info INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The computed *rcond* is never less than ρ (the reciprocal of the true condition number) and in practice is nearly always less than 10ρ . A call to this routine involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $2n^2$ floating-point operations for real flavors and $8n^2$ for complex flavors.

?pbcon

Estimates the reciprocal of the condition number of a symmetric (Hermitian) positive-definite band matrix.

```
call spbcon (uplo, n, kd, ab, ldab, anorm, rcond, work, iwork, info)
call dpbcon (uplo, n, kd, ab, ldab, anorm, rcond, work, iwork, info)
call cpbcon (uplo, n, kd, ab, ldab, anorm, rcond, work, rwork, info)
call zpbcon (uplo, n, kd, ab, ldab, anorm, rcond, work, rwork, info)
```

Discussion

This routine estimates the reciprocal of the condition number of a symmetric (Hermitian) positive-definite band matrix A :

$$\kappa_1(A) = \|A\|_1 \|A^{-1}\|_1 \quad (\text{since } A \text{ is symmetric or Hermitian, } \kappa_\infty(A) = \kappa_1(A)).$$

Before calling this routine:

- compute *anorm* (either $\|A\|_1 = \max_j \sum_i |a_{ij}|$ or $\|A\|_\infty = \max_i \sum_j |a_{ij}|$)
- call [?pbtrf](#) to compute the Cholesky factorization of A .

Input Parameters

<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. Indicates how the input matrix A has been factored: If <i>uplo</i> = 'U', the array <i>ab</i> stores the upper triangular factor U of the Cholesky factorization $A = U^H U$. If <i>uplo</i> = 'L', the array <i>ab</i> stores the lower triangular factor L of the factorization $A = L L^H$.
<i>n</i>	INTEGER. The order of the matrix A ($n \geq 0$).
<i>kd</i>	INTEGER. The number of super-diagonals or sub-diagonals in the matrix A ($kd \geq 0$).
<i>ldab</i>	INTEGER. The first dimension of the array <i>ab</i> . ($ldab \geq kd + 1$).
<i>ab, work</i>	REAL for <i>spbcon</i> DOUBLE PRECISION for <i>dpbcon</i> COMPLEX for <i>cpbcon</i> DOUBLE COMPLEX for <i>zpbcon</i> .

Arrays: $ab(ldab, *)$, $work(*)$.

The array ab contains the factored matrix A in band form, as returned by [?pbtrf](#).

The second dimension of ab must be at least $\max(1, n)$,

The array $work$ is a workspace for the routine.

The dimension of $work$ must be at least $\max(1, 3*n)$ for real flavors and $\max(1, 2*n)$ for complex flavors.

$anorm$	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. The norm of the <i>original</i> matrix A (see <i>Discussion</i>).
$iwork$	INTEGER. Workspace array, DIMENSION at least $\max(1, n)$.
$rwork$	REAL for $cpbcon$ DOUBLE PRECISION for $zpbcon$. Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

$rcond$	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. An estimate of the reciprocal of the condition number. The routine sets $rcond=0$ if the estimate underflows; in this case the matrix is singular (to working precision). However, anytime $rcond$ is small compared to 1.0, for the working precision, the matrix may be poorly conditioned or even singular.
$info$	INTEGER. If $info=0$, the execution is successful. If $info = -i$, the i th parameter had an illegal value.

Application Notes

The computed $rcond$ is never less than ρ (the reciprocal of the true condition number) and in practice is nearly always less than 10ρ . A call to this routine involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $4n(kd + 1)$ floating-point operations for real flavors and $16n(kd + 1)$ for complex flavors.

?ptcon

Estimates the reciprocal of the condition number of a symmetric (Hermitian) positive-definite tridiagonal matrix.

```
call sptcon (n, d, e, anorm, rcond, work, info)
call dptcon (n, d, e, anorm, rcond, work, info)
call cptcon (n, d, e, anorm, rcond, work, info)
call zptcon (n, d, e, anorm, rcond, work, info)
```

Discussion

This routine computes the reciprocal of the condition number (in the 1-norm) of a real symmetric or complex Hermitian positive-definite tridiagonal matrix using the factorization $A = LDL^H$ or $A = U^H DU$ computed by [?pttrf](#):

$$\kappa_1(A) = \|A\|_1 \|A^{-1}\|_1 \quad (\text{since } A \text{ is symmetric or Hermitian, } \kappa_\infty(A) = \kappa_1(A)).$$

The norm $\|A^{-1}\|_1$ is computed by a direct method, and the reciprocal of the condition number is computed as $rcond = 1 / (\|A\|_1 \|A^{-1}\|_1)$.

Before calling this routine:

- compute *anorm* as $\|A\|_1 = \max_j \sum_i |a_{ij}|$
- call [?pttrf](#) to compute the factorization of *A*.

Input Parameters

n **INTEGER**. The order of the matrix *A* ($n \geq 0$).

d, work **REAL** for single precision flavors
DOUBLE PRECISION for double precision flavors.
 Arrays, dimension (*n*).
 The array *d* contains the *n* diagonal elements of the diagonal matrix *D* from the factorization of *A*, as computed by [?pttrf](#) ;
work is a workspace array.

e REAL for *sptcon*
 DOUBLE PRECISION for *dptcon*
 COMPLEX for *cptcon*
 DOUBLE COMPLEX for *zptcon*.
 Array, DIMENSION ($n - 1$).
 Contains off-diagonal elements of the unit bidiagonal factor *U* or *L* from the factorization computed by [?pttrf](#).

anorm REAL for single precision flavors.
 DOUBLE PRECISION for double precision flavors.
 The 1- norm of the *original* matrix *A* (see *Discussion*).

Output Parameters

rcond REAL for single precision flavors.
 DOUBLE PRECISION for double precision flavors.
 An estimate of the reciprocal of the condition number.
 The routine sets *rcond*=0 if the estimate underflows; in this case the matrix is singular (to working precision).
 However, anytime *rcond* is small compared to 1.0, for the working precision, the matrix may be poorly conditioned or even singular.

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = $-i$, the *i*th parameter had an illegal value.

Application Notes

The computed *rcond* is never less than ρ (the reciprocal of the true condition number) and in practice is nearly always less than 10ρ . A call to this routine involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $4n(kd + 1)$ floating-point operations for real flavors and $16n(kd + 1)$ for complex flavors.

?sycon

Estimates the reciprocal of the condition number of a symmetric matrix.

```
call ssycon (uplo, n, a, lda, ipiv, anorm, rcond, work, iwork, info)
call dsycon (uplo, n, a, lda, ipiv, anorm, rcond, work, iwork, info)
call csycon (uplo, n, a, lda, ipiv, anorm, rcond, work, rwork, info)
call zsycon (uplo, n, a, lda, ipiv, anorm, rcond, work, rwork, info)
```

Discussion

This routine estimates the reciprocal of the condition number of a symmetric matrix A :

$$\kappa_1(A) = \| |A| \|_1 \|A^{-1}\|_1 \quad (\text{since } A \text{ is symmetric, } \kappa_\infty(A) = \kappa_1(A)).$$

Before calling this routine:

- compute *anorm* (either $\| |A| \|_1 = \max_j \sum_i |a_{ij}|$ or $\| |A| \|_\infty = \max_i \sum_j |a_{ij}|$)
- call [?sytrf](#) to compute the factorization of A .

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
 Indicates how the input matrix A has been factored:
 If *uplo* = 'U', the array *a* stores the upper triangular factor U of the factorization $A = PUDU^T P^T$.
 If *uplo* = 'L', the array *a* stores the lower triangular factor L of the factorization $A = PLDL^T P^T$.

n INTEGER. The order of matrix A ($n \geq 0$).

a, work REAL for *ssycon*
 DOUBLE PRECISION for *dsycon*
 COMPLEX for *csycon*
 DOUBLE COMPLEX for *zsycon*.
 Arrays: *a(lda, *)*, *work(*)*.
 The array *a* contains the factored matrix A , as returned by [?sytrf](#).
 The second dimension of *a* must be at least $\max(1, n)$.

The array *work* is a workspace for the routine.
The dimension of *work* must be at least $\max(1, 2 \cdot n)$.

<i>lda</i>	INTEGER. The first dimension of <i>a</i> ; $lda \geq \max(1, n)$.
<i>ipiv</i>	INTEGER. Array, DIMENSION at least $\max(1, n)$. The array <i>ipiv</i> , as returned by ?sytrf .
<i>anorm</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. The norm of the <i>original</i> matrix <i>A</i> (see <i>Discussion</i>).
<i>iwork</i>	INTEGER. Workspace array, DIMENSION at least $\max(1, n)$.
<i>rwork</i>	REAL for <i>csycon</i> DOUBLE PRECISION for <i>zsycon</i> . Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

<i>rcond</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. An estimate of the reciprocal of the condition number. The routine sets <i>rcond</i> =0 if the estimate underflows; in this case the matrix is singular (to working precision). However, anytime <i>rcond</i> is small compared to 1.0, for the working precision, the matrix may be poorly conditioned or even singular.
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value.

Application Notes

The computed *rcond* is never less than ρ (the reciprocal of the true condition number) and in practice is nearly always less than 10ρ . A call to this routine involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $2n^2$ floating-point operations for real flavors and $8n^2$ for complex flavors.

?hecon

Estimates the reciprocal of the condition number of a Hermitian matrix.

```
call checon (uplo, n, a, lda, ipiv, anorm, rcond, work, rwork, info)
call zhecon (uplo, n, a, lda, ipiv, anorm, rcond, work, rwork, info)
```

Discussion

This routine estimates the reciprocal of the condition number of a Hermitian matrix A :

$$\kappa_1(A) = \|A\|_1 \|A^{-1}\|_1 \quad (\text{since } A \text{ is Hermitian, } \kappa_\infty(A) = \kappa_1(A)).$$

Before calling this routine:

- compute `anorm` (either $\|A\|_1 = \max_j \sum_i |a_{ij}|$ or $\|A\|_\infty = \max_i \sum_j |a_{ij}|$)
- call [?hetrf](#) to compute the factorization of A .

Input Parameters

`uplo` CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If `uplo` = 'U', the array `a` stores the upper triangular factor U of the factorization $A = PUDU^H P^T$.
If `uplo` = 'L', the array `a` stores the lower triangular factor L of the factorization $A = PLDL^H P^T$.

`n` INTEGER. The order of matrix A ($n \geq 0$).

`a, work` COMPLEX for `checon`
DOUBLE COMPLEX for `zhecon`.
Arrays: `a(lda,*)`, `work(*)`.
The array `a` contains the factored matrix A , as returned by [?hetrf](#).
The second dimension of `a` must be at least $\max(1, n)$.
The array `work` is a workspace for the routine.
The dimension of `work` must be at least $\max(1, 2*n)$.

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, n)$.

ipiv INTEGER. Array, DIMENSION at least $\max(1, n)$.
The array *ipiv*, as returned by [?hetrf](#).

anorm REAL for single precision flavors.
DOUBLE PRECISION for double precision flavors.
The norm of the *original* matrix *A* (see *Discussion*).

rwork REAL for *checon*
DOUBLE PRECISION for *zhecon*
Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

rcond REAL for single precision flavors.
DOUBLE PRECISION for double precision flavors.
An estimate of the reciprocal of the condition number.
The routine sets *rcond*=0 if the estimate underflows; in this case the matrix is singular (to working precision).
However, anytime *rcond* is small compared to 1.0, for the working precision, the matrix may be poorly conditioned or even singular.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = *-i*, the *i*th parameter had an illegal value.

Application Notes

The computed *rcond* is never less than ρ (the reciprocal of the true condition number) and in practice is nearly always less than 10ρ . A call to this routine involves solving a number of systems of linear equations $Ax = b$; the number is usually 5 and never more than 11. Each solution requires approximately $8n^2$ floating-point operations.

?spcon

Estimates the reciprocal of the condition number of a packed symmetric matrix.

```
call sspcon ( uplo, n, ap, ipiv, anorm, rcond, work, iwork, info )
call dspcon ( uplo, n, ap, ipiv, anorm, rcond, work, iwork, info )
call cspcon ( uplo, n, ap, ipiv, anorm, rcond, work, rwork, info )
call zspcon ( uplo, n, ap, ipiv, anorm, rcond, work, rwork, info )
```

Discussion

This routine estimates the reciprocal of the condition number of a packed symmetric matrix A :

$$\kappa_1(A) = \|A\|_1 \|A^{-1}\|_1 \quad (\text{since } A \text{ is symmetric, } \kappa_\infty(A) = \kappa_1(A)).$$

Before calling this routine:

- compute *anorm* (either $\|A\|_1 = \max_j \sum_i |a_{ij}|$ or $\|A\|_\infty = \max_i \sum_j |a_{ij}|$)
- call [?sprtf](#) to compute the factorization of A .

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If *uplo* = 'U', the array *ap* stores the packed upper triangular factor U of the factorization $A = PUDU^T P^T$.
If *uplo* = 'L', the array *ap* stores the packed lower triangular factor L of the factorization $A = PLDL^T P^T$.

n INTEGER. The order of matrix A ($n \geq 0$).

ap, work REAL for sspcon
DOUBLE PRECISION for dspcon
COMPLEX for cspcon
DOUBLE COMPLEX for zspcon.
Arrays: *ap*(*), *work*(*).

The array *ap* contains the packed factored matrix A , as returned by [?sprtf](#).
The dimension of *ap* must be at least $\max(1, n(n+1)/2)$.

The array *work* is a workspace for the routine.
 The dimension of *work* must be at least $\max(1, 2 \cdot n)$.

ipiv INTEGER. Array, DIMENSION at least $\max(1, n)$.
 The array *ipiv*, as returned by [?sptf](#).

anorm REAL for single precision flavors.
 DOUBLE PRECISION for double precision flavors.
 The norm of the *original* matrix *A* (see *Discussion*).

iwork INTEGER.
 Workspace array, DIMENSION at least $\max(1, n)$.

rwork REAL for *cspcon*
 DOUBLE PRECISION for *zspcon*
 Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

rcond REAL for single precision flavors.
 DOUBLE PRECISION for double precision flavors.
 An estimate of the reciprocal of the condition number.
 The routine sets *rcond*=0 if the estimate underflows; in this case the matrix is singular (to working precision).
 However, anytime *rcond* is small compared to 1.0, for the working precision, the matrix may be poorly conditioned or even singular.

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The computed *rcond* is never less than ρ (the reciprocal of the true condition number) and in practice is nearly always less than 10ρ . A call to this routine involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $2n^2$ floating-point operations for real flavors and $8n^2$ for complex flavors.

?hpcon

Estimates the reciprocal of the condition number of a packed Hermitian matrix.

```
call chpcon ( uplo, n, ap, ipiv, anorm, rcond, work, rwork, info )
call zhpcon ( uplo, n, ap, ipiv, anorm, rcond, work, rwork, info )
```

Discussion

This routine estimates the reciprocal of the condition number of a Hermitian matrix A :

$$\kappa_1(A) = \|A\|_1 \|A^{-1}\|_1 \quad (\text{since } A \text{ is Hermitian, } \kappa_\infty(A) = \kappa_1(A)).$$

Before calling this routine:

- compute *anorm* (either $\|A\|_1 = \max_j \sum_i |a_{ij}|$ or $\|A\|_\infty = \max_i \sum_j |a_{ij}|$)
- call [?hptrf](#) to compute the factorization of A .

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
 Indicates how the input matrix A has been factored:
 If *uplo* = 'U', the array *ap* stores the packed upper triangular factor U of the factorization $A = PUDU^T P^T$.
 If *uplo* = 'L', the array *ap* stores the packed lower triangular factor L of the factorization $A = PLDL^T P^T$.

n INTEGER. The order of matrix A ($n \geq 0$).

ap, work COMPLEX for *chpcon*
 DOUBLE COMPLEX for *zhpcon*.
 Arrays: *ap*(*), *work*(*).

The array *ap* contains the packed factored matrix A , as returned by [?hptrf](#).
 The dimension of *ap* must be at least $\max(1, n(n+1)/2)$.

The array *work* is a workspace for the routine.
 The dimension of *work* must be at least $\max(1, 2*n)$.

<i>ipiv</i>	INTEGER. Array, DIMENSION at least $\max(1, n)$. The array <i>ipiv</i> , as returned by ?hptrf .
<i>anorm</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. The norm of the <i>original</i> matrix <i>A</i> (see <i>Discussion</i>).
<i>rwork</i>	REAL for <i>chpcon</i> DOUBLE PRECISION for <i>zhpcon</i> . Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

<i>rcond</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. An estimate of the reciprocal of the condition number. The routine sets <i>rcond</i> =0 if the estimate underflows; in this case the matrix is singular (to working precision). However, anytime <i>rcond</i> is small compared to 1.0, for the working precision, the matrix may be poorly conditioned or even singular.
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value.

Application Notes

The computed *rcond* is never less than ρ (the reciprocal of the true condition number) and in practice is nearly always less than 10ρ . A call to this routine involves solving a number of systems of linear equations $Ax = b$; the number is usually 5 and never more than 11. Each solution requires approximately $8n^2$ floating-point operations.

?trcon

Estimates the reciprocal of the condition number of a triangular matrix.

```
call strcon (norm, uplo, diag, n, a, lda, rcond, work, iwork, info)
call dtrcon (norm, uplo, diag, n, a, lda, rcond, work, iwork, info)
call ctrcon (norm, uplo, diag, n, a, lda, rcond, work, rwork, info)
call ztrcon (norm, uplo, diag, n, a, lda, rcond, work, rwork, info)
```

Discussion

This routine estimates the reciprocal of the condition number of a triangular matrix A in either the 1-norm or infinity-norm:

$$\begin{aligned} \kappa_1(A) &= \|A\|_1 \|A^{-1}\|_1 = \kappa_\infty(A^T) = \kappa_\infty(A^H) \\ \kappa_\infty(A) &= \|A\|_\infty \|A^{-1}\|_\infty = \kappa_1(A^T) = \kappa_1(A^H). \end{aligned}$$

Input Parameters

- norm* CHARACTER*1. Must be '1' or 'O' or 'I'.
 If *norm* = '1' or 'O', then the routine estimates $\kappa_1(A)$.
 If *norm* = 'I', then the routine estimates $\kappa_\infty(A)$.
- uplo* CHARACTER*1. Must be 'U' or 'L'.
 Indicates whether A is upper or lower triangular:
 If *uplo* = 'U', the array *a* stores the upper triangle of A , other array elements are not referenced.
 If *uplo* = 'L', the array *a* stores the lower triangle of A , other array elements are not referenced.
- diag* CHARACTER*1. Must be 'N' or 'U'.
 If *diag* = 'N', then A is not a unit triangular matrix.
 If *diag* = 'U', then A is unit triangular: diagonal elements are assumed to be 1 and not referenced in the array *a*.
- n* INTEGER. The order of the matrix A ($n \geq 0$).

a, *work* REAL for *strcon*
 DOUBLE PRECISION for *dtrcon*
 COMPLEX for *ctrcon*
 DOUBLE COMPLEX for *ztrcon*.
 Arrays: *a(lda,*)*, *work(*)*.
 The array *a* contains the matrix *A*.
 The second dimension of *a* must be at least $\max(1, n)$.
 The array *work* is a workspace for the routine.
 The dimension of *work* must be at least $\max(1, 3 * n)$ for real flavors and $\max(1, 2 * n)$ for complex flavors.

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, n)$.

iwork INTEGER.
 Workspace array, DIMENSION at least $\max(1, n)$.

rwork REAL for *ctrcon*
 DOUBLE PRECISION for *ztrcon*.
 Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

rcond REAL for single precision flavors.
 DOUBLE PRECISION for double precision flavors.
 An estimate of the reciprocal of the condition number.
 The routine sets *rcond*=0 if the estimate underflows; in this case the matrix is singular (to working precision).
 However, anytime *rcond* is small compared to 1.0, for the working precision, the matrix may be poorly conditioned or even singular.

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The computed *rcond* is never less than ρ (the reciprocal of the true condition number) and in practice is nearly always less than 10ρ . A call to this routine involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately n^2 floating-point operations for real flavors and $4n^2$ operations for complex flavors.

?tpcon

Estimates the reciprocal of the condition number of a packed triangular matrix.

```
call stpcon (norm,uplo,diag,n,ap,rcond,work,iwork,info)
call dtpcon (norm,uplo,diag,n,ap,rcond,work,iwork,info)
call ctpcon (norm,uplo,diag,n,ap,rcond,work,rwork,info)
call ztpcon (norm,uplo,diag,n,ap,rcond,work,rwork,info)
```

Discussion

This routine estimates the reciprocal of the condition number of a packed triangular matrix A in either the 1-norm or infinity-norm:

$$\kappa_1(A) = \| |A| \|_1 \| |A^{-1}| \|_1 = \kappa_\infty(A^T) = \kappa_\infty(A^H)$$

$$\kappa_\infty(A) = \| |A| \|_\infty \| |A^{-1}| \|_\infty = \kappa_1(A^T) = \kappa_1(A^H).$$

Input Parameters

norm CHARACTER*1. Must be '1' or 'O' or 'I'.
 If *norm* = '1' or 'O', then the routine estimates $\kappa_1(A)$.
 If *norm* = 'I', then the routine estimates $\kappa_\infty(A)$.

uplo CHARACTER*1. Must be 'U' or 'L'.
 Indicates whether A is upper or lower triangular:
 If *uplo* = 'U', the array *ap* stores the upper triangle of A in packed form.
 If *uplo* = 'L', the array *ap* stores the lower triangle of A in packed form.

diag CHARACTER*1. Must be 'N' or 'U'.
 If *diag* = 'N', then A is not a unit triangular matrix.
 If *diag* = 'U', then A is unit triangular: diagonal elements are assumed to be 1 and not referenced in the array *ap*.

n INTEGER. The order of the matrix A ($n \geq 0$).

ap, work REAL for *stpcon*
 DOUBLE PRECISION for *dtpcon*
 COMPLEX for *ctpcon*
 DOUBLE COMPLEX for *ztpcon*.
 Arrays: *ap*(*), *work*(*).
 The array *ap* contains the packed matrix *A*.
 The dimension of *ap* must be at least $\max(1, n(n+1)/2)$.
 The array *work* is a workspace for the routine.
 The dimension of *work* must be at least $\max(1, 3*n)$ for real flavors and $\max(1, 2*n)$ for complex flavors.

iwork INTEGER.
 Workspace array, DIMENSION at least $\max(1, n)$.

rwork REAL for *ctpcon*
 DOUBLE PRECISION for *ztpcon*
 Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

rcond REAL for single precision flavors.
 DOUBLE PRECISION for double precision flavors.
 An estimate of the reciprocal of the condition number.
 The routine sets *rcond* = 0 if the estimate underflows; in this case the matrix is singular (to working precision).
 However, anytime *rcond* is small compared to 1.0, for the working precision, the matrix may be poorly conditioned or even singular.

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The computed *rcond* is never less than ρ (the reciprocal of the true condition number) and in practice is nearly always less than 10ρ . A call to this routine involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately n^2 floating-point operations for real flavors and $4n^2$ operations for complex flavors.

?tbcon

Estimates the reciprocal of the condition number of a triangular band matrix.

```
call stbcon ( norm,uplo,diag,n,kd,ab,ldab,rcond,work,iwork,info )
call dtbcon ( norm,uplo,diag,n,kd,ab,ldab,rcond,work,iwork,info )
call ctbcon ( norm,uplo,diag,n,kd,ab,ldab,rcond,work,rwork,info )
call ztbcon ( norm,uplo,diag,n,kd,ab,ldab,rcond,work,rwork,info )
```

Discussion

This routine estimates the reciprocal of the condition number of a triangular band matrix A in either the 1-norm or infinity-norm:

$$\kappa_1(A) = \| |A| \|_1 \| |A^{-1}| \|_1 = \kappa_\infty(A^T) = \kappa_\infty(A^H)$$

$$\kappa_\infty(A) = \| |A| \|_\infty \| |A^{-1}| \|_\infty = \kappa_1(A^T) = \kappa_1(A^H).$$

Input Parameters

norm CHARACTER*1. Must be '1' or 'O' or 'I'.
 If *norm* = '1' or 'O', then the routine estimates $\kappa_1(A)$.
 If *norm* = 'I', then the routine estimates $\kappa_\infty(A)$.

uplo CHARACTER*1. Must be 'U' or 'L'.
 Indicates whether A is upper or lower triangular:
 If *uplo* = 'U', the array *ap* stores the upper triangle of A in packed form.
 If *uplo* = 'L', the array *ap* stores the lower triangle of A in packed form.

diag CHARACTER*1. Must be 'N' or 'U'.
 If *diag* = 'N', then A is not a unit triangular matrix.
 If *diag* = 'U', then A is unit triangular: diagonal elements are assumed to be 1 and not referenced in the array *ab*.

n INTEGER. The order of the matrix A ($n \geq 0$).

kd INTEGER. The number of super-diagonals or sub-diagonals in the matrix A ($kd \geq 0$).

ab, work REAL for *stbcon*
 DOUBLE PRECISION for *dtbcon*
 COMPLEX for *ctbcon*
 DOUBLE COMPLEX for *ztbcon*.
 Arrays: *ab(ldab,*)*, *work(*)*.
 The array *ab* contains the band matrix *A*.
 The second dimension of *ab* must be at least $\max(1, n)$.
 The array *work* is a workspace for the routine.
 The dimension of *work* must be at least $\max(1, 3 * n)$ for
 real flavors and $\max(1, 2 * n)$ for complex flavors.

ldab INTEGER. The first dimension of the array *ab*.
 ($ldab \geq kd + 1$).

iwork INTEGER.
 Workspace array, DIMENSION at least $\max(1, n)$.

rwork REAL for *ctbcon*
 DOUBLE PRECISION for *ztbcon*.
 Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

rcond REAL for single precision flavors.
 DOUBLE PRECISION for double precision flavors.
 An estimate of the reciprocal of the condition number.
 The routine sets *rcond*=0 if the estimate underflows; in
 this case the matrix is singular (to working precision).
 However, anytime *rcond* is small compared to 1.0,
 for the working precision, the matrix may be poorly
 conditioned or even singular.

info INTEGER. If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The computed *rcond* is never less than ρ (the reciprocal of the true condition number) and in practice is nearly always less than 10ρ . A call to this routine involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $2n(kd + 1)$ floating-point operations for real flavors and $8n(kd + 1)$ operations for complex flavors.

Refining the Solution and Estimating Its Error

This section describes the LAPACK routines for refining the computed solution of a system of linear equations and estimating the solution error. You can call these routines after factorizing the matrix of the system of equations and computing the solution (see [Routines for Matrix Factorization](#) and [Routines for Solving Systems of Linear Equations](#)).

?gerfs

Refines the solution of a system of linear equations with a general matrix and estimates its error.

```
call sgerfs (trans,n,nrhs,a,lda,af,ldaf,ipiv,b,ldb,
            x,ldx,ferr,berr,work,iwork,info)
```

```
call dgerfs (trans,n,nrhs,a,lda,af,ldaf,ipiv,b,ldb,
            x,ldx,ferr,berr,work,iwork,info)
```

```
call cgerfs (trans,n,nrhs,a,lda,af,ldaf,ipiv,b,ldb,
            x,ldx,ferr,berr,work,rwork,info)
```

```
call zgerfs (trans,n,nrhs,a,lda,af,ldaf,ipiv,b,ldb,
            x,ldx,ferr,berr,work,rwork,info)
```

Discussion

This routine performs an iterative refinement of the solution to a system of linear equations $AX = B$ or $A^T X = B$ or $A^H X = B$ with a general matrix A , with multiple right-hand sides. For each computed solution vector x , the routine computes the *component-wise backward error* β . This error is the smallest relative perturbation in elements of A and b such that x is the exact solution of the perturbed system:

$$|\delta a_{ij}| / |a_{ij}| \leq \beta |a_{ij}|, \quad |\delta b_i| / |b_i| \leq \beta |b_i| \text{ such that } (A + \delta A)x = (b + \delta b).$$

Finally, the routine estimates the *component-wise forward error* in the computed solution $\| |x - x_e| \|_\infty / \|x\|_\infty$ (here x_e is the exact solution).

Before calling this routine:

- call the factorization routine [?getrf](#)
- call the solver routine [?getrs](#).

Input Parameters

trans CHARACTER*1. Must be 'N' or 'T' or 'C'.
Indicates the form of the equations:
If *trans* = 'N', the system has the form $AX = B$.
If *trans* = 'T', the system has the form $A^T X = B$.
If *trans* = 'C', the system has the form $A^H X = B$.

n INTEGER. The order of the matrix A ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides ($nrhs \geq 0$).

a, *af*, *b*, *x*, *work* REAL for *sgerfs*
DOUBLE PRECISION for *dgerfs*
COMPLEX for *cgerfs*
DOUBLE COMPLEX for *zgerfs*.

Arrays:

a(*lda*,*) contains the original matrix A , as supplied to [?getrf](#).

af(*ldaf*,*) contains the factored matrix A , as returned by [?getrf](#).

b(*ldb*,*) contains the right-hand side matrix B .

x(*ldx*,*) contains the solution matrix X .

work(*) is a workspace array.

The second dimension of *a* and *af* must be at least $\max(1, n)$; the second dimension of *b* and *x* must be at least $\max(1, nrhs)$; the dimension of *work* must be at least $\max(1, 3 * n)$ for real flavors and $\max(1, 2 * n)$ for complex flavors.

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, n)$.

ldaf INTEGER. The first dimension of *af*; $ldaf \geq \max(1, n)$.

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

ldx INTEGER. The first dimension of *x*; $ldx \geq \max(1, n)$.

<i>ipiv</i>	INTEGER. Array, DIMENSION at least max(1,n). The <i>ipiv</i> array, as returned by ?getrf .
<i>iwork</i>	INTEGER. Workspace array, DIMENSION at least max(1, n).
<i>rwork</i>	REAL for <i>cgerfs</i> DOUBLE PRECISION for <i>zgerfs</i> . Workspace array, DIMENSION at least max(1, n).

Output Parameters

<i>x</i>	The refined solution matrix <i>X</i> .
<i>ferr, berr</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. Arrays, DIMENSION at least max(1, <i>nrhs</i>). Contain the component-wise forward and backward errors, respectively, for each solution vector.
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value.

Application Notes

The bounds returned in *ferr* are not rigorous, but in practice they almost always overestimate the actual error.

For each right-hand side, computation of the backward error involves a minimum of $4n^2$ floating-point operations (for real flavors) or $16n^2$ operations (for complex flavors). In addition, each step of iterative refinement involves $6n^2$ operations (for real flavors) or $24n^2$ operations (for complex flavors); the number of iterations may range from 1 to 5. Estimating the forward error involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $2n^2$ floating-point operations for real flavors or $8n^2$ for complex flavors.

?gbrfs

Refines the solution of a system of linear equations with a general band matrix and estimates its error.

```
call sgbtrfs (trans,n,kl,ku,nrhs,ab,ldab,afb,ldafb,ipiv,
             b,ldb,x,ldx,ferr,berr,work,iwork,info)
```

```
call dgbtrfs (trans,n,kl,ku,nrhs,ab,ldab,afb,ldafb,ipiv,
             b,ldb,x,ldx,ferr,berr,work,iwork,info)
```

```
call cgbtrfs (trans,n,kl,ku,nrhs,ab,ldab,afb,ldafb,ipiv,
             b,ldb,x,ldx,ferr,berr,work,rwork,info)
```

```
call zgbtrfs (trans,n,kl,ku,nrhs,ab,ldab,afb,ldafb,ipiv,
             b,ldb,x,ldx,ferr,berr,work,rwork,info)
```

Discussion

This routine performs an iterative refinement of the solution to a system of linear equations $AX = B$ or $A^T X = B$ or $A^H X = B$ with a band matrix A , with multiple right-hand sides. For each computed solution vector x , the routine computes the *component-wise backward error* β . This error is the smallest relative perturbation in elements of A and b such that x is the exact solution of the perturbed system:

$$|\delta a_{ij}| / |a_{ij}| \leq \beta |a_{ij}|, \quad |\delta b_i| / |b_i| \leq \beta |b_i| \text{ such that } (A + \delta A)x = (b + \delta b).$$

Finally, the routine estimates the *component-wise forward error* in the computed solution $\|x - x_e\|_\infty / \|x\|_\infty$ (here x_e is the exact solution).

Before calling this routine:

- call the factorization routine [?gbtrf](#)
- call the solver routine [?gbtrs](#).

Input Parameters

<i>trans</i>	CHARACTER*1. Must be 'N' or 'T' or 'C'. Indicates the form of the equations: If <i>trans</i> = 'N', the system has the form $AX = B$. If <i>trans</i> = 'T', the system has the form $A^T X = B$. If <i>trans</i> = 'C', the system has the form $A^H X = B$.
<i>n</i>	INTEGER. The order of the matrix A ($n \geq 0$).
<i>kl</i>	INTEGER. The number of sub-diagonals within the band of A ($kl \geq 0$).
<i>ku</i>	INTEGER. The number of super-diagonals within the band of A ($ku \geq 0$).
<i>nrhs</i>	INTEGER. The number of right-hand sides ($nrhs \geq 0$).
<i>ab,afb,b,x,work</i>	REAL for <i>sgbrfs</i> DOUBLE PRECISION for <i>dgbrfs</i> COMPLEX for <i>cgbrfs</i> DOUBLE COMPLEX for <i>zgbrfs</i> .

Arrays:

ab(ldab,)* contains the original band matrix A , as supplied to [?gbtrf](#), but stored in rows from 1 to $kl + ku + 1$.

afb(ldafb,)* contains the factored band matrix A , as returned by [?gbtrf](#).

b(ldb,)* contains the right-hand side matrix B .

x(ldx,)* contains the solution matrix X .

work()* is a workspace array.

The second dimension of *ab* and *afb* must be at least $\max(1, n)$; the second dimension of *b* and *x* must be at least $\max(1, nrhs)$; the dimension of *work* must be at least $\max(1, 3 * n)$ for real flavors and $\max(1, 2 * n)$ for complex flavors.

<i>ldab</i>	INTEGER. The first dimension of <i>ab</i> .
<i>ldafb</i>	INTEGER. The first dimension of <i>afb</i> .
<i>ldb</i>	INTEGER. The first dimension of <i>b</i> ; $ldb \geq \max(1, n)$.
<i>ldx</i>	INTEGER. The first dimension of <i>x</i> ; $ldx \geq \max(1, n)$.

ipiv INTEGER.
Array, DIMENSION at least $\max(1, n)$.
The *ipiv* array, as returned by [?gbtrf](#).

iwork INTEGER.
Workspace array, DIMENSION at least $\max(1, n)$.

rwork REAL for *cgbrfs*
DOUBLE PRECISION for *zgbrfs*
Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

x The refined solution matrix X .

ferr, berr REAL for single precision flavors.
DOUBLE PRECISION for double precision flavors.
Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and backward errors, respectively, for each solution vector.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = $-i$, the *i*th parameter had an illegal value.

Application Notes

The bounds returned in *ferr* are not rigorous, but in practice they almost always overestimate the actual error.

For each right-hand side, computation of the backward error involves a minimum of $4n(kl + ku)$ floating-point operations (for real flavors) or $16n(kl + ku)$ operations (for complex flavors). In addition, each step of iterative refinement involves $2n(4kl + 3ku)$ operations (for real flavors) or $8n(4kl + 3ku)$ operations (for complex flavors); the number of iterations may range from 1 to 5. Estimating the forward error involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $2n^2$ floating-point operations for real flavors or $8n^2$ for complex flavors.

?gtrfs

Refines the solution of a system of linear equations with a tridiagonal matrix and estimates its error.

```
call sgtrfs (trans, n, nrhs, dl, d, du, dlf, df, duf, du2, ipiv, b,
            ldb, x, ldx, ferr, berr, work, iwork, info)
call dgtrfs (trans, n, nrhs, dl, d, du, dlf, df, duf, du2, ipiv, b,
            ldb, x, ldx, ferr, berr, work, iwork, info)
call cgtrfs (trans, n, nrhs, dl, d, du, dlf, df, duf, du2, ipiv, b,
            ldb, x, ldx, ferr, berr, work, rwork, info)
call zgtrfs (trans, n, nrhs, dl, d, du, dlf, df, duf, du2, ipiv, b,
            ldb, x, ldx, ferr, berr, work, rwork, info)
```

Discussion

This routine performs an iterative refinement of the solution to a system of linear equations $AX = B$ or $A^T X = B$ or $A^H X = B$ with a tridiagonal matrix A , with multiple right-hand sides. For each computed solution vector x , the routine computes the *component-wise backward error* β . This error is the smallest relative perturbation in elements of A and b such that x is the exact solution of the perturbed system:

$$|\delta a_{ij}| / |a_{ij}| \leq \beta |a_{ij}|, \quad |\delta b_i| / |b_i| \leq \beta |b_i| \quad \text{such that } (A + \delta A)x = (b + \delta b).$$

Finally, the routine estimates the *component-wise forward error* in the computed solution $\|x - x_e\| \propto \|x\|_\infty$ (here x_e is the exact solution).

Before calling this routine:

- call the factorization routine [?gttrf](#)
- call the solver routine [?gttrs](#).

Input Parameters

trans CHARACTER*1. Must be 'N' or 'T' or 'C'.

Indicates the form of the equations:

If *trans* = 'N', the system has the form $AX = B$.

If *trans* = 'T', the system has the form $A^T X = B$.

If *trans* = 'C', the system has the form $A^H X = B$.

n **INTEGER**. The order of the matrix *A* ($n \geq 0$).
nrhs **INTEGER**. The number of right-hand sides, i.e., the number of columns of the matrix *B* ($nrhs \geq 0$).
dl, d, du, dlf, df, duf, du2, b, x, work **REAL** for *sgtrfs*
 DOUBLE PRECISION for *dgtrfs*
 COMPLEX for *cgtrfs*
 DOUBLE COMPLEX for *zgtrfs*.
 Arrays:
dl, dimension ($n - 1$), contains the subdiagonal elements of *A*.
d, dimension (n), contains the diagonal elements of *A*.
du, dimension ($n - 1$), contains the superdiagonal elements of *A*.
dlf, dimension ($n - 1$), contains the ($n - 1$) multipliers that define the matrix *L* from the *LU* factorization of *A* as computed by [?gttrf](#).
df, dimension (n), contains the n diagonal elements of the upper triangular matrix *U* from the *LU* factorization of *A*.
duf, dimension ($n - 1$), contains the ($n - 1$) elements of the first super-diagonal of *U*.
du2, dimension ($n - 2$), contains the ($n - 2$) elements of the second super-diagonal of *U*.
b(ldb, nrhs) contains the right-hand side matrix *B*.
x(ldx, nrhs) contains the solution matrix *X*, as computed by [?gttrs](#).
work ()* is a workspace array;
 the dimension of *work* must be at least $\max(1, 3 * n)$ for real flavors and $\max(1, 2 * n)$ for complex flavors.
ldb **INTEGER**. The first dimension of *b*; $ldb \geq \max(1, n)$.
ldx **INTEGER**. The first dimension of *x*; $ldx \geq \max(1, n)$.
ipiv **INTEGER**.
 Array, **DIMENSION** at least $\max(1, n)$.
 The *ipiv* array, as returned by [?gttrf](#).

iwork **INTEGER.**
 Workspace array, **DIMENSION** (*n*). Used for real flavors only.

rwork **REAL** for **cgtrfs**
DOUBLE PRECISION for **zgtrfs**.
 Workspace array, **DIMENSION** (*n*). Used for complex flavors only.

Output Parameters

x The refined solution matrix *X*.

ferr, berr **REAL** for single precision flavors.
DOUBLE PRECISION for double precision flavors.
 Arrays, **DIMENSION** at least $\max(1, nrhs)$. Contain the component-wise forward and backward errors, respectively, for each solution vector.

info **INTEGER.**
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

?porfs

Refines the solution of a system of linear equations with a symmetric (Hermitian) positive-definite matrix and estimates its error.

```
call sporfs (uplo,n,nrhs,a,lda,af,ldaf,b,ldb,
             x,ldx,ferr,berr,work,iwork,info)
```

```
call dporfs (uplo,n,nrhs,a,lda,af,ldaf,b,ldb,
             x,ldx,ferr,berr,work,iwork,info)
```

```
call cporfs (uplo,n,nrhs,a,lda,af,ldaf,b,ldb,
             x,ldx,ferr,berr,work,rwork,info)
```

```
call zporfs (uplo,n,nrhs,a,lda,af,ldaf,b,ldb,
             x,ldx,ferr,berr,work,rwork,info)
```

Discussion

This routine performs an iterative refinement of the solution to a system of linear equations $AX = B$ with a symmetric (Hermitian) positive definite matrix A , with multiple right-hand sides. For each computed solution vector x , the routine computes the *component-wise backward error* β . This error is the smallest relative perturbation in elements of A and b such that x is the exact solution of the perturbed system:

$$|\delta a_{ij}| / |a_{ij}| \leq \beta |a_{ij}|, \quad |\delta b_i| / |b_i| \leq \beta |b_i| \quad \text{such that } (A + \delta A)x = (b + \delta b).$$

Finally, the routine estimates the *component-wise forward error* in the computed solution $\|x - x_e\|_\infty / \|x\|_\infty$ (here x_e is the exact solution).

Before calling this routine:

- call the factorization routine [?potrf](#)
- call the solver routine [?potrs](#).

Input Parameters

<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. Indicates how the input matrix <i>A</i> has been factored: If <i>uplo</i> = 'U', the array <i>af</i> stores the factor <i>U</i> of the Cholesky factorization $A = U^H U$. If <i>uplo</i> = 'L', the array <i>af</i> stores the factor <i>L</i> of the Cholesky factorization $A = LL^H$.
<i>n</i>	INTEGER. The order of the matrix <i>A</i> ($n \geq 0$).
<i>nrhs</i>	INTEGER. The number of right-hand sides ($nrhs \geq 0$).
<i>a</i> , <i>af</i> , <i>b</i> , <i>x</i> , <i>work</i>	REAL for <i>sporfs</i> DOUBLE PRECISION for <i>dporfs</i> COMPLEX for <i>cporfs</i> DOUBLE COMPLEX for <i>zporfs</i> .
	Arrays: <i>a</i> (<i>lda</i> ,*) contains the original matrix <i>A</i> , as supplied to ?potrf . <i>af</i> (<i>ldaf</i> ,*) contains the factored matrix <i>A</i> , as returned by ?potrf . <i>b</i> (<i>ldb</i> ,*) contains the right-hand side matrix <i>B</i> . <i>x</i> (<i>ldx</i> ,*) contains the solution matrix <i>X</i> . <i>work</i> (*) is a workspace array. The second dimension of <i>a</i> and <i>af</i> must be at least $\max(1, n)$; the second dimension of <i>b</i> and <i>x</i> must be at least $\max(1, nrhs)$; the dimension of <i>work</i> must be at least $\max(1, 3 * n)$ for real flavors and $\max(1, 2 * n)$ for complex flavors.
<i>lda</i>	INTEGER. The first dimension of <i>a</i> ; $lda \geq \max(1, n)$.
<i>ldaf</i>	INTEGER. The first dimension of <i>af</i> ; $ldaf \geq \max(1, n)$.
<i>ldb</i>	INTEGER. The first dimension of <i>b</i> ; $ldb \geq \max(1, n)$.
<i>ldx</i>	INTEGER. The first dimension of <i>x</i> ; $ldx \geq \max(1, n)$.
<i>iwork</i>	INTEGER. Workspace array, DIMENSION at least $\max(1, n)$.
<i>rwork</i>	REAL for <i>cporfs</i> DOUBLE PRECISION for <i>zporfs</i> Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

<i>x</i>	The refined solution matrix X .
<i>ferr</i> , <i>berr</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and backward errors, respectively, for each solution vector.
<i>info</i>	INTEGER . If <i>info</i> = 0, the execution is successful. If <i>info</i> = $-i$, the <i>i</i> th parameter had an illegal value.

Application Notes

The bounds returned in *ferr* are not rigorous, but in practice they almost always overestimate the actual error.

For each right-hand side, computation of the backward error involves a minimum of $4n^2$ floating-point operations (for real flavors) or $16n^2$ operations (for complex flavors). In addition, each step of iterative refinement involves $6n^2$ operations (for real flavors) or $24n^2$ operations (for complex flavors); the number of iterations may range from 1 to 5.

Estimating the forward error involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $2n^2$ floating-point operations for real flavors or $8n^2$ for complex flavors.

?pprfs

Refines the solution of a system of linear equations with a packed symmetric (Hermitian) positive-definite matrix and estimates its error.

```
call sprfs (uplo,n,nrhs,ap,afp,b,ldb,x,ldx,
           ferr,berr,work,iwork,info)
call dpprfs (uplo,n,nrhs,ap,afp,b,ldb,x,ldx,
           ferr,berr,work,iwork,info)

call cpprfs (uplo,n,nrhs,ap,afp,b,ldb,x,ldx,
           ferr,berr,work,rwork,info)
call zpprfs (uplo,n,nrhs,ap,afp,b,ldb,x,ldx,
           ferr,berr,work,rwork,info)
```

Discussion

This routine performs an iterative refinement of the solution to a system of linear equations $AX = B$ with a packed symmetric (Hermitian) positive definite matrix A , with multiple right-hand sides. For each computed solution vector x , the routine computes the *component-wise backward error* β . This error is the smallest relative perturbation in elements of A and b such that x is the exact solution of the perturbed system:

$$|\delta a_{ij}| / |a_{ij}| \leq \beta |a_{ij}|, \quad |\delta b_i| / |b_i| \leq \beta |b_i| \quad \text{such that } (A + \delta A)x = (b + \delta b).$$

Finally, the routine estimates the *component-wise forward error* in the computed solution $\| |x - x_e| | \infty / \| |x| | \infty$ (here x_e is the exact solution).

Before calling this routine:

- call the factorization routine [?pptrf](#)
- call the solver routine [?pptrs](#).

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:

If *uplo* = 'U', the array *afp* stores the packed factor *U* of the Cholesky factorization $A = U^H U$.

If *uplo* = 'L', the array *afp* stores the packed factor *L* of the Cholesky factorization $A = LL^H$.

n INTEGER. The order of the matrix *A* ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides ($nrhs \geq 0$).

ap, *afp*, *b*, *x*, *work* REAL for *spprfs*
 DOUBLE PRECISION for *dpprfs*
 COMPLEX for *cpprfs*
 DOUBLE COMPLEX for *zpprfs*.

Arrays:

ap(*) contains the original packed matrix *A*, as supplied to [?pptrf](#).

afp(*) contains the factored packed matrix *A*, as returned by [?pptrf](#).

b(*ldb*,*) contains the right-hand side matrix *B*.

x(*ldx*,*) contains the solution matrix *X*.

work(*) is a workspace array.

The dimension of arrays *ap* and *afp* must be at least $\max(1, n(n+1)/2)$; the second dimension of *b* and *x* must be at least $\max(1, nrhs)$; the dimension of *work* must be at least $\max(1, 3*n)$ for real flavors and $\max(1, 2*n)$ for complex flavors.

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

ldx INTEGER. The first dimension of *x*; $ldx \geq \max(1, n)$.

iwork INTEGER.
 Workspace array, DIMENSION at least $\max(1, n)$.

rwork REAL for *cpprfs*
 DOUBLE PRECISION for *zpprfs*
 Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

x The refined solution matrix *X*.

ferr, berr REAL for single precision flavors.
DOUBLE PRECISION for double precision flavors.
Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and backward errors, respectively, for each solution vector.

info INTEGER. If *info*=0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The bounds returned in *ferr* are not rigorous, but in practice they almost always overestimate the actual error.

For each right-hand side, computation of the backward error involves a minimum of $4n^2$ floating-point operations (for real flavors) or $16n^2$ operations (for complex flavors). In addition, each step of iterative refinement involves $6n^2$ operations (for real flavors) or $24n^2$ operations (for complex flavors); the number of iterations may range from 1 to 5.

Estimating the forward error involves solving a number of systems of linear equations $Ax = b$; the number of systems is usually 4 or 5 and never more than 11. Each solution requires approximately $2n^2$ floating-point operations for real flavors or $8n^2$ for complex flavors.

?pbrfs

Refines the solution of a system of linear equations with a band symmetric (Hermitian) positive-definite matrix and estimates its error.

```
call spbrfs (uplo, n, kd, nrhs, ab, ldab, afb, ldafb,
             b, ldb, x, ldx, ferr, berr, work, iwork, info)
call dpbrfs (uplo, n, kd, nrhs, ab, ldab, afb, ldafb,
             b, ldb, x, ldx, ferr, berr, work, iwork, info)

call cpbrfs (uplo, n, kd, nrhs, ab, ldab, afb, ldafb,
             b, ldb, x, ldx, ferr, berr, work, rwork, info)
call zpbrfs (uplo, n, kd, nrhs, ab, ldab, afb, ldafb,
             b, ldb, x, ldx, ferr, berr, work, rwork, info)
```

Discussion

This routine performs an iterative refinement of the solution to a system of linear equations $AX = B$ with a symmetric (Hermitian) positive definite band matrix A , with multiple right-hand sides. For each computed solution vector x , the routine computes the *component-wise backward error* β . This error is the smallest relative perturbation in elements of A and b such that x is the exact solution of the perturbed system:

$$|\delta a_{ij}| / |a_{ij}| \leq \beta |a_{ij}|, \quad |\delta b_i| / |b_i| \leq \beta |b_i| \quad \text{such that } (A + \delta A)x = (b + \delta b).$$

Finally, the routine estimates the *component-wise forward error* in the computed solution $\|x - x_e\| / \|x\|$ (here x_e is the exact solution).

Before calling this routine:

- call the factorization routine [?pbtrf](#)
- call the solver routine [?pbtrs](#).

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:

If *uplo* = 'U', the array *afb* stores the factor *U* of the Cholesky factorization $A = U^H U$.

If *uplo* = 'L', the array *afb* stores the factor *L* of the Cholesky factorization $A = LL^H$.

<i>n</i>	INTEGER. The order of the matrix <i>A</i> ($n \geq 0$).
<i>kd</i>	INTEGER. The number of super-diagonals or sub-diagonals in the matrix <i>A</i> ($kd \geq 0$).
<i>nrhs</i>	INTEGER. The number of right-hand sides ($nrhs \geq 0$).
<i>ab,afb,b,x,work</i>	REAL for <i>spbrfs</i> DOUBLE PRECISION for <i>dpbrfs</i> COMPLEX for <i>cpbrfs</i> DOUBLE COMPLEX for <i>zpbrfs</i> .

Arrays:

ab(ldab,)* contains the original band matrix *A*, as supplied to [?pbtrf](#).

afb(ldafb,)* contains the factored band matrix *A*, as returned by [?pbtrf](#).

b(ldb,)* contains the right-hand side matrix *B*.

x(ldx,)* contains the solution matrix *X*.

work()* is a workspace array.

The second dimension of *ab* and *afb* must be at least $\max(1,n)$; the second dimension of *b* and *x* must be at least $\max(1,nrhs)$; the dimension of *work* must be at least $\max(1, 3*n)$ for real flavors and $\max(1, 2*n)$ for complex flavors.

<i>ldab</i>	INTEGER. The first dimension of <i>ab</i> ; $ldab \geq kd + 1$.
<i>ldafb</i>	INTEGER. The first dimension of <i>afb</i> ; $ldafb \geq kd + 1$.
<i>ldb</i>	INTEGER. The first dimension of <i>b</i> ; $ldb \geq \max(1, n)$.
<i>ldx</i>	INTEGER. The first dimension of <i>x</i> ; $ldx \geq \max(1, n)$.
<i>iwork</i>	INTEGER. Workspace array, DIMENSION at least $\max(1, n)$.
<i>rwork</i>	REAL for <i>cpbrfs</i> DOUBLE PRECISION for <i>zpbrfs</i> Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

<i>x</i>	The refined solution matrix X .
<i>ferr</i> , <i>berr</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and backward errors, respectively, for each solution vector.
<i>info</i>	INTEGER . If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value.

Application Notes

The bounds returned in *ferr* are not rigorous, but in practice they almost always overestimate the actual error.

For each right-hand side, computation of the backward error involves a minimum of $8n*kd$ floating-point operations (for real flavors) or $32n*kd$ operations (for complex flavors). In addition, each step of iterative refinement involves $12n*kd$ operations (for real flavors) or $48n*kd$ operations (for complex flavors); the number of iterations may range from 1 to 5.

Estimating the forward error involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $4n*kd$ floating-point operations for real flavors or $16n*kd$ for complex flavors.

?ptrfs

Refines the solution of a system of linear equations with a symmetric (Hermitian) positive-definite tridiagonal matrix and estimates its error.

```
call sptrfs (n,nrhs,d,e,df,ef,b,ldb,x,ldx,ferr,berr,work,info)
call dptrfs (n,nrhs,d,e,df,ef,b,ldb,x,ldx,ferr,berr,work,info)
call cptrfs (uplo,n,nrhs,d,e,df,ef,b,ldb,x,ldx,ferr,berr,
            work,rwork,info)
call zptrfs (uplo,n,nrhs,d,e,df,ef,b,ldb,x,ldx,ferr,berr,
            work,rwork,info)
```

Discussion

This routine performs an iterative refinement of the solution to a system of linear equations $AX = B$ with a symmetric (Hermitian) positive definite tridiagonal matrix A , with multiple right-hand sides. For each computed solution vector x , the routine computes the *component-wise backward error* β . This error is the smallest relative perturbation in elements of A and b such that x is the exact solution of the perturbed system:

$$|\delta a_{ij}| / |a_{ij}| \leq \beta |a_{ij}|, \quad |\delta b_i| / |b_i| \leq \beta |b_i| \quad \text{such that } (A + \delta A)x = (b + \delta b).$$

Finally, the routine estimates the *component-wise forward error* in the computed solution $\|x - x_e\| / \|x\|_\infty$ (here x_e is the exact solution).

Before calling this routine:

- call the factorization routine [?pttrf](#)
- call the solver routine [?pttrs](#).

Input Parameters

uplo CHARACTER*1. Used for complex flavors only. Must be 'U' or 'L'. Specifies whether the superdiagonal or the subdiagonal of the tridiagonal matrix A is stored and how A is factored:

If *uplo* = 'U', the array *e* stores the superdiagonal of *A*, and *A* is factored as $U^H D U$;

If *uplo* = 'L', the array *e* stores the subdiagonal of *A*, and *A* is factored as LDL^H .

n INTEGER. The order of the matrix *A* ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides ($nrhs \geq 0$).

d,df,rwork REAL for single precision flavors
 DOUBLE PRECISION for double precision flavors
 Arrays: *d*(*n*), *df*(*n*), *rwork*(*n*).
 The array *d* contains the *n* diagonal elements of the tridiagonal matrix *A*.
 The array *df* contains the *n* diagonal elements of the diagonal matrix *D* from the factorization of *A* as computed by [?ptrf](#).
 The array *rwork* is a workspace array used for complex flavors only.

e,ef,b,x,work REAL for [sptrfs](#)
 DOUBLE PRECISION for [dptrfs](#)
 COMPLEX for [cptrfs](#)
 DOUBLE COMPLEX for [zptrfs](#).
 Arrays: *e*(*n* - 1), *ef*(*n* - 1), *b*(*ldb*,*nrhs*), *x*(*ldx*,*nrhs*), *work*(*).
 The array *e* contains the (*n* - 1) off-diagonal elements of the tridiagonal matrix *A* (see *uplo*).
 The array *ef* contains the (*n* - 1) off-diagonal elements of the unit bidiagonal factor *U* or *L* from the factorization computed by [?ptrf](#) (see *uplo*).
 The array *b* contains the matrix *B* whose columns are the right-hand sides for the systems of equations.
 The array *x* contains the solution matrix *X* as computed by [?ptrs](#).
 The array *work* is a workspace array. The dimension of *work* must be at least $2 * n$ for real flavors, and at least *n* for complex flavors.

ldb INTEGER. The leading dimension of *b*; $ldb \geq \max(1, n)$.

ldx INTEGER. The leading dimension of *x*; $ldx \geq \max(1, n)$.

Output Parameters

<i>x</i>	The refined solution matrix <i>X</i> .
<i>ferr, berr</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and backward errors, respectively, for each solution vector.
<i>info</i>	INTEGER . If <i>info</i> = 0, the execution is successful. If <i>info</i> = <i>-i</i> , the <i>i</i> th parameter had an illegal value.

?syrfs

Refines the solution of a system of linear equations with a symmetric matrix and estimates its error.

```
call ssyrfs (uplo,n,nrhs,a,lda,af,ldaf,ipiv,b,ldb,
             x,ldx,ferr,berr,work,iwork,info)
call dsyrfs (uplo,n,nrhs,a,lda,af,ldaf,ipiv,b,ldb,
             x,ldx,ferr,berr,work,iwork,info)

call csyrfs (uplo,n,nrhs,a,lda,af,ldaf,ipiv,b,ldb,
             x,ldx,ferr,berr,work,rwork,info)
call zsyrfs (uplo,n,nrhs,a,lda,af,ldaf,ipiv,b,ldb,
             x,ldx,ferr,berr,work,rwork,info)
```

Discussion

This routine performs an iterative refinement of the solution to a system of linear equations $AX = B$ with a symmetric full-storage matrix A , with multiple right-hand sides. For each computed solution vector x , the routine computes the *component-wise backward error* β . This error is the smallest relative perturbation in elements of A and b such that x is the exact solution of the perturbed system:

$$|\delta a_{ij}| / |a_{ij}| \leq \beta |a_{ij}|, \quad |\delta b_i| / |b_i| \leq \beta |b_i| \quad \text{such that } (A + \delta A)x = (b + \delta b).$$

Finally, the routine estimates the *component-wise forward error* in the computed solution $\|x - x_e\|_\infty / \|x\|_\infty$ (here x_e is the exact solution).

Before calling this routine:

- call the factorization routine [?sytrf](#)
- call the solver routine [?sytrs](#).

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:

If *uplo* = 'U', the array *af* stores the Bunch-Kaufman factorization $A = PUDU^T P^T$.

If *uplo* = 'L', the array *af* stores the Bunch-Kaufman factorization $A = PLDL^T P^T$.

n INTEGER. The order of the matrix *A* ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides ($nrhs \geq 0$).

a, *af*, *b*, *x*, *work* REAL for *ssyrfs*
 DOUBLE PRECISION for *dsyrfs*
 COMPLEX for *csyrfs*
 DOUBLE COMPLEX for *zsyrfs*.

Arrays:

a(*lda*,*) contains the original matrix *A*, as supplied to [?sytrf](#).

af(*ldaf*,*) contains the factored matrix *A*, as returned by [?sytrf](#).

b(*ldb*,*) contains the right-hand side matrix *B*.

x(*ldx*,*) contains the solution matrix *X*.

work(*) is a workspace array.

The second dimension of *a* and *af* must be at least $\max(1, n)$; the second dimension of *b* and *x* must be at least $\max(1, nrhs)$; the dimension of *work* must be at least $\max(1, 3 * n)$ for real flavors and $\max(1, 2 * n)$ for complex flavors.

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, n)$.

ldaf INTEGER. The first dimension of *af*; $ldaf \geq \max(1, n)$.

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

ldx INTEGER. The first dimension of *x*; $ldx \geq \max(1, n)$.

ipiv INTEGER.
 Array, DIMENSION at least $\max(1, n)$.
 The *ipiv* array, as returned by [?sytrf](#).

iwork INTEGER.
 Workspace array, DIMENSION at least $\max(1, n)$.

rwork REAL for *csyrfs*
 DOUBLE PRECISION for *zsyrfs*.
 Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

<i>x</i>	The refined solution matrix X .
<i>ferr</i> , <i>berr</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and backward errors, respectively, for each solution vector.
<i>info</i>	INTEGER . If <i>info</i> = 0, the execution is successful. If <i>info</i> = $-i$, the i th parameter had an illegal value.

Application Notes

The bounds returned in *ferr* are not rigorous, but in practice they almost always overestimate the actual error.

For each right-hand side, computation of the backward error involves a minimum of $4n^2$ floating-point operations (for real flavors) or $16n^2$ operations (for complex flavors). In addition, each step of iterative refinement involves $6n^2$ operations (for real flavors) or $24n^2$ operations (for complex flavors); the number of iterations may range from 1 to 5.

Estimating the forward error involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $2n^2$ floating-point operations for real flavors or $8n^2$ for complex flavors.

?herfs

Refines the solution of a system of linear equations with a complex Hermitian matrix and estimates its error.

```
call cherfs (uplo,n,nrhs,a,lda,af,ldaf,ipiv,b,ldb,
             x,ldx,ferr,berr,work,rwork,info)
call zherfs (uplo,n,nrhs,a,lda,af,ldaf,ipiv,b,ldb,
             x,ldx,ferr,berr,work,rwork,info)
```

Discussion

This routine performs an iterative refinement of the solution to a system of linear equations $AX = B$ with a complex Hermitian full-storage matrix A , with multiple right-hand sides. For each computed solution vector x , the routine computes the *component-wise backward error* β . This error is the smallest relative perturbation in elements of A and b such that x is the exact solution of the perturbed system:

$$|\delta a_{ij}| / |a_{ij}| \leq \beta |a_{ij}|, \quad |\delta b_i| / |b_i| \leq \beta |b_i| \quad \text{such that } (A + \delta A)x = (b + \delta b).$$

Finally, the routine estimates the *component-wise forward error* in the computed solution $\|x - x_e\| / \|x\|$ (here x_e is the exact solution).

Before calling this routine:

- call the factorization routine [?hetrf](#)
- call the solver routine [?hetrs](#).

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If **uplo** = 'U', the array **af** stores the Bunch-Kaufman factorization $A = PUDU^H P^T$.
If **uplo** = 'L', the array **af** stores the Bunch-Kaufman factorization $A = PLDL^H P^T$.

n INTEGER. The order of the matrix A ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides ($nrhs \geq 0$).
a, af, b, x, work COMPLEX for *cherfs*
 DOUBLE COMPLEX for *zherfs*.

Arrays:

a(lda,)* contains the original matrix *A*, as supplied to [?hetrf](#).

af(ldaf,)* contains the factored matrix *A*, as returned by [?hetrf](#).

b(ldb,)* contains the right-hand side matrix *B*.

x(ldx,)* contains the solution matrix *X*.

work()* is a workspace array.

The second dimension of *a* and *af* must be at least $\max(1, n)$; the second dimension of *b* and *x* must be at least $\max(1, nrhs)$; the dimension of *work* must be at least $\max(1, 2 * n)$.

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, n)$.

ldaf INTEGER. The first dimension of *af*; $ldaf \geq \max(1, n)$.

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

ldx INTEGER. The first dimension of *x*; $ldx \geq \max(1, n)$.

ipiv INTEGER.

Array, DIMENSION at least $\max(1, n)$.

The *ipiv* array, as returned by [?hetrf](#).

rwork REAL for *cherfs*

DOUBLE PRECISION for *zherfs*.

Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

<i>x</i>	The refined solution matrix <i>X</i> .
<i>ferr</i> , <i>berr</i>	REAL for <i>cherfs</i> DOUBLE PRECISION for <i>zherfs</i> . Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and backward errors, respectively, for each solution vector.
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = <i>-i</i> , the <i>i</i> th parameter had an illegal value.

Application Notes

The bounds returned in *ferr* are not rigorous, but in practice they almost always overestimate the actual error.

For each right-hand side, computation of the backward error involves a minimum of $16n^2$ operations. In addition, each step of iterative refinement involves $24n^2$ operations; the number of iterations may range from 1 to 5.

Estimating the forward error involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $8n^2$ floating-point operations.

The real counterpart of this routine is *ssyrfs* / *dsyrfs*.

?sprfs

Refines the solution of a system of linear equations with a packed symmetric matrix and estimates the solution error.

```
call ssprfs (uplo,n,nrhs,ap,afp,ipiv,b,ldb,x,ldx,
             ferr,berr,work,iwork,info)
call dsprfs (uplo,n,nrhs,ap,afp,ipiv,b,ldb,x,ldx,
             ferr,berr,work,iwork,info)

call csprfs (uplo,n,nrhs,ap,afp,ipiv,b,ldb,x,ldx,
             ferr,berr,work,rwork,info)
call zsprfs (uplo,n,nrhs,ap,afp,ipiv,b,ldb,x,ldx,
             ferr,berr,work,rwork,info)
```

Discussion

This routine performs an iterative refinement of the solution to a system of linear equations $AX = B$ with a packed symmetric matrix A , with multiple right-hand sides. For each computed solution vector x , the routine computes the *component-wise backward error* β . This error is the smallest relative perturbation in elements of A and b such that x is the exact solution of the perturbed system:

$$|\delta a_{ij}| / |a_{ij}| \leq \beta |a_{ij}|, \quad |\delta b_i| / |b_i| \leq \beta |b_i| \quad \text{such that } (A + \delta A)x = (b + \delta b).$$

Finally, the routine estimates the *component-wise forward error* in the computed solution $\|x - x_e\|_\infty / \|x\|_\infty$ (here x_e is the exact solution).

Before calling this routine:

- call the factorization routine [?sptf](#)
- call the solver routine [?sptrs](#).

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:

If *uplo* = 'U', the array *afp* stores the packed Bunch-Kaufman factorization $A = PUDU^T P^T$.

If *uplo* = 'L', the array *afp* stores the packed Bunch-Kaufman factorization $A = PLDL^T P^T$.

n INTEGER. The order of the matrix *A* ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides ($nrhs \geq 0$).

ap, *afp*, *b*, *x*, *work* REAL for *ssprfs*
 DOUBLE PRECISION for *dsprfs*
 COMPLEX for *csprfs*
 DOUBLE COMPLEX for *zsprfs*.

Arrays:

ap(*) contains the original packed matrix *A*, as supplied to [?sprtf](#).

afp(*) contains the factored packed matrix *A*, as returned by [?sprtf](#).

b(*ldb*,*) contains the right-hand side matrix *B*.

x(*ldx*,*) contains the solution matrix *X*.

work(*) is a workspace array.

The dimension of arrays *ap* and *afp* must be at least $\max(1, n(n+1)/2)$; the second dimension of *b* and *x* must be at least $\max(1, nrhs)$; the dimension of *work* must be at least $\max(1, 3*n)$ for real flavors and $\max(1, 2*n)$ for complex flavors..

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

ldx INTEGER. The first dimension of *x*; $ldx \geq \max(1, n)$.

ipiv INTEGER.

Array, DIMENSION at least $\max(1, n)$.

The *ipiv* array, as returned by [?sprtf](#).

iwork INTEGER.

Workspace array, DIMENSION at least $\max(1, n)$.

rwork REAL for *csprfs*

DOUBLE PRECISION for *zsprfs*

Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

<i>x</i>	The refined solution matrix X .
<i>ferr</i> , <i>berr</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and backward errors, respectively, for each solution vector.
<i>info</i>	INTEGER . If <i>info</i> = 0, the execution is successful. If <i>info</i> = $-i$, the <i>i</i> th parameter had an illegal value.

Application Notes

The bounds returned in *ferr* are not rigorous, but in practice they almost always overestimate the actual error.

For each right-hand side, computation of the backward error involves a minimum of $4n^2$ floating-point operations (for real flavors) or $16n^2$ operations (for complex flavors). In addition, each step of iterative refinement involves $6n^2$ operations (for real flavors) or $24n^2$ operations (for complex flavors); the number of iterations may range from 1 to 5.

Estimating the forward error involves solving a number of systems of linear equations $Ax = b$; the number of systems is usually 4 or 5 and never more than 11. Each solution requires approximately $2n^2$ floating-point operations for real flavors or $8n^2$ for complex flavors.

?hprfs

Refines the solution of a system of linear equations with a packed complex Hermitian matrix and estimates the solution error.

```
call chprfs (uplo, n, nrhs, ap, afp, ipiv, b, ldb, x, ldx,
             ferr, berr, work, rwork, info)
call zhprfs (uplo, n, nrhs, ap, afp, ipiv, b, ldb, x, ldx,
             ferr, berr, work, rwork, info)
```

Discussion

This routine performs an iterative refinement of the solution to a system of linear equations $AX = B$ with a packed complex Hermitian matrix A , with multiple right-hand sides. For each computed solution vector x , the routine computes the *component-wise backward error* β . This error is the smallest relative perturbation in elements of A and b such that x is the exact solution of the perturbed system:

$$|\delta a_{ij}| / |a_{ij}| \leq \beta |a_{ij}|, \quad |\delta b_i| / |b_i| \leq \beta |b_i| \text{ such that } (A + \delta A)x = (b + \delta b).$$

Finally, the routine estimates the *component-wise forward error* in the computed solution $\|x - x_e\|_\infty / \|x\|_\infty$ (here x_e is the exact solution).

Before calling this routine:

- call the factorization routine [?hptrf](#)
- call the solver routine [?hptrs](#).

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If **uplo** = 'U', the array **afp** stores the packed Bunch-Kaufman factorization $A = PUDU^H P^T$.
If **uplo** = 'L', the array **afp** stores the packed Bunch-Kaufman factorization $A = PLDL^H P^T$.

n INTEGER. The order of the matrix A ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides ($nrhs \geq 0$).

ap, *afp*, *b*, *x*, *work* COMPLEX for *chprfs*
 DOUBLE COMPLEX for *zhprfs*.

Arrays:

ap(*) contains the original packed matrix *A*, as supplied to [?hptraf](#).

afp(*) contains the factored packed matrix *A*, as returned by [?hptraf](#).

b(*ldb*,*) contains the right-hand side matrix *B*.

x(*ldx*,*) contains the solution matrix *X*.

work(*) is a workspace array.

The dimension of arrays *ap* and *afp* must be at least $\max(1, n(n+1)/2)$; the second dimension of *b* and *x* must be at least $\max(1, nrhs)$; the dimension of *work* must be at least $\max(1, 2 * n)$.

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

ldx INTEGER. The first dimension of *x*; $ldx \geq \max(1, n)$.

ipiv INTEGER.

Array, DIMENSION at least $\max(1, n)$.
 The *ipiv* array, as returned by [?hptraf](#).

rwork REAL for *chprfs*
 DOUBLE PRECISION for *zhprfs*
 Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

x The refined solution matrix *X*.

ferr, *berr* REAL for *chprfs*.

DOUBLE PRECISION for *zhprfs*.

Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and backward errors, respectively, for each solution vector.

info INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The bounds returned in *ferr* are not rigorous, but in practice they almost always overestimate the actual error.

For each right-hand side, computation of the backward error involves a minimum of $16n^2$ operations. In addition, each step of iterative refinement involves $24n^2$ operations; the number of iterations may range from 1 to 5.

Estimating the forward error involves solving a number of systems of linear equations $Ax = b$; the number is usually 4 or 5 and never more than 11. Each solution requires approximately $8n^2$ floating-point operations.

The real counterpart of this routine is *ssprfs* / *dsprfs*.

?trrfs

Estimates the error in the solution of a system of linear equations with a triangular matrix.

```
call strrfs (uplo,trans,diag,n,nrhs,a,lda,b,ldb,
            x,ldx,ferr,berr,work,iwork,info)
```

```
call dtrrfs (uplo,trans,diag,n,nrhs,a,lda,b,ldb,
            x,ldx,ferr,berr,work,iwork,info)
```

```
call ctrrfs (uplo,trans,diag,n,nrhs,a,lda,b,ldb,
            x,ldx,ferr,berr,work,rwork,info)
```

```
call ztrrfs (uplo,trans,diag,n,nrhs,a,lda,b,ldb,
            x,ldx,ferr,berr,work,rwork,info)
```

Discussion

This routine estimates the errors in the solution to a system of linear equations $AX = B$ or $A^T X = B$ or $A^H X = B$ with a triangular matrix A , with multiple right-hand sides. For each computed solution vector x , the routine computes the *component-wise backward error* β . This error is the smallest relative perturbation in elements of A and b such that x is the exact solution of the perturbed system:

$$|\delta a_{ij}| / |a_{ij}| \leq \beta |a_{ij}|, \quad |\delta b_i| / |b_i| \leq \beta |b_i| \quad \text{such that } (A + \delta A)x = (b + \delta b).$$

The routine also estimates the *component-wise forward error* in the computed solution $\| |x - x_e| | \infty / \|x\| \infty$ (here x_e is the exact solution).

Before calling this routine, call the solver routine [?trtrs](#).

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates whether A is upper or lower triangular:
If *uplo* = 'U', then A is upper triangular.
If *uplo* = 'L', then A is lower triangular.

trans CHARACTER*1. Must be 'N' or 'T' or 'C'.
Indicates the form of the equations:
If *trans* = 'N', the system has the form $AX = B$.
If *trans* = 'T', the system has the form $A^T X = B$.
If *trans* = 'C', the system has the form $A^H X = B$.

diag CHARACTER*1. Must be 'N' or 'U'.
If *diag* = 'N', then A is not a unit triangular matrix.
If *diag* = 'U', then A is unit triangular: diagonal elements of A are assumed to be 1 and not referenced in the array *a*.

n INTEGER. The order of the matrix A ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides ($nrhs \geq 0$).

a, b, x, work REAL for *strrfs*
DOUBLE PRECISION for *dtrrfs*
COMPLEX for *ctrdfs*
DOUBLE COMPLEX for *ztrdfs*.

Arrays:

a(lda,)* contains the upper or lower triangular matrix A, as specified by *uplo*.

b(ldb,)* contains the right-hand side matrix B.

x(ldx,)* contains the solution matrix X.

work()* is a workspace array.

The second dimension of *a* must be at least $\max(1, n)$; the second dimension of *b* and *x* must be at least $\max(1, nrhs)$; the dimension of *work* must be at least $\max(1, 3 * n)$ for real flavors and $\max(1, 2 * n)$ for complex flavors.

<i>lda</i>	INTEGER. The first dimension of <i>a</i> ; $lda \geq \max(1, n)$.
<i>ldb</i>	INTEGER. The first dimension of <i>b</i> ; $ldb \geq \max(1, n)$.
<i>ldx</i>	INTEGER. The first dimension of <i>x</i> ; $ldx \geq \max(1, n)$.
<i>iwork</i>	INTEGER. Workspace array, DIMENSION at least $\max(1, n)$.
<i>rwork</i>	REAL for <i>ctrdfs</i> DOUBLE PRECISION for <i>ztrdfs</i> Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

<i>ferr, berr</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and backward errors, respectively, for each solution vector.
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value.

Application Notes

The bounds returned in *ferr* are not rigorous, but in practice they almost always overestimate the actual error.

A call to this routine involves, for each right-hand side, solving a number of systems of linear equations $Ax = b$; the number of systems is usually 4 or 5 and never more than 11. Each solution requires approximately n^2 floating-point operations for real flavors or $4n^2$ for complex flavors.

?tprfs

Estimates the error in the solution of a system of linear equations with a packed triangular matrix.

```
call stprfs (uplo,trans,diag,n,nrhs,ap,b,ldb,
            x,ldx,ferr,berr,work,iwork,info)
```

```
call dtprfs (uplo,trans,diag,n,nrhs,ap,b,ldb,
            x,ldx,ferr,berr,work,iwork,info)
```

```
call ctprfs (uplo,trans,diag,n,nrhs,ap,b,ldb,
            x,ldx,ferr,berr,work,rwork,info)
```

```
call ztprfs (uplo,trans,diag,n,nrhs,ap,b,ldb,
            x,ldx,ferr,berr,work,rwork,info)
```

Discussion

This routine estimates the errors in the solution to a system of linear equations $AX = B$ or $A^T X = B$ or $A^H X = B$ with a packed triangular matrix A , with multiple right-hand sides. For each computed solution vector x , the routine computes the *component-wise backward error* β . This error is the smallest relative perturbation in elements of A and b such that x is the exact solution of the perturbed system:

$$|\delta a_{ij}| / |a_{ij}| \leq \beta |a_{ij}|, \quad |\delta b_i| / |b_i| \leq \beta |b_i| \quad \text{such that } (A + \delta A)x = (b + \delta b).$$

The routine also estimates the *component-wise forward error* in the computed solution $\|x - x_e\|_\infty / \|x\|_\infty$ (here x_e is the exact solution).

Before calling this routine, call the solver routine [?tprts](#).

Input Parameters

<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. Indicates whether <i>A</i> is upper or lower triangular: If <i>uplo</i> = 'U', then <i>A</i> is upper triangular. If <i>uplo</i> = 'L', then <i>A</i> is lower triangular.
<i>trans</i>	CHARACTER*1. Must be 'N' or 'T' or 'C'. Indicates the form of the equations: If <i>trans</i> = 'N', the system has the form $AX = B$. If <i>trans</i> = 'T', the system has the form $A^T X = B$. If <i>trans</i> = 'C', the system has the form $A^H X = B$.
<i>diag</i>	CHARACTER*1. Must be 'N' or 'U'. If <i>diag</i> = 'N', <i>A</i> is not a unit triangular matrix. If <i>diag</i> = 'U', <i>A</i> is unit triangular: diagonal elements of <i>A</i> are assumed to be 1 and not referenced in the array <i>ap</i> .
<i>n</i>	INTEGER. The order of the matrix <i>A</i> ($n \geq 0$).
<i>nrhs</i>	INTEGER. The number of right-hand sides ($nrhs \geq 0$).
<i>ap, b, x, work</i>	REAL for <i>strrfs</i> DOUBLE PRECISION for <i>dtrrfs</i> COMPLEX for <i>ctrdfs</i> DOUBLE COMPLEX for <i>ztrrfs</i> .

Arrays:

ap(*) contains the upper or lower triangular matrix *A*, as specified by *uplo*.

b(*ldb*,*) contains the right-hand side matrix *B*.

x(*ldx*,*) contains the solution matrix *X*.

work(*) is a workspace array.

The dimension of *ap* must be at least $\max(1, n(n+1)/2)$; the second dimension of *b* and *x* must be at least $\max(1, nrhs)$; the dimension of *work* must be at least $\max(1, 3*n)$ for real flavors and $\max(1, 2*n)$ for complex flavors.

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

ldx INTEGER. The first dimension of *x*; $ldx \geq \max(1, n)$.

iwork INTEGER.

Workspace array, DIMENSION at least $\max(1, n)$.

rwork REAL for `ctrdfs`
DOUBLE PRECISION for `ztrdfs`
Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

ferr, berr REAL for single precision flavors.
DOUBLE PRECISION for double precision flavors.
Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and backward errors, respectively, for each solution vector.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The bounds returned in *ferr* are not rigorous, but in practice they almost always overestimate the actual error.

A call to this routine involves, for each right-hand side, solving a number of systems of linear equations $Ax = b$; the number of systems is usually 4 or 5 and never more than 11. Each solution requires approximately n^2 floating-point operations for real flavors or $4n^2$ for complex flavors.

?tbrfs

Estimates the error in the solution of a system of linear equations with a triangular band matrix.

```
call stbrfs (uplo,trans,diag,n,kd,nrhs,ab,ldab,b,ldb,
            x,ldx,ferr,berr,work,iwork,info)
```

```
call dtbrfs (uplo,trans,diag,n,kd,nrhs,ab,ldab,b,ldb,
            x,ldx,ferr,berr,work,iwork,info)
```

```
call ctbrfs (uplo,trans,diag,n,kd,nrhs,ab,ldab,b,ldb,
            x,ldx,ferr,berr,work,rwork,info)
```

```
call ztbrfs (uplo,trans,diag,n,kd,nrhs,ab,ldab,b,ldb,
            x,ldx,ferr,berr,work,rwork,info)
```

Discussion

This routine estimates the errors in the solution to a system of linear equations $AX = B$ or $A^T X = B$ or $A^H X = B$ with a triangular band matrix A , with multiple right-hand sides. For each computed solution vector x , the routine computes the *component-wise backward error* β . This error is the smallest relative perturbation in elements of A and b such that x is the exact solution of the perturbed system:

$$|\delta a_{ij}| / |a_{ij}| \leq \beta |a_{ij}|, \quad |\delta b_i| / |b_i| \leq \beta |b_i| \quad \text{such that } (A + \delta A)x = (b + \delta b).$$

The routine also estimates the *component-wise forward error* in the computed solution $\|x - x_e\|_\infty / \|x\|_\infty$ (here x_e is the exact solution).

Before calling this routine, call the solver routine [?tbtrs](#).

Input Parameters

<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. Indicates whether <i>A</i> is upper or lower triangular: If <i>uplo</i> = 'U', then <i>A</i> is upper triangular. If <i>uplo</i> = 'L', then <i>A</i> is lower triangular.
<i>trans</i>	CHARACTER*1. Must be 'N' or 'T' or 'C'. Indicates the form of the equations: If <i>trans</i> = 'N', the system has the form $AX = B$. If <i>trans</i> = 'T', the system has the form $A^T X = B$. If <i>trans</i> = 'C', the system has the form $A^H X = B$.
<i>diag</i>	CHARACTER*1. Must be 'N' or 'U'. If <i>diag</i> = 'N', <i>A</i> is not a unit triangular matrix. If <i>diag</i> = 'U', <i>A</i> is unit triangular: diagonal elements of <i>A</i> are assumed to be 1 and not referenced in the array <i>ab</i> .
<i>n</i>	INTEGER. The order of the matrix <i>A</i> ($n \geq 0$).
<i>kd</i>	INTEGER. The number of super-diagonals or sub-diagonals in the matrix <i>A</i> ($kd \geq 0$).
<i>nrhs</i>	INTEGER. The number of right-hand sides ($nrhs \geq 0$).
<i>ab, b, x, work</i>	REAL for <i>stbrfs</i> DOUBLE PRECISION for <i>dtbrfs</i> COMPLEX for <i>ctbrfs</i> DOUBLE COMPLEX for <i>ztbrfs</i> .

Arrays:

ab(*ldab*,*) contains the upper or lower triangular matrix *A*, as specified by *uplo*, in band storage format.

b(*ldb*,*) contains the right-hand side matrix *B*.

x(*ldx*,*) contains the solution matrix *X*.

work(*) is a workspace array.

The second dimension of *a* must be at least $\max(1, n)$;
the second dimension of *b* and *x* must be at least $\max(1, nrhs)$.

The dimension of *work* must be at least $\max(1, 3*n)$ for real flavors and $\max(1, 2*n)$ for complex flavors.

ldab INTEGER. The first dimension of the array *ab*.
(*ldab* $\geq kd + 1$).

<i>ldb</i>	INTEGER. The first dimension of <i>b</i> ; $ldb \geq \max(1, n)$.
<i>ldx</i>	INTEGER. The first dimension of <i>x</i> ; $ldx \geq \max(1, n)$.
<i>iwork</i>	INTEGER. Workspace array, DIMENSION at least $\max(1, n)$.
<i>rwork</i>	REAL for <i>ctbrfs</i> DOUBLE PRECISION for <i>ztbrfs</i> Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

<i>ferr</i> , <i>berr</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and backward errors, respectively, for each solution vector.
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value.

Application Notes

The bounds returned in *ferr* are not rigorous, but in practice they almost always overestimate the actual error.

A call to this routine involves, for each right-hand side, solving a number of systems of linear equations $Ax = b$; the number of systems is usually 4 or 5 and never more than 11. Each solution requires approximately $2n \cdot kd$ floating-point operations for real flavors or $8n \cdot kd$ operations for complex flavors.

Routines for Matrix Inversion

It is seldom necessary to compute an explicit inverse of a matrix.

In particular, do not attempt to solve a system of equations $Ax = b$ by first computing A^{-1} and then forming the matrix-vector product $x = A^{-1}b$.

Call a solver routine instead (see [Routines for Solving Systems of Linear Equations](#)); this is more efficient and more accurate.

However, matrix inversion routines are provided for the rare occasions when an explicit inverse matrix is needed.

?getri

Computes the inverse of an LU-factored general matrix.

```
call sgetri (n, a, lda, ipiv, work, lwork, info)
call dgetri (n, a, lda, ipiv, work, lwork, info)
call cgetri (n, a, lda, ipiv, work, lwork, info)
call zgetri (n, a, lda, ipiv, work, lwork, info)
```

Discussion

This routine computes the inverse (A^{-1}) of a general matrix A .

Before calling this routine, call [?getrf](#) to factorize A .

Input Parameters

n **INTEGER**. The order of the matrix A ($n \geq 0$).

$a, work$ **REAL** for `sgetri`
DOUBLE PRECISION for `dgetri`
COMPLEX for `cgetri`
DOUBLE COMPLEX for `zgetri`.

Arrays: $a(lda, *)$, $work(lwork)$.
 $a(lda, *)$ contains the factorization of the matrix A , as returned by [?getrf](#): $A = PLU$.
The second dimension of a must be at least $\max(1, n)$.
 $work(lwork)$ is a workspace array.

<i>lda</i>	INTEGER. The first dimension of <i>a</i> ; $lda \geq \max(1, n)$.
<i>ipiv</i>	INTEGER. Array, DIMENSION at least $\max(1, n)$. The <i>ipiv</i> array, as returned by <code>?getrf</code> .
<i>lwork</i>	INTEGER. The size of the <i>work</i> array ($lwork \geq n$) See <i>Application notes</i> for the suggested value of <i>lwork</i> .

Output Parameters

<i>a</i>	Overwritten by the <i>n</i> by <i>n</i> matrix A^{-1} .
<i>work(1)</i>	If <i>info</i> = 0, on exit <i>work(1)</i> contains the minimum value of <i>lwork</i> required for optimum performance. Use this <i>lwork</i> for subsequent runs.
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = <i>-i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = <i>i</i> , the <i>i</i> th diagonal element of the factor <i>U</i> is zero, <i>U</i> is singular, and the inversion could not be completed.

Application Notes

For better performance, try using $lwork = n * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*.

If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

The computed inverse *X* satisfies the following error bound:

$$|XA - I| \leq c(n) \varepsilon |X|P|L||U|$$

where $c(n)$ is a modest linear function of *n*; ε is the machine precision; *I* denotes the identity matrix; *P*, *L*, and *U* are the factors of the matrix factorization $A = PLU$.

The total number of floating-point operations is approximately $(4/3)n^3$ for real flavors and $(16/3)n^3$ for complex flavors.

?potri

Computes the inverse of a symmetric
(Hermitian) positive-definite matrix.

```
call spotri (uplo, n, a, lda, info)
call dpotri (uplo, n, a, lda, info)
call cpotri (uplo, n, a, lda, info)
call zpotri (uplo, n, a, lda, info)
```

Discussion

This routine computes the inverse (A^{-1}) of a symmetric positive definite or, for complex flavors, Hermitian positive-definite matrix A . Before calling this routine, call [?potrf](#) to factorize A .

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If **uplo** = 'U', the array **a** stores the factor U of the Cholesky factorization $A = U^H U$.
If **uplo** = 'L', the array **a** stores the factor L of the Cholesky factorization $A = L L^H$.

n INTEGER. The order of the matrix A ($n \geq 0$).

a REAL for **spotri**
DOUBLE PRECISION for **dpotri**
COMPLEX for **cpotri**
DOUBLE COMPLEX for **zpotri**.
Array: **a(lda,*)**.
Contains the factorization of the matrix A , as returned by [?potrf](#).
The second dimension of **a** must be at least $\max(1, n)$.

lda INTEGER. The first dimension of **a**; $lda \geq \max(1, n)$.

Output Parameters

<i>a</i>	Overwritten by the <i>n</i> by <i>n</i> matrix A^{-1} .
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = <i>i</i> , the <i>i</i> th diagonal element of the Cholesky factor (and hence the factor itself) is zero, and the inversion could not be completed.

Application Notes

The computed inverse X satisfies the following error bounds:

$$\|XA - I\|_2 \leq c(n)\varepsilon\kappa_2(A), \quad \|AX - I\|_2 \leq c(n)\varepsilon\kappa_2(A)$$

where $c(n)$ is a modest linear function of n , and ε is the machine precision; I denotes the identity matrix.

The 2-norm $\|A\|_2$ of a matrix A is defined by $\|A\|_2 = \max_{x,x=1} (Ax \cdot Ax)^{1/2}$, and the condition number $\kappa_2(A)$ is defined by $\kappa_2(A) = \|A\|_2 \|A^{-1}\|_2$.

The total number of floating-point operations is approximately $(2/3)n^3$ for real flavors and $(8/3)n^3$ for complex flavors.

?pptri

Computes the inverse of a packed symmetric (Hermitian) positive-definite matrix

```
call spptri (uplo, n, ap, info)
call dpptri (uplo, n, ap, info)
call cpptri (uplo, n, ap, info)
call zpptri (uplo, n, ap, info)
```

Discussion

This routine computes the inverse (A^{-1}) of a symmetric positive definite or, for complex flavors, Hermitian positive-definite matrix A in *packed* form. Before calling this routine, call [?pptrf](#) to factorize A .

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If $uplo = 'U'$, the array ap stores the packed factor U of the Cholesky factorization $A = U^H U$.
If $uplo = 'L'$, the array ap stores the packed factor L of the Cholesky factorization $A = LL^H$.

n INTEGER. The order of the matrix A ($n \geq 0$).

ap REAL for `spptri`
DOUBLE PRECISION for `dpptri`
COMPLEX for `cpptri`
DOUBLE COMPLEX for `zpptri`.
Array, DIMENSION at least $\max(1, n(n+1)/2)$.
Contains the factorization of the packed matrix A , as returned by [?pptrf](#).
The dimension ap must be at least $\max(1, n(n+1)/2)$.

Output Parameters

<i>ap</i>	Overwritten by the packed n by n matrix A^{-1} .
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = $-i$, the i th parameter had an illegal value. If <i>info</i> = i , the i th diagonal element of the Cholesky factor (and hence the factor itself) is zero, and the inversion could not be completed.

Application Notes

The computed inverse X satisfies the following error bounds:

$$\|XA - I\|_2 \leq c(n)\varepsilon \kappa_2(A), \quad \|AX - I\|_2 \leq c(n)\varepsilon \kappa_2(A)$$

where $c(n)$ is a modest linear function of n , and ε is the machine precision; I denotes the identity matrix.

The 2-norm $\|A\|_2$ of a matrix A is defined by $\|A\|_2 = \max_{x,x=1} (Ax \cdot Ax)^{1/2}$, and the condition number $\kappa_2(A)$ is defined by $\kappa_2(A) = \|A\|_2 \|A^{-1}\|_2$.

The total number of floating-point operations is approximately $(2/3)n^3$ for real flavors and $(8/3)n^3$ for complex flavors.

?sytri

Computes the inverse of a symmetric matrix.

```
call ssytri (uplo, n, a, lda, ipiv, work, info)
call dsytri (uplo, n, a, lda, ipiv, work, info)
call csytri (uplo, n, a, lda, ipiv, work, info)
call zsytri (uplo, n, a, lda, ipiv, work, info)
```

Discussion

This routine computes the inverse (A^{-1}) of a symmetric matrix A . Before calling this routine, call [?sytrf](#) to factorize A .

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If *uplo* = 'U', the array *a* stores the Bunch-Kaufman factorization $A = PUDU^T P^T$.
If *uplo* = 'L', the array *a* stores the Bunch-Kaufman factorization $A = PLDL^T P^T$.

n INTEGER. The order of the matrix A ($n \geq 0$).

a, *work* REAL for *ssytri*
DOUBLE PRECISION for *dsytri*
COMPLEX for *csytri*
DOUBLE COMPLEX for *zsytri*.
Arrays:
a(*lda*,*) contains the factorization of the matrix A , as returned by [?sytrf](#).
The second dimension of *a* must be at least $\max(1, n)$.
work(*) is a workspace array.
The dimension of *work* must be at least $\max(1, 2*n)$.

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, n)$.

ipiv **INTEGER.**
 Array, **DIMENSION** at least $\max(1, n)$.
 The *ipiv* array, as returned by [?sytrf](#).

Output Parameters

a Overwritten by the *n* by *n* matrix A^{-1} .

info **INTEGER.**
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.
 If *info* = *i*, the *i*th diagonal element of *D* is zero, *D* is singular, and the inversion could not be completed.

Application Notes

The computed inverse *X* satisfies the following error bounds:

$$|DU^T P^T X P U - I| \leq c(n) \epsilon (|D| |U^T| P^T |X| P |U| + |D| |D^{-1}|)$$

for *uplo* = 'U', and

$$|DL^T P^T X P L - I| \leq c(n) \epsilon (|D| |L^T| P^T |X| P |L| + |D| |D^{-1}|)$$

for *uplo* = 'L'. Here $c(n)$ is a modest linear function of *n*, and ϵ is the machine precision; *I* denotes the identity matrix.

The total number of floating-point operations is approximately $(2/3)n^3$ for real flavors and $(8/3)n^3$ for complex flavors.

?hetri

Computes the inverse of a complex Hermitian matrix.

```
call chetri (uplo, n, a, lda, ipiv, work, info)
call zhetri (uplo, n, a, lda, ipiv, work, info)
```

Discussion

This routine computes the inverse (A^{-1}) of a complex Hermitian matrix A . Before calling this routine, call [?hetrf](#) to factorize A .

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If *uplo* = 'U', the array *a* stores the Bunch-Kaufman factorization $A = PUDU^H P^T$.
If *uplo* = 'L', the array *a* stores the Bunch-Kaufman factorization $A = PLDL^H P^T$.

n INTEGER. The order of the matrix A ($n \geq 0$).

a, work COMPLEX for *chetri*
DOUBLE COMPLEX for *zhetri*.
Arrays:
a(lda,)* contains the factorization of the matrix A , as returned by [?hetrf](#).
The second dimension of *a* must be at least $\max(1,n)$.
work()* is a workspace array.
The dimension of *work* must be at least $\max(1,n)$.

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, n)$.

ipiv INTEGER.
Array, DIMENSION at least $\max(1,n)$.
The *ipiv* array, as returned by [?hetrf](#).

Output Parameters

<i>a</i>	Overwritten by the <i>n</i> by <i>n</i> matrix A^{-1} .
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = <i>i</i> , the <i>i</i> th diagonal element of <i>D</i> is zero, <i>D</i> is singular, and the inversion could not be completed.

Application Notes

The computed inverse *X* satisfies the following error bounds:

$$|DU^H P^T X P U - I| \leq c(n) \varepsilon (|D| |U^H| |P^T| |X| |P| |U| + |D| |D^{-1}|)$$

for *uplo* = 'U', and

$$|DL^H P^T X P L - I| \leq c(n) \varepsilon (|D| |L^H| |P^T| |X| |P| |L| + |D| |D^{-1}|)$$

for *uplo* = 'L'. Here $c(n)$ is a modest linear function of *n*, and ε is the machine precision; *I* denotes the identity matrix.

The total number of floating-point operations is approximately $(2/3)n^3$ for real flavors and $(8/3)n^3$ for complex flavors.

The real counterpart of this routine is `?sytri`.

?sptri

Computes the inverse of a symmetric matrix using packed storage.

```
call ssptri (uplo, n, ap, ipiv, work, info)
call dsptri (uplo, n, ap, ipiv, work, info)
call csptri (uplo, n, ap, ipiv, work, info)
call zsptri (uplo, n, ap, ipiv, work, info)
```

Discussion

This routine computes the inverse (A^{-1}) of a packed symmetric matrix A . Before calling this routine, call [?sptf](#) to factorize A .

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If *uplo* = 'U', the array *ap* stores the Bunch-Kaufman factorization $A = PUDU^T P^T$.
If *uplo* = 'L', the array *ap* stores the Bunch-Kaufman factorization $A = PLDL^T P^T$.

n INTEGER. The order of the matrix A ($n \geq 0$).

ap, work REAL for *ssptri*
DOUBLE PRECISION for *dsptri*
COMPLEX for *csptri*
DOUBLE COMPLEX for *zsptri*.
Arrays:
ap(*) contains the factorization of the matrix A , as returned by [?sptf](#).
The dimension of *ap* must be at least $\max(1, n(n+1)/2)$.
work(*) is a workspace array.
The dimension of *work* must be at least $\max(1, n)$.

ipiv INTEGER.
 Array, DIMENSION at least max(1,*n*).
 The *ipiv* array, as returned by [?sptf](#).

Output Parameters

ap Overwritten by the *n* by *n* matrix A^{-1} in packed form.

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.
 If *info* = *i*, the *i*th diagonal element of *D* is zero, *D* is singular, and the inversion could not be completed.

Application Notes

The computed inverse *X* satisfies the following error bounds:

$$|DU^T P^T X P U - I| \leq c(n) \epsilon (|D| |U^T| P^T |X| P |U| + |D| |D^{-1}|)$$

for *uplo* = 'U', and

$$|DL^T P^T X P L - I| \leq c(n) \epsilon (|D| |L^T| P^T |X| P |L| + |D| |D^{-1}|)$$

for *uplo* = 'L'. Here $c(n)$ is a modest linear function of *n*, and ϵ is the machine precision; *I* denotes the identity matrix.

The total number of floating-point operations is approximately $(2/3)n^3$ for real flavors and $(8/3)n^3$ for complex flavors.

?hptri

Computes the inverse of a complex Hermitian matrix using packed storage.

```
call chptri (uplo, n, ap, ipiv, work, info)
call zhptri (uplo, n, ap, ipiv, work, info)
```

Discussion

This routine computes the inverse (A^{-1}) of a complex Hermitian matrix A using packed storage.

Before calling this routine, call [?hptrf](#) to factorize A .

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates how the input matrix A has been factored:
If *uplo* = 'U', the array *ap* stores the packed Bunch-Kaufman factorization $A = PUDU^H P^T$.
If *uplo* = 'L', the array *ap* stores the packed Bunch-Kaufman factorization $A = PLDL^H P^T$.

n INTEGER. The order of the matrix A ($n \geq 0$).

ap COMPLEX for [chptri](#)
DOUBLE COMPLEX for [zhptri](#).
Arrays:
ap(*) contains the factorization of the matrix A , as returned by [?hptrf](#).
The dimension of *ap* must be at least $\max(1, n(n+1)/2)$.
work(*) is a workspace array.
The dimension of *work* must be at least $\max(1, n)$.

ipiv INTEGER.
Array, DIMENSION at least $\max(1, n)$.
The *ipiv* array, as returned by [?hptrf](#).

Output Parameters

<i>ap</i>	Overwritten by the <i>n</i> by <i>n</i> matrix A^{-1} .
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = <i>i</i> , the <i>i</i> th diagonal element of <i>D</i> is zero, <i>D</i> is singular, and the inversion could not be completed.

Application Notes

The computed inverse *X* satisfies the following error bounds:

$$|DU^H P^T X P U - I| \leq c(n) \varepsilon (|D| |U^H| |P^T| |X| |P| |U| + |D| |D^{-1}|)$$

for *uplo* = 'U', and

$$|DL^H P^T X P L - I| \leq c(n) \varepsilon (|D| |L^H| |P^T| |X| |P| |L| + |D| |D^{-1}|)$$

for *uplo* = 'L'. Here $c(n)$ is a modest linear function of *n*, and ε is the machine precision; *I* denotes the identity matrix.

The total number of floating-point operations is approximately $(2/3)n^3$ for real flavors and $(8/3)n^3$ for complex flavors.

The real counterpart of this routine is `?sptri`.

?trtri

Computes the inverse of a triangular matrix.

```

call strtri (uplo, diag, n, a, lda, info)
call dtrtri (uplo, diag, n, a, lda, info)
call ctrtri (uplo, diag, n, a, lda, info)
call ztrtri (uplo, diag, n, a, lda, info)

```

Discussion

This routine computes the inverse (A^{-1}) of a triangular matrix A .

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates whether A is upper or lower triangular:
If *uplo* = 'U', then A is upper triangular.
If *uplo* = 'L', then A is lower triangular.

diag CHARACTER*1. Must be 'N' or 'U'.
If *diag* = 'N', then A is not a unit triangular matrix.
If *diag* = 'U', A is unit triangular: diagonal elements of A are assumed to be 1 and not referenced in the array *a*.

n INTEGER. The order of the matrix A ($n \geq 0$).

a REAL for strtri
DOUBLE PRECISION for dtrtri
COMPLEX for ctrtri
DOUBLE COMPLEX for ztrtri.
Array: DIMENSION (*lda*, *).
Contains the matrix A .
The second dimension of *a* must be at least $\max(1, n)$.

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, n)$.

Output Parameters

<i>a</i>	Overwritten by the <i>n</i> by <i>n</i> matrix A^{-1} .
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = <i>i</i> , the <i>i</i> th diagonal element of <i>A</i> is zero, <i>A</i> is singular, and the inversion could not be completed.

Application Notes

The computed inverse *X* satisfies the following error bounds:

$$|XA - I| \leq c(n)\epsilon |X||A|$$

$$|X - A^{-1}| \leq c(n)\epsilon |A^{-1}||A||X|$$

where $c(n)$ is a modest linear function of *n*; ϵ is the machine precision; *I* denotes the identity matrix.

The total number of floating-point operations is approximately $(1/3)n^3$ for real flavors and $(4/3)n^3$ for complex flavors.

?tptri

Computes the inverse of a triangular matrix using packed storage.

```
call stptri (uplo, diag, n, ap, info)
call dtptri (uplo, diag, n, ap, info)
call ctptri (uplo, diag, n, ap, info)
call ztptri (uplo, diag, n, ap, info)
```

Discussion

This routine computes the inverse (A^{-1}) of a packed triangular matrix *A*.

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates whether A is upper or lower triangular:
If *uplo* = 'U', then A is upper triangular.
If *uplo* = 'L', then A is lower triangular.

diag CHARACTER*1. Must be 'N' or 'U'.
If *diag* = 'N', then A is not a unit triangular matrix.
If *diag* = 'U', A is unit triangular: diagonal elements of A are assumed to be 1 and not referenced in the array *ap*.

n INTEGER. The order of the matrix A ($n \geq 0$).

ap REAL for *stptri*
DOUBLE PRECISION for *dtptri*
COMPLEX for *ctptri*
DOUBLE COMPLEX for *ztptri*.
Array: DIMENSION at least $\max(1, n(n+1)/2)$.
Contains the packed triangular matrix A .

Output Parameters

ap Overwritten by the packed n by n matrix A^{-1} .

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = $-i$, the i th parameter had an illegal value.
If *info* = i , the i th diagonal element of A is zero, A is singular, and the inversion could not be completed.

Application Notes

The computed inverse X satisfies the following error bounds:

$$|XA - I| \leq c(n)\epsilon |X||A|$$

$$|X - A^{-1}| \leq c(n)\epsilon |A^{-1}||A||X|$$

where $c(n)$ is a modest linear function of n ; ϵ is the machine precision; I denotes the identity matrix.

The total number of floating-point operations is approximately $(1/3)n^3$ for real flavors and $(4/3)n^3$ for complex flavors.

Routines for Matrix Equilibration

Routines described in this section are used to compute scaling factors needed to equilibrate a matrix. Note that these routines do not actually scale the matrices.

?gseequ

Computes row and column scaling factors intended to equilibrate a matrix and reduce its condition number.

```
call sseequ (m, n, a, lda, r, c, rowcnd, colcnd, amax, info)
call dseequ (m, n, a, lda, r, c, rowcnd, colcnd, amax, info)
call cseequ (m, n, a, lda, r, c, rowcnd, colcnd, amax, info)
call zseequ (m, n, a, lda, r, c, rowcnd, colcnd, amax, info)
```

Discussion

This routine computes row and column scalings intended to equilibrate an m -by- n matrix A and reduce its condition number. The output array r returns the row scale factors and the array c the column scale factors. These factors are chosen to try to make the largest element in each row and column of the matrix B with elements $b_{ij}=r(i)*a_{ij}*c(j)$ have absolute value 1.

Input Parameters

m **INTEGER**. The number of rows of the matrix A , $m \geq 0$.

n **INTEGER**. The number of columns of the matrix A , $n \geq 0$.

a **REAL** for **sseequ**
DOUBLE PRECISION for **dseequ**
COMPLEX for **cseequ**
DOUBLE COMPLEX for **zseequ**.

Array: `DIMENSION (lda, *)`.

Contains the m -by- n matrix A whose equilibration factors are to be computed.

The second dimension of a must be at least $\max(1, n)$.

`lda` `INTEGER`. The leading dimension of a ; $lda \geq \max(1, m)$.

Output Parameters

`r, c` `REAL` for single precision flavors;
 `DOUBLE PRECISION` for double precision flavors.
 Arrays: $r(m)$, $c(n)$.
 If $info = 0$, or $info > m$, the array r contains the row scale factors of the matrix A .
 If $info = 0$, the array c contains the column scale factors of the matrix A .

`rowcnd` `REAL` for single precision flavors;
 `DOUBLE PRECISION` for double precision flavors.
 If $info = 0$ or $info > m$, `rowcnd` contains the ratio of the smallest $r(i)$ to the largest $r(i)$.

`colcnd` `REAL` for single precision flavors;
 `DOUBLE PRECISION` for double precision flavors.
 If $info = 0$, `colcnd` contains the ratio of the smallest $c(i)$ to the largest $c(i)$.

`amax` `REAL` for single precision flavors;
 `DOUBLE PRECISION` for double precision flavors.
 Absolute value of the largest element of the matrix A .

`info` `INTEGER`.
 If $info = 0$, the execution is successful.
 If $info = -i$, the i th parameter had an illegal value.
 If $info = i$ and
 $i \leq m$, the i th row of A is exactly zero;
 $i > m$, the $(i-m)$ th column of A is exactly zero.

Application Notes

All the components of r and c are restricted to be between SMLNUM = smallest safe number and BIGNUM = largest safe number. Use of these scaling factors is not guaranteed to reduce the condition number of A but works well in practice.

If $rowcnd \geq 0.1$ and $amax$ is neither too large nor too small, it is not worth scaling by r . If $colcnd \geq 0.1$, it is not worth scaling by c .

If $amax$ is very close to overflow or very close to underflow, the matrix A should be scaled.

?gbequ

Computes row and column scaling factors intended to equilibrate a band matrix and reduce its condition number.

```
call sgbequ (m, n, kl, ku, ab, ldab, r, c, rowcnd, colcnd,amax,info)
call dgbequ (m, n, kl, ku, ab, ldab, r, c, rowcnd, colcnd,amax,info)
call cgbequ (m, n, kl, ku, ab, ldab, r, c, rowcnd, colcnd,amax,info)
call zgbequ (m, n, kl, ku, ab, ldab, r, c, rowcnd, colcnd,amax,info)
```

Discussion

This routine computes row and column scalings intended to equilibrate an m -by- n band matrix A and reduce its condition number. The output array r returns the row scale factors and the array c the column scale factors. These factors are chosen to try to make the largest element in each row and column of the matrix B with elements $b_{ij}=r(i)*a_{ij}*c(j)$ have absolute value 1.

Input Parameters

m **INTEGER**. The number of rows of the matrix A , $m \geq 0$.

n **INTEGER**. The number of columns of the matrix A , $n \geq 0$.

kl **INTEGER**. The number of sub-diagonals within the band of A ($kl \geq 0$).

ku **INTEGER**. The number of super-diagonals within the band of A ($ku \geq 0$).

ab **REAL** for `sgbequ`
DOUBLE PRECISION for `dgbequ`
COMPLEX for `cgbequ`
DOUBLE COMPLEX for `zgbequ`.
 Array, **DIMENSION** ($ldab, *$).
 Contains the original band matrix A stored in rows from 1 to $kl + ku + 1$.

The second dimension of *ab* must be at least $\max(1,n)$;
ldab INTEGER. The leading dimension of *ab*,
 $ldab \geq kl+ku+1$.

Output Parameters

r, c REAL for single precision flavors;
DOUBLE PRECISION for double precision flavors.
Arrays: *r*(*m*), *c*(*n*).
If *info* = 0, or *info* > *m*, the array *r* contains the row
scale factors of the matrix *A*.
If *info* = 0, the array *c* contains the column scale
factors of the matrix *A*.

rowcnd REAL for single precision flavors;
DOUBLE PRECISION for double precision flavors.
If *info* = 0 or *info* > *m*, *rowcnd* contains the ratio of
the smallest *r*(*i*) to the largest *r*(*i*).

colcnd REAL for single precision flavors;
DOUBLE PRECISION for double precision flavors.
If *info* = 0, *colcnd* contains the ratio of the smallest
c(*i*) to the largest *c*(*i*).

amax REAL for single precision flavors;
DOUBLE PRECISION for double precision flavors.
Absolute value of the largest element of the matrix *A*.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i* and
 i ≤ *m*, the *i*th row of *A* is exactly zero;
 i > *m*, the (*i*-*m*)th column of *A* is exactly zero.

Application Notes

All the components of *r* and *c* are restricted to be between SMLNUM = smallest safe number and BIGNUM = largest safe number. Use of these scaling factors is not guaranteed to reduce the condition number of *A* but works well in practice.

If $\text{rowcnd} \geq 0.1$ and amax is neither too large nor too small, it is not worth scaling by r . If $\text{colcnd} \geq 0.1$, it is not worth scaling by c .

If amax is very close to overflow or very close to underflow, the matrix A should be scaled.

?poequ

Computes row and column scaling factors intended to equilibrate a symmetric (Hermitian) positive definite matrix and reduce its condition number.

```
call spoequ (n, a, lda, s, scond, amax, info)
call dpoequ (n, a, lda, s, scond, amax, info)
call cpoequ (n, a, lda, s, scond, amax, info)
call zpoequ (n, a, lda, s, scond, amax, info)
```

Discussion

This routine computes row and column scalings intended to equilibrate a symmetric (Hermitian) positive definite matrix A and reduce its condition number (with respect to the two-norm). The output array s returns scale factors computed as

$$s(i) = \frac{1}{\sqrt{a_{i,i}}}$$

These factors are chosen so that the scaled matrix B with elements $b_{ij} = s(i) * a_{ij} * s(j)$ has diagonal elements equal to 1.

This choice of s puts the condition number of B within a factor n of the smallest possible condition number over all possible diagonal scalings.

Input Parameters

n **INTEGER**. The order of the matrix A , $n \geq 0$.

a REAL for `spoequ`
DOUBLE PRECISION for `dpoequ`
COMPLEX for `cpoequ`
DOUBLE COMPLEX for `zpoequ`.
Array: `DIMENSION (lda, *)`.
Contains the n -by- n symmetric or Hermitian positive definite matrix A whose scaling factors are to be computed. Only diagonal elements of A are referenced. The second dimension of *a* must be at least $\max(1, n)$.

lda INTEGER. The leading dimension of *a*; $lda \geq \max(1, m)$.

Output Parameters

s REAL for single precision flavors;
DOUBLE PRECISION for double precision flavors.
Array, `DIMENSION (n)`.
If *info* = 0, the array *s* contains the scale factors for A .

scond REAL for single precision flavors;
DOUBLE PRECISION for double precision flavors.
If *info* = 0, *scond* contains the ratio of the smallest $s(i)$ to the largest $s(i)$.

amax REAL for single precision flavors;
DOUBLE PRECISION for double precision flavors.
Absolute value of the largest element of the matrix A .

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = $-i$, the i th parameter had an illegal value.
If *info* = i , the i th diagonal element of A is nonpositive.

Application Notes

If $scond \geq 0.1$ and *amax* is neither too large nor too small, it is not worth scaling by *s*.

If *amax* is very close to overflow or very close to underflow, the matrix A should be scaled.

?ppequ

Computes row and column scaling factors intended to equilibrate a symmetric (Hermitian) positive definite matrix in packed storage and reduce its condition number.

```
call sppequ (uplo, n, ap, s, scond, amax, info)
call dppequ (uplo, n, ap, s, scond, amax, info)
call cppequ (uplo, n, ap, s, scond, amax, info)
call zppequ (uplo, n, ap, s, scond, amax, info)
```

Discussion

This routine computes row and column scalings intended to equilibrate a symmetric (Hermitian) positive definite matrix A in packed storage and reduce its condition number (with respect to the two-norm). The output array s returns scale factors computed as

$$s(i) = \frac{1}{\sqrt{a_{i,i}}}$$

These factors are chosen so that the scaled matrix B with elements $b_{ij}=s(i)*a_{ij}*s(j)$ has diagonal elements equal to 1.

This choice of s puts the condition number of B within a factor n of the smallest possible condition number over all possible diagonal scalings.

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates whether the upper or lower triangular part of A is packed in the array ap :
If $uplo = 'U'$, the array ap stores the upper triangular part of the matrix A .
If $uplo = 'L'$, the array ap stores the lower triangular part of the matrix A .

n INTEGER. The order of matrix *A* ($n \geq 0$).

ap REAL for `spequ`
 DOUBLE PRECISION for `dpequ`
 COMPLEX for `cppequ`
 DOUBLE COMPLEX for `zpequ`.
 Array, DIMENSION at least $\max(1, n(n+1)/2)$.
 The array *ap* contains either the upper or the lower triangular part of the matrix *A* (as specified by *uplo*) in *packed storage* (see [Matrix Storage Schemes](#)).

Output Parameters

s REAL for single precision flavors;
 DOUBLE PRECISION for double precision flavors.
 Array, DIMENSION (*n*).
 If *info* = 0, the array *s* contains the scale factors for *A*.

scond REAL for single precision flavors;
 DOUBLE PRECISION for double precision flavors.
 If *info* = 0, *scond* contains the ratio of the smallest *s*(*i*) to the largest *s*(*i*).

amax REAL for single precision flavors;
 DOUBLE PRECISION for double precision flavors.
 Absolute value of the largest element of the matrix *A*.

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = *-i*, the *i*th parameter had an illegal value.
 If *info* = *i*, the *i*th diagonal element of *A* is nonpositive.

Application Notes

If *scond* ≥ 0.1 and *amax* is neither too large nor too small, it is not worth scaling by *s*.

If *amax* is very close to overflow or very close to underflow, the matrix *A* should be scaled.

?pbequ

Computes row and column scaling factors intended to equilibrate a symmetric (Hermitian) positive definite band matrix and reduce its condition number.

```
call spbequ (uplo, n, kd, ab, ldab, s, scond, amax, info)
call dpbequ (uplo, n, kd, ab, ldab, s, scond, amax, info)
call cpbequ (uplo, n, kd, ab, ldab, s, scond, amax, info)
call zpbequ (uplo, n, kd, ab, ldab, s, scond, amax, info)
```

Discussion

This routine computes row and column scalings intended to equilibrate a symmetric (Hermitian) positive definite matrix A in packed storage and reduce its condition number (with respect to the two-norm). The output array s returns scale factors computed as

$$s(i) = \frac{1}{\sqrt{a_{i,i}}}$$

These factors are chosen so that the scaled matrix B with elements $b_{ij} = s(i) * a_{ij} * s(j)$ has diagonal elements equal to 1.

This choice of s puts the condition number of B within a factor n of the smallest possible condition number over all possible diagonal scalings.

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates whether the upper or lower triangular part of A is packed in the array **ab**:
If **uplo** = 'U', the array **ab** stores the upper triangular part of the matrix A .
If **uplo** = 'L', the array **ab** stores the lower triangular part of the matrix A .

n INTEGER. The order of matrix A ($n \geq 0$).

kd INTEGER. The number of super-diagonals or sub-diagonals in the matrix A ($kd \geq 0$).

<i>ab</i>	<p>REAL for <i>spbequ</i> DOUBLE PRECISION for <i>dpbequ</i> COMPLEX for <i>cpbequ</i> DOUBLE COMPLEX for <i>zpbequ</i>. Array, DIMENSION (<i>ldab</i>,*). The array <i>ap</i> contains either the upper or the lower triangular part of the matrix <i>A</i> (as specified by <i>uplo</i>) in <i>band storage</i> (see Matrix Storage Schemes). The second dimension of <i>ab</i> must be at least $\max(1, n)$.</p>
<i>ldab</i>	<p>INTEGER. The leading dimension of the array <i>ab</i>. ($ldab \geq kd + 1$).</p>

Output Parameters

<i>s</i>	<p>REAL for single precision flavors; DOUBLE PRECISION for double precision flavors. Array, DIMENSION (<i>n</i>). If <i>info</i> = 0, the array <i>s</i> contains the scale factors for <i>A</i>.</p>
<i>scond</i>	<p>REAL for single precision flavors; DOUBLE PRECISION for double precision flavors. If <i>info</i> = 0, <i>scond</i> contains the ratio of the smallest <i>s</i>(<i>i</i>) to the largest <i>s</i>(<i>i</i>).</p>
<i>amax</i>	<p>REAL for single precision flavors; DOUBLE PRECISION for double precision flavors. Absolute value of the largest element of the matrix <i>A</i>.</p>
<i>info</i>	<p>INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = <i>-i</i>, the <i>i</i>th parameter had an illegal value. If <i>info</i> = <i>i</i>, the <i>i</i>th diagonal element of <i>A</i> is nonpositive.</p>

Application Notes

If $scond \geq 0.1$ and *amax* is neither too large nor too small, it is not worth scaling by *s*.

If *amax* is very close to overflow or very close to underflow, the matrix *A* should be scaled.

Driver Routines

[Table 4-3](#) lists the LAPACK driver routines for solving systems of linear equations with real or complex matrices.

Table 4-3 Driver Routines for Solving Systems of Linear Equations

Matrix type, storage scheme	Simple Driver	Expert Driver
general	?gesv	?gesvx
general band	?gbsv	?gbsvx
general tridiagonal	?gtsv	?gtsvx
symmetric/Hermitian positive-definite	?posv	?posvx
symmetric/Hermitian positive-definite, packed storage	?ppsv	?ppsvx
symmetric/Hermitian positive-definite, band	?pbsv	?pbsvx
symmetric/Hermitian positive-definite, tridiagonal	?ptsv	?ptsvx
symmetric/Hermitian indefinite	?sysv / ?hesv	?sysvx / ?hesvx
symmetric/Hermitian indefinite, packed storage	?spsv / ?hpsv	?spsvx / ?hpsvx
complex symmetric	?sysv	?sysvx
complex symmetric, packed storage	?spsv	?spsvx

In this table ? stands for **s** (single precision real), **d** (double precision real), **c** (single precision complex), or **z** (double precision complex).

?gesv

Computes the solution to the system of linear equations with a square matrix A and multiple right-hand sides.

```
call sgesv (n, nrhs, a, lda, ipiv, b, ldb, info)
call dgesv (n, nrhs, a, lda, ipiv, b, ldb, info)
call cgesv (n, nrhs, a, lda, ipiv, b, ldb, info)
call zgesv (n, nrhs, a, lda, ipiv, b, ldb, info)
```

Discussion

This routine solves for X the system of linear equations $AX = B$, where A is an n -by- n matrix, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

The LU decomposition with partial pivoting and row interchanges is used to factor A as $A = PLU$, where P is a permutation matrix, L is unit lower triangular, and U is upper triangular. The factored form of A is then used to solve the system of equations $AX = B$.

Input Parameters

n **INTEGER.** The order of A ; the number of rows in B ($n \geq 0$).

$nrhs$ **INTEGER.** The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).

a, b **REAL** for **sgesv**
DOUBLE PRECISION for **dgesv**
COMPLEX for **cgesv**
DOUBLE COMPLEX for **zgesv**.
 Arrays: $a(lda, *)$, $b(ldb, *)$.
 The array a contains the matrix A .
 The array b contains the matrix B whose columns are the right-hand sides for the systems of equations.
 The second dimension of a must be at least $\max(1, n)$, the second dimension of b at least $\max(1, nrhs)$.

lda **INTEGER**. The first dimension of *a*; $lda \geq \max(1, n)$.
ldb **INTEGER**. The first dimension of *b*; $ldb \geq \max(1, n)$.

Output Parameters

a Overwritten by the factors *L* and *U* from the factorization of $A = P L U$; the unit diagonal elements of *L* are not stored .
b Overwritten by the solution matrix *X*.
ipiv **INTEGER**.
 Array, **DIMENSION** at least $\max(1, n)$.
 The pivot indices that define the permutation matrix *P*; row *i* of the matrix was interchanged with row *ipiv*(*i*).
info **INTEGER**. If *info*=0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.
 If *info* = *i*, $U(i, i)$ is exactly zero. The factorization has been completed, but the factor *U* is exactly singular, so the solution could not be computed.

?gesvx

Computes the solution to the system of linear equations with a square matrix A and multiple right-hand sides, and provides error bounds on the solution.

```
call sgesvx (fact, trans, n, nrhs, a, lda, af, ldaf, ipiv, equed, r,
            c, b, ldb, x, ldx, rcond, ferr, berr, work, iwork, info)
call dgesvx (fact, trans, n, nrhs, a, lda, af, ldaf, ipiv, equed, r,
            c, b, ldb, x, ldx, rcond, ferr, berr, work, iwork, info)
call cgesvx (fact, trans, n, nrhs, a, lda, af, ldaf, ipiv, equed, r,
            c, b, ldb, x, ldx, rcond, ferr, berr, work, rwork, info)
call zgesvx (fact, trans, n, nrhs, a, lda, af, ldaf, ipiv, equed, r,
            c, b, ldb, x, ldx, rcond, ferr, berr, work, rwork, info)
```

Discussion

This routine uses the LU factorization to compute the solution to a real or complex system of linear equations $AX = B$, where A is an n -by- n matrix, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

Error bounds on the solution and a condition estimate are also provided.

The routine `?gesvx` performs the following steps:

1. If `fact = 'E'`, real scaling factors `r` and `c` are computed to equilibrate the system:

$$\text{trans} = \text{'N'}: \quad \text{diag}(r) * A * \text{diag}(c) * \text{diag}(c)^{-1} * X = \text{diag}(r) * B$$

$$\text{trans} = \text{'T'}: \quad (\text{diag}(r) * A * \text{diag}(c))^T * \text{diag}(r)^{-1} * X = \text{diag}(c) * B$$

$$\text{trans} = \text{'C'}: \quad (\text{diag}(r) * A * \text{diag}(c))^H * \text{diag}(r)^{-1} * X = \text{diag}(c) * B$$

Whether or not the system will be equilibrated depends on the scaling of the matrix A , but if equilibration is used, A is overwritten by $\text{diag}(r) * A * \text{diag}(c)$ and B by $\text{diag}(r) * B$ (if `trans = 'N'`) or $\text{diag}(c) * B$ (if `trans = 'T'` or `'C'`).

2. If `fact = 'N'` or `'E'`, the LU decomposition is used to factor the matrix A (after equilibration if `fact = 'E'`) as $A = P L U$, where P is a permutation matrix, L is a unit lower triangular matrix, and U is upper triangular.

3. If some $U_{i,i} = 0$, so that U is exactly singular, then the routine returns with `info = i`. Otherwise, the factored form of A is used to estimate the condition number of the matrix A . If the reciprocal of the condition number is less than machine precision, `info = n + 1` is returned as a warning, but the routine still goes on to solve for X and compute error bounds as described below.

4. The system of equations is solved for X using the factored form of A .
5. Iterative refinement is applied to improve the computed solution matrix and calculate error bounds and backward error estimates for it.
6. If equilibration was used, the matrix X is premultiplied by $\text{diag}(c)$ (if `trans = 'N'`) or $\text{diag}(r)$ (if `trans = 'T'` or `'C'`) so that it solves the original system before equilibration.

Input Parameters

- fact* CHARACTER*1. Must be 'F', 'N', or 'E'.
 Specifies whether or not the factored form of the matrix A is supplied on entry, and if not, whether the matrix A should be equilibrated before it is factored.
 If *fact* = 'F': on entry, *af* and *ipiv* contain the factored form of A . If *equed* is not 'N', the matrix A has been equilibrated with scaling factors given by *r* and *c*. *a*, *af*, and *ipiv* are not modified.
 If *fact* = 'N', the matrix A will be copied to *af* and factored.
 If *fact* = 'E', the matrix A will be equilibrated if necessary, then copied to *af* and factored.
- trans* CHARACTER*1. Must be 'N', 'T', or 'C'.
 Specifies the form of the system of equations:
 If *trans* = 'N', the system has the form $A X = B$ (No transpose);
 If *trans* = 'T', the system has the form $A^T X = B$ (Transpose);
 If *trans* = 'C', the system has the form $A^H X = B$ (Conjugate transpose);
- n* INTEGER. The number of linear equations; the order of the matrix A ($n \geq 0$).
- nrhs* INTEGER. The number of right hand sides; the number of columns of the matrices B and X ($nrhs \geq 0$).
- a, af, b, work* REAL for sgesvx
 DOUBLE PRECISION for dgesvx
 COMPLEX for cgesvx
 DOUBLE COMPLEX for zgesvx.
 Arrays: *a*(*lda*, *), *af*(*ldaf*, *), *b*(*ldb*, *), *work*(*).
 The array *a* contains the matrix A . If *fact* = 'F' and *equed* is not 'N', then A must have been equilibrated by the scaling factors in *r* and/or *c*. The second dimension

of a must be at least $\max(1,n)$.

The array af is an input argument if $fact = 'F'$. It contains the factored form of the matrix A , i.e., the factors L and U from the factorization $A = P L U$ as computed by `?getrf`. If $equed$ is not 'N', then af is the factored form of the equilibrated matrix A . The second dimension of af must be at least $\max(1,n)$.

The array b contains the matrix B whose columns are the right-hand sides for the systems of equations. The second dimension of b must be at least $\max(1,nrhs)$.

$work(*)$ is a workspace array.

The dimension of $work$ must be at least $\max(1,4*n)$ for real flavors, and at least $\max(1,2*n)$ for complex flavors.

lda **INTEGER.** The first dimension of a ; $lda \geq \max(1, n)$.

$ldaf$ **INTEGER.** The first dimension of af ; $ldaf \geq \max(1, n)$.

ldb **INTEGER.** The first dimension of b ; $ldb \geq \max(1, n)$.

$ipiv$ **INTEGER.**

Array, **DIMENSION** at least $\max(1,n)$.

The array $ipiv$ is an input argument if $fact = 'F'$.

It contains the pivot indices from the factorization $A = P L U$ as computed by `?getrf`; row i of the matrix was interchanged with row $ipiv(i)$.

$equed$ **CHARACTER*1.** Must be 'N', 'R', 'C', or 'B'.
 $equed$ is an input argument if $fact = 'F'$. It specifies the form of equilibration that was done:
 If $equed = 'N'$, no equilibration was done (always true if $fact = 'N'$);
 If $equed = 'R'$, row equilibration was done and A has been premultiplied by $diag(r)$;
 If $equed = 'C'$, column equilibration was done and A has been postmultiplied by $diag(c)$;
 If $equed = 'B'$, both row and column equilibration was done; A has been replaced by $diag(r)*A*diag(c)$.

r, c **REAL** for single precision flavors;
DOUBLE PRECISION for double precision flavors.
 Arrays: $r(n)$, $c(n)$.
 The array r contains the row scale factors for A , and the array c contains the column scale factors for A . These arrays are input arguments if $fact = 'F'$ only; otherwise they are output arguments.
 If $equed = 'R'$ or $'B'$, A is multiplied on the left by $diag(r)$; if $equed = 'N'$ or $'C'$, r is not accessed.
 If $fact = 'F'$ and $equed = 'R'$ or $'B'$, each element of r must be positive.
 If $equed = 'C'$ or $'B'$, A is multiplied on the right by $diag(c)$; if $equed = 'N'$ or $'R'$, c is not accessed.
 If $fact = 'F'$ and $equed = 'C'$ or $'B'$, each element of c must be positive.

ldx **INTEGER**. The first dimension of the output array x ;
 $ldx \geq \max(1, n)$.

iwork **INTEGER**.
 Workspace array, **DIMENSION** at least $\max(1, n)$; used in real flavors only.

rwork **REAL** for single precision flavors;
DOUBLE PRECISION for double precision flavors.
 Workspace array, **DIMENSION** at least $\max(1, 2*n)$; used in complex flavors only.

Output Parameters

x **REAL** for `sgesvx`
DOUBLE PRECISION for `dgesvx`
COMPLEX for `cgesvx`
DOUBLE COMPLEX for `zgesvx`.
 Array, **DIMENSION** ($ldx, *$).
 If $info = 0$ or $info = n+1$, the array x contains the solution matrix X to the *original* system of equations. Note that A and B are modified on exit if $equed \neq 'N'$, and the solution to the *equilibrated* system is: $diag(c)^{-1} * X$, if $trans = 'N'$ and $equed = 'C'$ or $'B'$;

- $\text{diag}(x)^{-1} * X$, if *trans* = 'T' or 'C' and *equed* = 'R' or 'B'.
- The second dimension of *x* must be at least $\max(1, nrhs)$.
- a* Array *a* is not modified on exit if *fact* = 'F' or 'N', or if *fact* = 'E' and *equed* = 'N'.
If *equed* ≠ 'N', *A* is scaled on exit as follows:
equed = 'R': $A = \text{diag}(x) * A$
equed = 'C': $A = A * \text{diag}(c)$
equed = 'B': $A = \text{diag}(x) * A * \text{diag}(c)$
- af* If *fact* = 'N' or 'E', then *af* is an output argument and on exit returns the factors *L* and *U* from the factorization $A = P L U$ of the original matrix *A* (if *fact* = 'N') or of the equilibrated matrix *A* (if *fact* = 'E'). See the description of *a* for the form of the equilibrated matrix.
- b* Overwritten by $\text{diag}(x) * B$ if *trans* = 'N' and *equed* = 'R' or 'B';
overwritten by $\text{diag}(c) * B$ if *trans* = 'T' and *equed* = 'C' or 'B';
not changed if *equed* = 'N'.
- r, c* These arrays are output arguments if *fact* ≠ 'F'.
See the description of *r, c* in *Input Arguments* section.
- rcond* REAL for single precision flavors.
DOUBLE PRECISION for double precision flavors.
An estimate of the reciprocal condition number of the matrix *A* after equilibration (if done). The routine sets *rcond* = 0 if the estimate underflows; in this case the matrix is singular (to working precision). However, anytime *rcond* is small compared to 1.0, for the working precision, the matrix may be poorly conditioned or even singular.
- ferr, berr* REAL for single precision flavors.
DOUBLE PRECISION for double precision flavors.
Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and relative backward errors, respectively, for each solution vector.

- ipiv* If *fact* = 'N' or 'E', then *ipiv* is an output argument and on exit contains the pivot indices from the factorization $A = P L U$ of the original matrix A (if *fact* = 'N') or of the equilibrated matrix A (if *fact* = 'E').
- equed* If *fact* \neq 'F', then *equed* is an output argument. It specifies the form of equilibration that was done (see the description of *equed* in *Input Arguments* section).
- work*, *rwork* On exit, *work*(1) for real flavors, or *rwork*(1) for complex flavors, contains the reciprocal pivot growth factor $\text{norm}(A)/\text{norm}(U)$. The "max absolute element" norm is used. If *work*(1) for real flavors, or *rwork*(1) for complex flavors is much less than 1, then the stability of the LU factorization of the (equilibrated) matrix A could be poor. This also means that the solution x , condition estimator *rcond*, and forward error bound *ferr* could be unreliable. If factorization fails with $0 < \text{info} \leq n$, then *work*(1) for real flavors, or *rwork*(1) for complex flavors contains the reciprocal pivot growth factor for the leading *info* columns of A .
- info* INTEGER. If *info*=0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i*, and $i \leq n$, then $U(i,i)$ is exactly zero. The factorization has been completed, but the factor U is exactly singular, so the solution and error bounds could not be computed; *rcond* = 0 is returned.
If *info* = *i*, and $i = n + 1$, then U is nonsingular, but *rcond* is less than machine precision, meaning that the matrix is singular to working precision. Nevertheless, the solution and error bounds are computed because there are a number of situations where the computed solution can be more accurate than the value of *rcond* would suggest.

?gbsv

Computes the solution to the system of linear equations with a band matrix A and multiple right-hand sides.

```
call sgbsv (n, kl, ku, nrhs, ab, ldab, ipiv, b, ldb, info)
call dgbsv (n, kl, ku, nrhs, ab, ldab, ipiv, b, ldb, info)
call cgbsv (n, kl, ku, nrhs, ab, ldab, ipiv, b, ldb, info)
call zgbsv (n, kl, ku, nrhs, ab, ldab, ipiv, b, ldb, info)
```

Discussion

This routine solves for X the real or complex system of linear equations $AX = B$, where A is an n -by- n band matrix with kl subdiagonals and ku superdiagonals, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

The LU decomposition with partial pivoting and row interchanges is used to factor A as $A = LU$, where L is a product of permutation and unit lower triangular matrices with kl subdiagonals, and U is upper triangular with $kl+ku$ superdiagonals. The factored form of A is then used to solve the system of equations $AX = B$.

Input Parameters

<i>n</i>	INTEGER. The order of A ; the number of rows in B ($n \geq 0$).
<i>kl</i>	INTEGER. The number of sub-diagonals within the band of A ($kl \geq 0$).
<i>ku</i>	INTEGER. The number of super-diagonals within the band of A ($ku \geq 0$).
<i>nrhs</i>	INTEGER. The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).
<i>ab, b</i>	REAL for sgbsv DOUBLE PRECISION for dgbsv COMPLEX for cgbsv

DOUBLE COMPLEX for `zgbstv`.

Arrays: `ab(ldab,*)`, `b(ldb,*)`.

The array `ab` contains the matrix A in band storage (see [Matrix Storage Schemes](#)).

The second dimension of `ab` must be at least $\max(1, n)$.

The array `b` contains the matrix B whose columns are the right-hand sides for the systems of equations.

The second dimension of `b` must be at least $\max(1, nrhs)$.

`ldab` INTEGER. The first dimension of the array `ab`.
($ldab \geq 2kl + ku + 1$)

`ldb` INTEGER. The first dimension of `b`; $ldb \geq \max(1, n)$.

Output Parameters

`ab` Overwritten by L and U . The diagonal and $kl + ku$ super-diagonals of U are stored in the first $1 + kl + ku$ rows of `ab`. The multipliers used to form L are stored in the next kl rows.

`b` Overwritten by the solution matrix X .

`ipiv` INTEGER.
Array, DIMENSION at least $\max(1, n)$.
The pivot indices: row i was interchanged with row `ipiv(i)`.

`info` INTEGER. If `info`=0, the execution is successful.
If `info` = $-i$, the i th parameter had an illegal value.
If `info` = i , $U(i, i)$ is exactly zero. The factorization has been completed, but the factor U is exactly singular, so the solution could not be computed.

?gbsvx

Computes the solution to the real or complex system of linear equations with a band matrix A and multiple right-hand sides, and provides error bounds on the solution.

```
call sgbvx (fact, trans, n, kl, ku, nrhs, ab, ldab, afb, ldafb,
           ipiv, equed, r, c, b, ldb, x, ldx, rcond, ferr, berr,
           work, iwork, info)
call dgbvx (fact, trans, n, kl, ku, nrhs, ab, ldab, afb, ldafb,
           ipiv, equed, r, c, b, ldb, x, ldx, rcond, ferr, berr,
           work, iwork, info)
call cgbvx (fact, trans, n, kl, ku, nrhs, ab, ldab, afb, ldafb,
           ipiv, equed, r, c, b, ldb, x, ldx, rcond, ferr, berr,
           work, rwork, info)
call zgbvx (fact, trans, n, kl, ku, nrhs, ab, ldab, afb, ldafb,
           ipiv, equed, r, c, b, ldb, x, ldx, rcond, ferr, berr,
           work, rwork, info)
```

Discussion

This routine uses the LU factorization to compute the solution to a real or complex system of linear equations $AX = B$, $A^T X = B$, or $A^H X = B$, where A is a band matrix of order n with kl subdiagonals and ku superdiagonals, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

Error bounds on the solution and a condition estimate are also provided.

The routine ?gbsvx performs the following steps:

1. If $fact = 'E'$, real scaling factors r and c are computed to equilibrate the system:

$$trans = 'N': \quad \text{diag}(r) * A * \text{diag}(c) * \text{diag}(c)^{-1} * X = \text{diag}(r) * B$$

$$trans = 'T': \quad (\text{diag}(r) * A * \text{diag}(c))^T * \text{diag}(r)^{-1} * X = \text{diag}(c) * B$$

$$trans = 'C': \quad (\text{diag}(r) * A * \text{diag}(c))^H * \text{diag}(r)^{-1} * X = \text{diag}(c) * B$$

Whether or not the system will be equilibrated depends on the scaling of the matrix A , but if equilibration is used, A is overwritten by $\text{diag}(r) * A * \text{diag}(c)$ and B by $\text{diag}(r) * B$ (if $trans = 'N'$) or $\text{diag}(c) * B$ (if $trans = 'T'$ or $'C'$).

2. If $fact = 'N'$ or $'E'$, the LU decomposition is used to factor the matrix A (after equilibration if $fact = 'E'$) as $A = L U$, where L is a product of permutation and unit lower triangular matrices with kl subdiagonals, and U is upper triangular with $kl+ku$ superdiagonals.

3. If some $U_{i,i} = 0$, so that U is exactly singular, then the routine returns with $info = i$. Otherwise, the factored form of A is used to estimate the condition number of the matrix A . If the reciprocal of the condition number is less than machine precision, $info = n + 1$ is returned as a warning, but the routine still goes on to solve for X and compute error bounds as described below.

4. The system of equations is solved for X using the factored form of A .

5. Iterative refinement is applied to improve the computed solution matrix and calculate error bounds and backward error estimates for it.

6. If equilibration was used, the matrix X is premultiplied by $\text{diag}(c)$ (if $trans = 'N'$) or $\text{diag}(r)$ (if $trans = 'T'$ or $'C'$) so that it solves the original system before equilibration.

Input Parameters

fact CHARACTER*1. Must be 'F', 'N', or 'E'.

Specifies whether or not the factored form of the matrix A is supplied on entry, and if not, whether the matrix A should be equilibrated before it is factored.

If $fact = 'F'$: on entry, *afb* and *ipiv* contain the factored form of A . If *equed* is not 'N', the matrix A has been equilibrated with scaling factors given by r and c . *ab*, *afb*, and *ipiv* are not modified.

If $fact = 'N'$, the matrix A will be copied to *afb* and factored.

If $fact = 'E'$, the matrix A will be equilibrated if necessary, then copied to *afb* and factored.

trans CHARACTER*1. Must be 'N', 'T', or 'C'.

Specifies the form of the system of equations:

If *trans* = 'N', the system has the form $A X = B$
(No transpose);

If *trans* = 'T', the system has the form $A^T X = B$
(Transpose);

If *trans* = 'C', the system has the form $A^H X = B$
(Conjugate transpose);

n **INTEGER**. The number of linear equations; the order of the matrix *A* ($n \geq 0$).

kl **INTEGER**. The number of sub-diagonals within the band of *A* ($kl \geq 0$).

ku **INTEGER**. The number of super-diagonals within the band of *A* ($ku \geq 0$).

nrhs **INTEGER**. The number of right hand sides; the number of columns of the matrices *B* and *X* ($nrhs \geq 0$).

ab,afb,b,work **REAL** for *sgevsx*

DOUBLE PRECISION for *dgevsx*

COMPLEX for *cgevsx*

DOUBLE COMPLEX for *zgevsx*.

Arrays: *a(lda,*)*, *af(ldaf,*)*, *b(ldb,*)*,
work()*.

The array *ab* contains the matrix *A* in band storage (see [Matrix Storage Schemes](#)).

The second dimension of *ab* must be at least $\max(1, n)$.

If *fact* = 'F' and *equed* is not 'N', then *A* must have been equilibrated by the scaling factors in *r* and/or *c*.

The array *afb* is an input argument if *fact* = 'F'.

The second dimension of *afb* must be at least $\max(1, n)$.

It contains the factored form of the matrix *A*, i.e., the factors *L* and *U* from the factorization $A = L U$ as computed by [?gbtrf](#). *U* is stored as an upper triangular band matrix with $kl + ku$ super-diagonals in the first $1 + kl + ku$ rows of *afb*. The multipliers used during

the factorization are stored in the next kl rows.

If $equed$ is not 'N', then afb is the factored form of the equilibrated matrix A .

The array b contains the matrix B whose columns are the right-hand sides for the systems of equations. The second dimension of b must be at least $\max(1, nrhs)$.

$work(*)$ is a workspace array.

The dimension of $work$ must be at least $\max(1, 3*n)$ for real flavors, and at least $\max(1, 2*n)$ for complex flavors.

$ldab$ INTEGER. The first dimension of ab ; $ldab \geq kl + ku + 1$.

$ldafb$ INTEGER. The first dimension of afb ;
 $ldafb \geq 2*kl + ku + 1$.

ldb INTEGER. The first dimension of b ; $ldb \geq \max(1, n)$.

$ipiv$ INTEGER.
Array, DIMENSION at least $\max(1, n)$.
The array $ipiv$ is an input argument if $fact = 'F'$.
It contains the pivot indices from the factorization $A = LU$ as computed by `?gbtrf`; row i of the matrix was interchanged with row $ipiv(i)$.

$equed$ CHARACTER*1. Must be 'N', 'R', 'C', or 'B'.
 $equed$ is an input argument if $fact = 'F'$. It specifies the form of equilibration that was done:
If $equed = 'N'$, no equilibration was done (always true if $fact = 'N'$);
If $equed = 'R'$, row equilibration was done and A has been premultiplied by $\text{diag}(r)$;
If $equed = 'C'$, column equilibration was done and A has been postmultiplied by $\text{diag}(c)$;
If $equed = 'B'$, both row and column equilibration was done; A has been replaced by $\text{diag}(r)*A*\text{diag}(c)$.

r, c REAL for single precision flavors;
DOUBLE PRECISION for double precision flavors.
Arrays: $r(n)$, $c(n)$.
The array r contains the row scale factors for A , and the

array c contains the column scale factors for A . These arrays are input arguments if $fact = 'F'$ only; otherwise they are output arguments.

If $equed = 'R'$ or $'B'$, A is multiplied on the left by $\text{diag}(r)$; if $equed = 'N'$ or $'C'$, r is not accessed.

If $fact = 'F'$ and $equed = 'R'$ or $'B'$, each element of r must be positive.

If $equed = 'C'$ or $'B'$, A is multiplied on the right by $\text{diag}(c)$; if $equed = 'N'$ or $'R'$, c is not accessed.

If $fact = 'F'$ and $equed = 'C'$ or $'B'$, each element of c must be positive.

ldx	INTEGER. The first dimension of the output array x ; $ldx \geq \max(1, n)$.
$iwork$	INTEGER. Workspace array, DIMENSION at least $\max(1, n)$; used in real flavors only.
$rwork$	REAL for single precision flavors; DOUBLE PRECISION for double precision flavors. Workspace array, DIMENSION at least $\max(1, n)$; used in complex flavors only.

Output Parameters

x	REAL for <code>sgbsvx</code> DOUBLE PRECISION for <code>dgbsvx</code> COMPLEX for <code>cgbsvx</code> DOUBLE COMPLEX for <code>zgbsvx</code> . Array, DIMENSION ($ldx, *$). If $info = 0$ or $info = n+1$, the array x contains the solution matrix X to the <i>original</i> system of equations. Note that A and B are modified on exit if $equed \neq 'N'$, and the solution to the <i>equilibrated</i> system is: $\text{diag}(c)^{-1} * X$, if $trans = 'N'$ and $equed = 'C'$ or $'B'$; $\text{diag}(r)^{-1} * X$, if $trans = 'T'$ or $'C'$ and $equed = 'R'$ or $'B'$. The second dimension of x must be at least $\max(1, nrhs)$.
-----	--

<i>ab</i>	<p>Array <i>ab</i> is not modified on exit if <i>fact</i> = 'F' or 'N', or if <i>fact</i> = 'E' and <i>equed</i> = 'N'.</p> <p>If <i>equed</i> ≠ 'N', <i>A</i> is scaled on exit as follows:</p> <p><i>equed</i> = 'R': $A = \text{diag}(r) * A$</p> <p><i>equed</i> = 'C': $A = A * \text{diag}(c)$</p> <p><i>equed</i> = 'B': $A = \text{diag}(r) * A * \text{diag}(c)$</p>
<i>afb</i>	<p>If <i>fact</i> = 'N' or 'E', then <i>afb</i> is an output argument and on exit returns details of the <i>LU</i> factorization of the original matrix <i>A</i> (if <i>fact</i> = 'N') or of the equilibrated matrix <i>A</i> (if <i>fact</i> = 'E'). See the description of <i>ab</i> for the form of the equilibrated matrix.</p>
<i>b</i>	<p>Overwritten by $\text{diag}(r) * b$ if <i>trans</i> = 'N' and <i>equed</i> = 'R' or 'B';</p> <p>overwritten by $\text{diag}(c) * b$ if <i>trans</i> = 'T' and <i>equed</i> = 'C' or 'B';</p> <p>not changed if <i>equed</i> = 'N'.</p>
<i>r, c</i>	<p>These arrays are output arguments if <i>fact</i> ≠ 'F'.</p> <p>See the description of <i>r, c</i> in <i>Input Arguments</i> section.</p>
<i>rcond</i>	<p>REAL for single precision flavors.</p> <p>DOUBLE PRECISION for double precision flavors.</p> <p>An estimate of the reciprocal condition number of the matrix <i>A</i> after equilibration (if done).</p> <p>If <i>rcond</i> is less than the machine precision (in particular, if <i>rcond</i> = 0), the matrix is singular to working precision. This condition is indicated by a return code of <i>info</i> > 0.</p>
<i>ferr, berr</i>	<p>REAL for single precision flavors.</p> <p>DOUBLE PRECISION for double precision flavors.</p> <p>Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and relative backward errors, respectively, for each solution vector.</p>
<i>ipiv</i>	<p>If <i>fact</i> = 'N' or 'E', then <i>ipiv</i> is an output argument and on exit contains the pivot indices from the factorization $A = LU$ of the original matrix <i>A</i> (if <i>fact</i> = 'N') or of the equilibrated matrix <i>A</i> (if <i>fact</i> = 'E').</p>

<i>equed</i>	If <i>fact</i> ≠ 'F' , then <i>equed</i> is an output argument. It specifies the form of equilibration that was done (see the description of <i>equed</i> in <i>Input Arguments</i> section).
<i>work, rwork</i>	On exit, <i>work</i> (1) for real flavors, or <i>rwork</i> (1) for complex flavors, contains the reciprocal pivot growth factor $\text{norm}(A)/\text{norm}(U)$. The "max absolute element" norm is used. If <i>work</i> (1) for real flavors, or <i>rwork</i> (1) for complex flavors is much less than 1, then the stability of the <i>LU</i> factorization of the (equilibrated) matrix <i>A</i> could be poor. This also means that the solution <i>x</i> , condition estimator <i>rcond</i> , and forward error bound <i>ferr</i> could be unreliable. If factorization fails with $0 < \text{info} \leq n$, then <i>work</i> (1) for real flavors, or <i>rwork</i> (1) for complex flavors contains the reciprocal pivot growth factor for the leading <i>info</i> columns of <i>A</i> .
<i>info</i>	INTEGER. If <i>info</i> =0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = <i>i</i> , and $i \leq n$, then $U(i,i)$ is exactly zero. The factorization has been completed, but the factor <i>U</i> is exactly singular, so the solution and error bounds could not be computed; <i>rcond</i> = 0 is returned. If <i>info</i> = <i>i</i> , and $i = n + 1$, then <i>U</i> is nonsingular, but <i>rcond</i> is less than machine precision, meaning that the matrix is singular to working precision. Nevertheless, the solution and error bounds are computed because there are a number of situations where the computed solution can be more accurate than the value of <i>rcond</i> would suggest.

?gtsv

Computes the solution to the system of linear equations with a tridiagonal matrix A and multiple right-hand sides.

```
call sgtsv (n, nrhs, dl, d, du, b, ldb, info)
call dgtsv (n, nrhs, dl, d, du, b, ldb, info)
call cgtsv (n, nrhs, dl, d, du, b, ldb, info)
call zgtsv (n, nrhs, dl, d, du, b, ldb, info)
```

Discussion

This routine solves for X the system of linear equations $AX = B$, where A is an n -by- n tridiagonal matrix, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions. The routine uses Gaussian elimination with partial pivoting.

Note that the equation $A^T X = B$ may be solved by interchanging the order of the arguments du and dl .

Input Parameters

n **INTEGER**. The order of A ; the number of rows in B ($n \geq 0$).

$nrhs$ **INTEGER**. The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).

dl, d, du, b **REAL** for **sgtsv**
DOUBLE PRECISION for **dgtsv**
COMPLEX for **cgtsv**
DOUBLE COMPLEX for **zgtsv**.
 Arrays: $dl(n-1)$, $d(n)$, $du(n-1)$, $b(ldb,*)$.
 The array dl contains the $(n-1)$ subdiagonal elements of A .
 The array d contains the diagonal elements of A .
 The array du contains the $(n-1)$ superdiagonal elements of A .

The array *b* contains the matrix *B* whose columns are the right-hand sides for the systems of equations. The second dimension of *b* must be at least $\max(1, nrhs)$.

ldb **INTEGER.** The first dimension of *b*; $ldb \geq \max(1, n)$.

Output Parameters

dl Overwritten by the $(n-2)$ elements of the second superdiagonal of the upper triangular matrix *U* from the *LU* factorization of *A*. These elements are stored in *dl*(1), ... , *dl*(*n*-2).

d Overwritten by the *n* diagonal elements of *U*.

du Overwritten by the $(n-1)$ elements of the first superdiagonal of *U*.

b Overwritten by the solution matrix *X*.

info **INTEGER.** If *info*=0, the execution is successful. If *info* = -*i*, the *i*th parameter had an illegal value. If *info* = *i*, *U*(*i*,*i*) is exactly zero, and the solution has not been computed. The factorization has not been completed unless *i* = *n*.

?gtsvx

Computes the solution to the real or complex system of linear equations with a tridiagonal matrix A and multiple right-hand sides, and provides error bounds on the solution.

```
call sgtsvx (fact, trans, n, nrhs, dl, d, du, dlf, df, duf, du2,
            ipiv, b, ldb, x, ldx, rcond, ferr, berr, work,
            iwork, info)
call dgtsvx (fact, trans, n, nrhs, dl, d, du, dlf, df, duf, du2,
            ipiv, b, ldb, x, ldx, rcond, ferr, berr, work,
            iwork, info)
call cgtsvx (fact, trans, n, nrhs, dl, d, du, dlf, df, duf, du2,
            ipiv, b, ldb, x, ldx, rcond, ferr, berr, work,
            rwork, info)
call zgtsvx (fact, trans, n, nrhs, dl, d, du, dlf, df, duf, du2,
            ipiv, b, ldb, x, ldx, rcond, ferr, berr, work,
            rwork, info)
```

Discussion

This routine uses the LU factorization to compute the solution to a real or complex system of linear equations $AX = B$, $A^T X = B$, or $A^H X = B$, where A is a tridiagonal matrix of order n , the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

Error bounds on the solution and a condition estimate are also provided.

The routine ?gtsvx performs the following steps:

1. If `fact = 'N'`, the LU decomposition is used to factor the matrix A as $A = LU$, where L is a product of permutation and unit lower bidiagonal matrices and U is an upper triangular matrix with nonzeros in only the main diagonal and first two superdiagonals.
2. If some $U_{i,i} = 0$, so that U is exactly singular, then the routine returns with `info = i`. Otherwise, the factored form of A is used to estimate the condition number of the matrix A . If the reciprocal of the condition number

is less than machine precision, $info = n + 1$ is returned as a warning, but the routine still goes on to solve for X and compute error bounds as described below.

3. The system of equations is solved for X using the factored form of A .
4. Iterative refinement is applied to improve the computed solution matrix and calculate error bounds and backward error estimates for it.

Input Parameters

fact CHARACTER*1. Must be 'F' or 'N'.
 Specifies whether or not the factored form of the matrix A has been supplied on entry.
 If *fact* = 'F': on entry, *dlf*, *df*, *duf*, *du2*, and *ipiv* contain the factored form of A ; arrays *dl*, *d*, *du*, *dlf*, *df*, *duf*, *du2*, and *ipiv* will not be modified.
 If *fact* = 'N', the matrix A will be copied to *dlf*, *df*, and *duf* and factored.

trans CHARACTER*1. Must be 'N', 'T', or 'C'.
 Specifies the form of the system of equations:
 If *trans* = 'N', the system has the form $A X = B$ (No transpose);
 If *trans* = 'T', the system has the form $A^T X = B$ (Transpose);
 If *trans* = 'C', the system has the form $A^H X = B$ (Conjugate transpose);

n INTEGER. The number of linear equations; the order of the matrix A ($n \geq 0$).

nrhs INTEGER. The number of right hand sides; the number of columns of the matrices B and X ($nrhs \geq 0$).

dl, d, du, dlf, df, duf, du2, b, x, work REAL for *sgtsvx*
 DOUBLE PRECISION for *dgtsvx*
 COMPLEX for *cgtsvx*
 DOUBLE COMPLEX for *zgtsvx*.

Arrays:

d1, dimension $(n - 1)$, contains the subdiagonal elements of A .

d, dimension (n) , contains the diagonal elements of A .

du, dimension $(n - 1)$, contains the superdiagonal elements of A .

d1f, dimension $(n - 1)$. If *fact* = 'F', then *d1f* is an input argument and on entry contains the $(n - 1)$ multipliers that define the matrix L from the LU factorization of A as computed by [?gttrf](#).

df, dimension (n) . If *fact* = 'F', then *df* is an input argument and on entry contains the n diagonal elements of the upper triangular matrix U from the LU factorization of A .

duf, dimension $(n - 1)$. If *fact* = 'F', then *duf* is an input argument and on entry contains the $(n - 1)$ elements of the first super-diagonal of U .

du2, dimension $(n - 2)$. If *fact* = 'F', then *du2* is an input argument and on entry contains the $(n - 2)$ elements of the second super-diagonal of U .

b(ldb,)* contains the right-hand side matrix B . The second dimension of *b* must be at least $\max(1, nrhs)$.

x(ldx,)* contains the solution matrix X . The second dimension of *x* must be at least $\max(1, nrhs)$.

work()* is a workspace array;

the dimension of *work* must be at least $\max(1, 3*n)$ for real flavors and $\max(1, 2*n)$ for complex flavors.

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

ldx INTEGER. The first dimension of *x*; $ldx \geq \max(1, n)$.

ipiv INTEGER.

Array, DIMENSION at least $\max(1, n)$. If *fact* = 'F', then *ipiv* is an input argument and on entry contains the pivot indices, as returned by [?gttrf](#).

iwork INTEGER.

Workspace array, DIMENSION (n) . Used for real flavors only.

rwork REAL for `cgtsvx`
 DOUBLE PRECISION for `zgtsvx`.
 Workspace array, DIMENSION (*n*). Used for complex flavors only.

Output Parameters

x REAL for `sgtsvx`
 DOUBLE PRECISION for `dgtsvx`
 COMPLEX for `cgtsvx`
 DOUBLE COMPLEX for `zgtsvx`.
 Array, DIMENSION (*ldx*, *).
 If *info* = 0 or *info* = *n*+1, the array *x* contains the solution matrix *X*. The second dimension of *x* must be at least max(1,*nrhs*).

dlf If *fact* = 'N', then *dlf* is an output argument and on exit contains the (*n* - 1) multipliers that define the matrix *L* from the *LU* factorization of *A*.

df If *fact* = 'N', then *df* is an output argument and on exit contains the *n* diagonal elements of the upper triangular matrix *U* from the *LU* factorization of *A*.

duf If *fact* = 'N', then *duf* is an output argument and on exit contains the (*n* - 1) elements of the first super-diagonal of *U*.

du2 If *fact* = 'N', then *du2* is an output argument and on exit contains the (*n* - 2) elements of the second super-diagonal of *U*.

ipiv The array *ipiv* is an output argument if *fact* = 'N' and, on exit, contains the pivot indices from the factorization $A = LU$; row *i* of the matrix was interchanged with row *ipiv*(*i*). The value of *ipiv*(*i*) will always be either *i* or *i*+1; *ipiv*(*i*)=*i* indicates a row interchange was not required.

rcond REAL for single precision flavors.
 DOUBLE PRECISION for double precision flavors.
 An estimate of the reciprocal condition number of the

matrix A .

If $rcond$ is less than the machine precision (in particular, if $rcond = 0$), the matrix is singular to working precision. This condition is indicated by a return code of $info > 0$.

ferr, berr

REAL for single precision flavors.

DOUBLE PRECISION for double precision flavors.

Arrays, **DIMENSION** at least $\max(1, nrhs)$. Contain the component-wise forward and backward errors, respectively, for each solution vector.

info

INTEGER. If $info=0$, the execution is successful.

If $info = -i$, the i th parameter had an illegal value.

If $info = i$, and $i \leq n$, then $U(i,i)$ is exactly zero. The factorization has not been completed unless $i = n$, but the factor U is exactly singular, so the solution and error bounds could not be computed; $rcond = 0$ is returned.

If $info = i$, and $i = n + 1$, then U is nonsingular, but $rcond$ is less than machine precision, meaning that the matrix is singular to working precision. Nevertheless, the solution and error bounds are computed because there are a number of situations where the computed solution can be more accurate than the value of $rcond$ would suggest.

?posv

Computes the solution to the system of linear equations with a symmetric or Hermitian positive definite matrix A and multiple right-hand sides.

```
call sposv (uplo, n, nrhs, a, lda, b, ldb, info)
call dposv (uplo, n, nrhs, a, lda, b, ldb, info)
call cposv (uplo, n, nrhs, a, lda, b, ldb, info)
call zposv (uplo, n, nrhs, a, lda, b, ldb, info)
```

Discussion

This routine solves for X the real or complex system of linear equations $AX = B$, where A is an n -by- n symmetric/Hermitian positive definite matrix, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

The Cholesky decomposition is used to factor A as $A = U^H U$ if `uplo = 'U'` or $A = LL^H$ if `uplo = 'L'`, where U is an upper triangular matrix and L is a lower triangular matrix. The factored form of A is then used to solve the system of equations $AX = B$.

Input Parameters

<code>uplo</code>	<code>CHARACTER*1</code> . Must be 'U' or 'L'. Indicates whether the upper or lower triangular part of A is stored and how A is factored: If <code>uplo = 'U'</code> , the array <code>a</code> stores the upper triangular part of the matrix A , and A is factored as $U^H U$. If <code>uplo = 'L'</code> , the array <code>a</code> stores the lower triangular part of the matrix A ; A is factored as LL^H .
<code>n</code>	<code>INTEGER</code> . The order of matrix A ($n \geq 0$).
<code>nrhs</code>	<code>INTEGER</code> . The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).

a, *b* REAL for *sposv*
DOUBLE PRECISION for *dposv*
COMPLEX for *cposv*
DOUBLE COMPLEX for *zposv*.
Arrays: *a(lda,*)*, *b ldb,*)*.
The array *a* contains either the upper or the lower triangular part of the matrix *A* (see *uplo*).
The second dimension of *a* must be at least $\max(1, n)$.
The array *b* contains the matrix *B* whose columns are the right-hand sides for the systems of equations.
The second dimension of *b* must be at least $\max(1, nrhs)$.

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, n)$.
ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

Output Parameters

a If *info*=0, the upper or lower triangular part of *a* is overwritten by the Cholesky factor *U* or *L*, as specified by *uplo*.

b Overwritten by the solution matrix *X*.

info INTEGER. If *info*=0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i*, the leading minor of order *i* (and hence the matrix *A* itself) is not positive definite, so the factorization could not be completed, and the solution has not been computed.

?posvx

Uses the Cholesky factorization to compute the solution to the system of linear equations with a symmetric or Hermitian positive definite matrix A , and provides error bounds on the solution.

```
call sposvx (fact, uplo, n, nrhs, a, lda, af, ldaf, equed, s, b,
            ldb, x, ldx, rcond, ferr, berr, work, iwork, info)
call dposvx (fact, uplo, n, nrhs, a, lda, af, ldaf, equed, s, b,
            ldb, x, ldx, rcond, ferr, berr, work, iwork, info)
call cposvx (fact, uplo, n, nrhs, a, lda, af, ldaf, equed, s, b,
            ldb, x, ldx, rcond, ferr, berr, work, rwork, info)
call zposvx (fact, uplo, n, nrhs, a, lda, af, ldaf, equed, s, b,
            ldb, x, ldx, rcond, ferr, berr, work, rwork, info)
```

Discussion

This routine uses the Cholesky factorization $A=U^H U$ or $A=LL^H$ to compute the solution to a real or complex system of linear equations $AX=B$, where A is a n -by- n real symmetric/Hermitian positive definite matrix, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

Error bounds on the solution and a condition estimate are also provided.

The routine ?posvx performs the following steps:

1. If $fact = 'E'$, real scaling factors s are computed to equilibrate the system:

$$\text{diag}(s) * A * \text{diag}(s) * \text{diag}(s)^{-1} * X = \text{diag}(s) * B$$

Whether or not the system will be equilibrated depends on the scaling of the matrix A , but if equilibration is used, A is overwritten by $\text{diag}(s) * A * \text{diag}(s)$ and B by $\text{diag}(s) * B$.

2. If $fact = 'N'$ or $'E'$, the Cholesky decomposition is used to factor the matrix A (after equilibration if $fact = 'E'$) as

$A = U^H U$, if `uplo = 'U'`, or

$A = L L^H$, if `uplo = 'L'`,

where U is an upper triangular matrix and L is a lower triangular matrix.

3. If the leading i -by- i principal minor is not positive definite, then the routine returns with `info = i`. Otherwise, the factored form of A is used to estimate the condition number of the matrix A . If the reciprocal of the condition number is less than machine precision, `info = n + 1` is returned as a warning, but the routine still goes on to solve for X and compute error bounds as described below.

4. The system of equations is solved for X using the factored form of A .

5. Iterative refinement is applied to improve the computed solution matrix and calculate error bounds and backward error estimates for it.

6. If equilibration was used, the matrix X is premultiplied by `diag(s)` so that it solves the original system before equilibration.

Input Parameters

`fact` CHARACTER*1. Must be 'F', 'N', or 'E'.

Specifies whether or not the factored form of the matrix A is supplied on entry, and if not, whether the matrix A should be equilibrated before it is factored.

If `fact = 'F'`: on entry, `af` contains the factored form of A . If `equed = 'Y'`, the matrix A has been equilibrated with scaling factors given by `s`.

`a` and `af` will not be modified.

If `fact = 'N'`, the matrix A will be copied to `af` and factored.

If `fact = 'E'`, the matrix A will be equilibrated if necessary, then copied to `af` and factored.

`uplo` CHARACTER*1. Must be 'U' or 'L'.

Indicates whether the upper or lower triangular part of A is stored and how A is factored:

If `uplo = 'U'`, the array `a` stores the upper triangular part of the matrix A , and A is factored as $U^H U$.

If `uplo = 'L'`, the array `a` stores the lower triangular part of the matrix A ; A is factored as LL^H .

<i>n</i>	INTEGER. The order of matrix <i>A</i> ($n \geq 0$).
<i>nrhs</i>	INTEGER. The number of right-hand sides; the number of columns in <i>B</i> ($nrhs \geq 0$).
<i>a, af, b, work</i>	REAL for <i>sposvx</i> DOUBLE PRECISION for <i>dposvx</i> COMPLEX for <i>cposvx</i> DOUBLE COMPLEX for <i>zposvx</i> . Arrays: <i>a(lda, *)</i> , <i>af(ldaf, *)</i> , <i>b(ldb, *)</i> , <i>work(*)</i> . The array <i>a</i> contains the matrix <i>A</i> as specified by <i>uplo</i> . If <i>fact</i> = 'F' and <i>equed</i> = 'Y', then <i>A</i> must have been equilibrated by the scaling factors in <i>s</i> , and <i>a</i> must contain the equilibrated matrix $\text{diag}(s) * A * \text{diag}(s)$. The second dimension of <i>a</i> must be at least $\max(1, n)$. The array <i>af</i> is an input argument if <i>fact</i> = 'F'. It contains the triangular factor <i>U</i> or <i>L</i> from the Cholesky factorization of <i>A</i> in the same storage format as <i>A</i> . If <i>equed</i> is not 'N', then <i>af</i> is the factored form of the equilibrated matrix $\text{diag}(s) * A * \text{diag}(s)$. The second dimension of <i>af</i> must be at least $\max(1, n)$. The array <i>b</i> contains the matrix <i>B</i> whose columns are the right-hand sides for the systems of equations. The second dimension of <i>b</i> must be at least $\max(1, nrhs)$. <i>work(*)</i> is a workspace array. The dimension of <i>work</i> must be at least $\max(1, 3 * n)$ for real flavors, and at least $\max(1, 2 * n)$ for complex flavors.
<i>lda</i>	INTEGER. The first dimension of <i>a</i> ; $lda \geq \max(1, n)$.
<i>ldaf</i>	INTEGER. The first dimension of <i>af</i> ; $ldaf \geq \max(1, n)$.
<i>ldb</i>	INTEGER. The first dimension of <i>b</i> ; $ldb \geq \max(1, n)$.
<i>equed</i>	CHARACTER*1. Must be 'N' or 'Y'. <i>equed</i> is an input argument if <i>fact</i> = 'F'. It specifies the form of equilibration that was done: If <i>equed</i> = 'N', no equilibration was done (always

true if *fact* = 'N');

If *equed* = 'Y', equilibration was done and *A* has been replaced by $\text{diag}(s) * A * \text{diag}(s)$.

s **REAL** for single precision flavors;
DOUBLE PRECISION for double precision flavors.
 Array, **DIMENSION** (*n*).
 The array *s* contains the scale factors for *A*. This array is an input argument if *fact* = 'F' only; otherwise it is an output argument.
 If *equed* = 'N', *s* is not accessed.
 If *fact* = 'F' and *equed* = 'Y', each element of *s* must be positive.

ldx **INTEGER**. The first dimension of the output array *x*;
 $ldx \geq \max(1, n)$.

iwork **INTEGER**.
 Workspace array, **DIMENSION** at least $\max(1, n)$; used in real flavors only.

rwork **REAL** for *cposvx*;
DOUBLE PRECISION for *zposvx*.
 Workspace array, **DIMENSION** at least $\max(1, n)$; used in complex flavors only.

Output Parameters

x **REAL** for *sposvx*
DOUBLE PRECISION for *dposvx*
COMPLEX for *cposvx*
DOUBLE COMPLEX for *zposvx*.
 Array, **DIMENSION** (*ldx*, *).
 If *info* = 0 or *info* = *n*+1, the array *x* contains the solution matrix *X* to the *original* system of equations. Note that if *equed* = 'Y', *A* and *B* are modified on exit, and the solution to the equilibrated system is $\text{diag}(s)^{-1} * X$.
 The second dimension of *x* must be at least $\max(1, nrhs)$.

<i>a</i>	Array <i>a</i> is not modified on exit if <i>fact</i> = 'F' or 'N', or if <i>fact</i> = 'E' and <i>equed</i> = 'N'. If <i>fact</i> = 'E' and <i>equed</i> = 'Y', <i>A</i> is overwritten by $\text{diag}(s)*A*\text{diag}(s)$
<i>af</i>	If <i>fact</i> = 'N' or 'E', then <i>af</i> is an output argument and on exit returns the triangular factor <i>U</i> or <i>L</i> from the Cholesky factorization $A=U^H U$ or $A=LL^H$ of the original matrix <i>A</i> (if <i>fact</i> = 'N'), or of the equilibrated matrix <i>A</i> (if <i>fact</i> = 'E'). See the description of <i>a</i> for the form of the equilibrated matrix.
<i>b</i>	Overwritten by $\text{diag}(s)*B$, if <i>equed</i> = 'Y'; not changed if <i>equed</i> = 'N'.
<i>s</i>	This array is an output argument if <i>fact</i> ≠ 'F'. See the description of <i>s</i> in <i>Input Arguments</i> section.
<i>rcond</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. An estimate of the reciprocal condition number of the matrix <i>A</i> after equilibration (if done). If <i>rcond</i> is less than the machine precision (in particular, if <i>rcond</i> = 0), the matrix is singular to working precision. This condition is indicated by a return code of <i>info</i> > 0.
<i>ferr, berr</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and relative backward errors, respectively, for each solution vector.
<i>equed</i>	If <i>fact</i> ≠ 'F', then <i>equed</i> is an output argument. It specifies the form of equilibration that was done (see the description of <i>equed</i> in <i>Input Arguments</i> section).
<i>info</i>	INTEGER. If <i>info</i> =0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = <i>i</i> , and $i \leq n$, the leading minor of order <i>i</i> (and hence the matrix <i>A</i> itself) is not positive definite, so the factorization could not be completed, and the solution and error bounds could not be computed; <i>rcond</i> = 0 is

returned.

If $info = i$, and $i = n + 1$, then U is nonsingular, but $rcond$ is less than machine precision, meaning that the matrix is singular to working precision. Nevertheless, the solution and error bounds are computed because there are a number of situations where the computed solution can be more accurate than the value of $rcond$ would suggest.

?ppsv

Computes the solution to the system of linear equations with a symmetric (Hermitian) positive definite packed matrix A and multiple right-hand sides.

```
call sppsv (uplo, n, nrhs, ap, b, ldb, info)
call dppsv (uplo, n, nrhs, ap, b, ldb, info)
call cppsv (uplo, n, nrhs, ap, b, ldb, info)
call zppsv (uplo, n, nrhs, ap, b, ldb, info)
```

Discussion

This routine solves for X the real or complex system of linear equations $AX = B$, where A is an n -by- n real symmetric/Hermitian positive definite matrix stored in packed format, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

The Cholesky decomposition is used to factor A as $A = U^H U$ if $uplo = 'U'$ or $A = LL^H$ if $uplo = 'L'$, where U is an upper triangular matrix and L is a lower triangular matrix. The factored form of A is then used to solve the system of equations $AX = B$.

Input Parameters

$uplo$ CHARACTER*1. Must be 'U' or 'L'.

	Indicates whether the upper or lower triangular part of A is stored and how A is factored: If $uplo = 'U'$, the array a stores the upper triangular part of the matrix A , and A is factored as $U^H U$. If $uplo = 'L'$, the array a stores the lower triangular part of the matrix A ; A is factored as LL^H .
n	INTEGER. The order of matrix A ($n \geq 0$).
$nrhs$	INTEGER. The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).
ap, b	REAL for <code>sppsv</code> DOUBLE PRECISION for <code>dppsv</code> COMPLEX for <code>cppsv</code> DOUBLE COMPLEX for <code>zppsv</code> . Arrays: $ap(*)$, $b(l db, *)$. The array ap contains either the upper or the lower triangular part of the matrix A (as specified by $uplo$) in <i>packed storage</i> (see Matrix Storage Schemes). The dimension of ap must be at least $\max(1, n(n+1)/2)$. The array b contains the matrix B whose columns are the right-hand sides for the systems of equations. The second dimension of b must be at least $\max(1, nrhs)$.
ldb	INTEGER. The first dimension of b ; $ldb \geq \max(1, n)$.

Output Parameters

ap	If $info=0$, the upper or lower triangular part of A in packed storage is overwritten by the Cholesky factor U or L , as specified by $uplo$.
b	Overwritten by the solution matrix X .
$info$	INTEGER. If $info=0$, the execution is successful. If $info = -i$, the i th parameter had an illegal value. If $info = i$, the leading minor of order i (and hence the matrix A itself) is not positive definite, so the factorization could not be completed, and the solution has not been computed.

?ppsvx

Uses the Cholesky factorization to compute the solution to the system of linear equations with a symmetric (Hermitian) positive definite packed matrix A , and provides error bounds on the solution.

```
call sppsvx (fact, uplo, n, nrhs, ap, AFP, equed, s, b, ldb, x, ldx,
            rcond, ferr, berr, work, iwork, info)
call dppsvx (fact, uplo, n, nrhs, ap, AFP, equed, s, b, ldb, x, ldx,
            rcond, ferr, berr, work, iwork, info)
call cppsvx (fact, uplo, n, nrhs, ap, AFP, equed, s, b, ldb, x, ldx,
            rcond, ferr, berr, work, rwork, info)
call zppsvx (fact, uplo, n, nrhs, ap, AFP, equed, s, b, ldb, x, ldx,
            rcond, ferr, berr, work, rwork, info)
```

Discussion

This routine uses the Cholesky factorization $A=U^H U$ or $A=LL^H$ to compute the solution to a real or complex system of linear equations $AX=B$, where A is a n -by- n symmetric or Hermitian positive definite matrix stored in packed format, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

Error bounds on the solution and a condition estimate are also provided.

The routine ?ppsvx performs the following steps:

1. If $fact = 'E'$, real scaling factors s are computed to equilibrate the system:

$$\text{diag}(s) * A * \text{diag}(s) * \text{diag}(s)^{-1} * X = \text{diag}(s) * B$$

Whether or not the system will be equilibrated depends on the scaling of the matrix A , but if equilibration is used, A is overwritten by $\text{diag}(s) * A * \text{diag}(s)$ and B by $\text{diag}(s) * B$.

2. If $fact = 'N'$ or $'E'$, the Cholesky decomposition is used to factor the matrix A (after equilibration if $fact = 'E'$) as

$A = U^H U$, if *uplo* = 'U', or

$A = L L^H$, if *uplo* = 'L',

where U is an upper triangular matrix and L is a lower triangular matrix.

3. If the leading i -by- i principal minor is not positive definite, then the routine returns with *info* = i . Otherwise, the factored form of A is used to estimate the condition number of the matrix A . If the reciprocal of the condition number is less than machine precision, *info* = $n + 1$ is returned as a warning, but the routine still goes on to solve for X and compute error bounds as described below.

4. The system of equations is solved for X using the factored form of A .

5. Iterative refinement is applied to improve the computed solution matrix and calculate error bounds and backward error estimates for it.

6. If equilibration was used, the matrix X is premultiplied by $\text{diag}(s)$ so that it solves the original system before equilibration.

Input Parameters

fact CHARACTER*1. Must be 'F', 'N', or 'E'.
 Specifies whether or not the factored form of the matrix A is supplied on entry, and if not, whether the matrix A should be equilibrated before it is factored.
 If *fact* = 'F': on entry, *afp* contains the factored form of A . If *equed* = 'Y', the matrix A has been equilibrated with scaling factors given by *s*.
ap and *afp* will not be modified.
 If *fact* = 'N', the matrix A will be copied to *afp* and factored.
 If *fact* = 'E', the matrix A will be equilibrated if necessary, then copied to *afp* and factored.

uplo CHARACTER*1. Must be 'U' or 'L'.
 Indicates whether the upper or lower triangular part of A is stored and how A is factored:
 If *uplo* = 'U', the array *ap* stores the upper triangular part of the matrix A , and A is factored as $U^H U$.
 If *uplo* = 'L', the array *ap* stores the lower triangular part of the matrix A ; A is factored as LL^H .

n INTEGER. The order of matrix *A* ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides; the number of columns in *B* ($nrhs \geq 0$).

ap,afp,b,work REAL for *sppsvx*
 DOUBLE PRECISION for *dppsvx*
 COMPLEX for *cppsvx*
 DOUBLE COMPLEX for *zppsvx*.
 Arrays: *ap*(*), *afp*(*), *b*(*ldb*,*), *work*(*).

The array *ap* contains the upper or lower triangle of the original symmetric/Hermitian matrix *A* in *packed storage* (see [Matrix Storage Schemes](#)). In case when *fact* = 'F' and *equed* = 'Y', *ap* must contain the equilibrated matrix $\text{diag}(s)*A*\text{diag}(s)$.

The array *afp* is an input argument if *fact* = 'F' and contains the triangular factor *U* or *L* from the Cholesky factorization of *A* in the same storage format as *A*. If *equed* is not 'N', then *afp* is the factored form of the equilibrated matrix *A*.

The array *b* contains the matrix *B* whose columns are the right-hand sides for the systems of equations.

work(*) is a workspace array.

The dimension of arrays *ap* and *afp* must be at least $\max(1, n(n+1)/2)$; the second dimension of *b* must be at least $\max(1, nrhs)$; the dimension of *work* must be at least $\max(1, 3*n)$ for real flavors and $\max(1, 2*n)$ for complex flavors.

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

equed CHARACTER*1. Must be 'N' or 'Y'.
equed is an input argument if *fact* = 'F'. It specifies the form of equilibration that was done:
 If *equed* = 'N', no equilibration was done (always true if *fact* = 'N');
 If *equed* = 'Y', equilibration was done and *A* has been replaced by $\text{diag}(s)*A*\text{diag}(s)$.

<i>s</i>	<p>REAL for single precision flavors; DOUBLE PRECISION for double precision flavors. Array, DIMENSION (<i>n</i>). The array <i>s</i> contains the scale factors for <i>A</i>. This array is an input argument if <i>fact</i> = 'F' only; otherwise it is an output argument. If <i>equed</i> = 'N' , <i>s</i> is not accessed. If <i>fact</i> = 'F' and <i>equed</i> = 'Y' , each element of <i>s</i> must be positive.</p>
<i>ldx</i>	<p>INTEGER. The first dimension of the output array <i>x</i>; $ldx \geq \max(1, n)$.</p>
<i>iwork</i>	<p>INTEGER. Workspace array, DIMENSION at least $\max(1, n)$; used in real flavors only.</p>
<i>rwork</i>	<p>REAL for <i>cppsvox</i>; DOUBLE PRECISION for <i>zppsvox</i>. Workspace array, DIMENSION at least $\max(1, n)$; used in complex flavors only.</p>

Output Parameters

<i>x</i>	<p>REAL for <i>sppsvox</i> DOUBLE PRECISION for <i>dppsvox</i> COMPLEX for <i>cppsvox</i> DOUBLE COMPLEX for <i>zppsvox</i>. Array, DIMENSION (<i>ldx</i>, *). If <i>info</i> = 0 or <i>info</i> = <i>n</i>+1, the array <i>x</i> contains the solution matrix <i>X</i> to the <i>original</i> system of equations. Note that if <i>equed</i> = 'Y' , <i>A</i> and <i>B</i> are modified on exit, and the solution to the equilibrated system is $\text{diag}(s)^{-1} * X$. The second dimension of <i>x</i> must be at least $\max(1, nrhs)$.</p>
<i>ap</i>	<p>Array <i>ap</i> is not modified on exit if <i>fact</i> = 'F' or 'N', or if <i>fact</i> = 'E' and <i>equed</i> = 'N'. If <i>fact</i> = 'E' and <i>equed</i> = 'Y' , <i>A</i> is overwritten by $\text{diag}(s) * A * \text{diag}(s)$</p>

<i>afp</i>	If <i>fact</i> = 'N' or 'E', then <i>afp</i> is an output argument and on exit returns the triangular factor <i>U</i> or <i>L</i> from the Cholesky factorization $A=U^H U$ or $A=LL^H$ of the original matrix <i>A</i> (if <i>fact</i> = 'N'), or of the equilibrated matrix <i>A</i> (if <i>fact</i> = 'E'). See the description of <i>ap</i> for the form of the equilibrated matrix.
<i>b</i>	Overwritten by $\text{diag}(s)*B$, if <i>equed</i> = 'Y'; not changed if <i>equed</i> = 'N'.
<i>s</i>	This array is an output argument if <i>fact</i> ≠ 'F'. See the description of <i>s</i> in <i>Input Arguments</i> section.
<i>rcond</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. An estimate of the reciprocal condition number of the matrix <i>A</i> after equilibration (if done). If <i>rcond</i> is less than the machine precision (in particular, if <i>rcond</i> = 0), the matrix is singular to working precision. This condition is indicated by a return code of <i>info</i> > 0.
<i>ferr, berr</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and relative backward errors, respectively, for each solution vector.
<i>equed</i>	If <i>fact</i> ≠ 'F', then <i>equed</i> is an output argument. It specifies the form of equilibration that was done (see the description of <i>equed</i> in <i>Input Arguments</i> section).
<i>info</i>	INTEGER. If <i>info</i> =0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = <i>i</i> , and $i \leq n$, the leading minor of order <i>i</i> (and hence the matrix <i>A</i> itself) is not positive definite, so the factorization could not be completed, and the solution and error bounds could not be computed; <i>rcond</i> = 0 is returned. If <i>info</i> = <i>i</i> , and $i = n + 1$, then <i>U</i> is nonsingular, but <i>rcond</i> is less than machine precision, meaning that the matrix is singular to working precision. Nevertheless, the solution and error bounds are computed because there are a number of situations where the computed solution can be more accurate than the value of <i>rcond</i> would suggest.

?pbsv

Computes the solution to the system of linear equations with a symmetric or Hermitian positive definite band matrix A and multiple right-hand sides.

```
call spbsv (uplo, n, kd, nrhs, ab, ldab, b, ldb, info)
call dpbsv (uplo, n, kd, nrhs, ab, ldab, b, ldb, info)
call cpbsv (uplo, n, kd, nrhs, ab, ldab, b, ldb, info)
call zpbsv (uplo, n, kd, nrhs, ab, ldab, b, ldb, info)
```

Discussion

This routine solves for X the real or complex system of linear equations $AX = B$, where A is an n -by- n symmetric/Hermitian positive definite band matrix, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

The Cholesky decomposition is used to factor A as $A = U^H U$ if `uplo = 'U'` or $A = LL^H$ if `uplo = 'L'`, where U is an upper triangular band matrix and L is a lower triangular band matrix, with the same number of superdiagonals or subdiagonals as A . The factored form of A is then used to solve the system of equations $AX = B$.

Input Parameters

<code>uplo</code>	CHARACTER*1. Must be 'U' or 'L'. Indicates whether the upper or lower triangular part of A is stored in the array <code>ab</code> , and how A is factored: If <code>uplo = 'U'</code> , the array <code>ab</code> stores the upper triangular part of the matrix A , and A is factored as $U^H U$. If <code>uplo = 'L'</code> , the array <code>ab</code> stores the lower triangular part of the matrix A ; A is factored as LL^H .
<code>n</code>	INTEGER. The order of matrix A ($n \geq 0$).

<i>kd</i>	INTEGER. The number of superdiagonals of the matrix <i>A</i> if <i>uplo</i> = 'U', or the number of subdiagonals if <i>uplo</i> = 'L' (<i>kd</i> ≥ 0).
<i>nrhs</i>	INTEGER. The number of right-hand sides; the number of columns in <i>B</i> (<i>nrhs</i> ≥ 0).
<i>ab, b</i>	REAL for <i>spbsv</i> DOUBLE PRECISION for <i>dpbsv</i> COMPLEX for <i>cpbsv</i> DOUBLE COMPLEX for <i>zpbsv</i> . Arrays: <i>ab</i> (<i>ldab</i> , *), <i>b</i> (<i>ldb</i> , *). The array <i>ab</i> contains either the upper or the lower triangular part of the matrix <i>A</i> (as specified by <i>uplo</i>) in <i>band storage</i> (see Matrix Storage Schemes). The second dimension of <i>ab</i> must be at least max(1, <i>n</i>). The array <i>b</i> contains the matrix <i>B</i> whose columns are the right-hand sides for the systems of equations. The second dimension of <i>b</i> must be at least max(1, <i>nrhs</i>).
<i>ldab</i>	INTEGER. The first dimension of the array <i>ab</i> . (<i>ldab</i> ≥ <i>kd</i> + 1)
<i>ldb</i>	INTEGER. The first dimension of <i>b</i> ; <i>ldb</i> ≥ max(1, <i>n</i>).

Output Parameters

<i>ab</i>	The upper or lower triangular part of <i>A</i> (in band storage) is overwritten by the Cholesky factor <i>U</i> or <i>L</i> , as specified by <i>uplo</i> , in the same storage format as <i>A</i> .
<i>b</i>	Overwritten by the solution matrix <i>X</i> .
<i>info</i>	INTEGER. If <i>info</i> =0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = <i>i</i> , the leading minor of order <i>i</i> (and hence the matrix <i>A</i> itself) is not positive definite, so the factorization could not be completed, and the solution has not been computed.

?pbsvx

Uses the Cholesky factorization to compute the solution to the system of linear equations with a symmetric (Hermitian) positive definite band matrix A , and provides error bounds on the solution.

```
call spbsvx (fact, uplo, n, kd, nrhs, ab, ldab, afb, ldafb, equed,
            s, b, ldb, x, ldx, rcond, ferr, berr, work, iwork, info)
call dpbsvx (fact, uplo, n, kd, nrhs, ab, ldab, afb, ldafb, equed,
            s, b, ldb, x, ldx, rcond, ferr, berr, work, iwork, info)
call cpbsvx (fact, uplo, n, kd, nrhs, ab, ldab, afb, ldafb, equed,
            s, b, ldb, x, ldx, rcond, ferr, berr, work, iwork, info)
call zpbsvx (fact, uplo, n, kd, nrhs, ab, ldab, afb, ldafb, equed,
            s, b, ldb, x, ldx, rcond, ferr, berr, work, iwork, info)
```

Discussion

This routine uses the Cholesky factorization $A=U^H U$ or $A=LL^H$ to compute the solution to a real or complex system of linear equations $AX=B$, where A is a n -by- n symmetric or Hermitian positive definite band matrix, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

Error bounds on the solution and a condition estimate are also provided.

The routine ?pbsvx performs the following steps:

1. If $fact = 'E'$, real scaling factors s are computed to equilibrate the system:

$$\text{diag}(s) * A * \text{diag}(s) * \text{diag}(s)^{-1} * X = \text{diag}(s) * B$$

Whether or not the system will be equilibrated depends on the scaling of the matrix A , but if equilibration is used, A is overwritten by $\text{diag}(s) * A * \text{diag}(s)$ and B by $\text{diag}(s) * B$.

2. If $fact = 'N'$ or $'E'$, the Cholesky decomposition is used to factor the matrix A (after equilibration if $fact = 'E'$) as

$A = U^H U$, if *uplo* = 'U', or
 $A = L L^H$, if *uplo* = 'L',

where U is an upper triangular band matrix and L is a lower triangular band matrix.

3. If the leading i -by- i principal minor is not positive definite, then the routine returns with *info* = i . Otherwise, the factored form of A is used to estimate the condition number of the matrix A . If the reciprocal of the condition number is less than machine precision, *info* = $n + 1$ is returned as a warning, but the routine still goes on to solve for X and compute error bounds as described below.

4. The system of equations is solved for X using the factored form of A .

5. Iterative refinement is applied to improve the computed solution matrix and calculate error bounds and backward error estimates for it.

6. If equilibration was used, the matrix X is premultiplied by $\text{diag}(s)$ so that it solves the original system before equilibration.

Input Parameters

fact CHARACTER*1. Must be 'F', 'N', or 'E'.

Specifies whether or not the factored form of the matrix A is supplied on entry, and if not, whether the matrix A should be equilibrated before it is factored.

If *fact* = 'F': on entry, *afb* contains the factored form of A . If *equed* = 'Y', the matrix A has been equilibrated with scaling factors given by *s*.

ab and *afb* will not be modified.

If *fact* = 'N', the matrix A will be copied to *afb* and factored.

If *fact* = 'E', the matrix A will be equilibrated if necessary, then copied to *afb* and factored.

uplo CHARACTER*1. Must be 'U' or 'L'.

Indicates whether the upper or lower triangular part of A is stored and how A is factored:

If *uplo* = 'U', the array *ab* stores the upper triangular part of the matrix *A*, and *A* is factored as $U^H U$.

If *uplo* = 'L', the array *ab* stores the lower triangular part of the matrix *A*; *A* is factored as LL^H .

n **INTEGER.** The order of matrix *A* ($n \geq 0$).

kd **INTEGER.** The number of super-diagonals or sub-diagonals in the matrix *A* ($kd \geq 0$).

nrhs **INTEGER.** The number of right-hand sides; the number of columns in *B* ($nrhs \geq 0$).

ab,afb,b,work **REAL** for *spbsvx*
DOUBLE PRECISION for *dpbsvx*
COMPLEX for *cpbsvx*
DOUBLE COMPLEX for *zpbsvx*.
 Arrays: *ab(ldab,*)*, *afb(ldab,*)*, *b(ldb,*)*, *work(*)*.

The array *ab* contains the upper or lower triangle of the matrix *A* in band storage (see [Matrix Storage Schemes](#)).

If *fact* = 'F' and *equed* = 'Y', then *ab* must contain the equilibrated matrix $\text{diag}(s) * A * \text{diag}(s)$. The second dimension of *ab* must be at least $\max(1, n)$.

The array *afb* is an input argument if *fact* = 'F'. It contains the triangular factor *U* or *L* from the Cholesky factorization of the band matrix *A* in the same storage format as *A*. If *equed* = 'Y', then *afb* is the factored form of the equilibrated matrix *A*.

The second dimension of *afb* must be at least $\max(1, n)$.

The array *b* contains the matrix *B* whose columns are the right-hand sides for the systems of equations. The second dimension of *b* must be at least $\max(1, nrhs)$.

work()* is a workspace array.

The dimension of *work* must be at least $\max(1, 3 * n)$ for real flavors, and at least $\max(1, 2 * n)$ for complex flavors.

ldab **INTEGER.** The first dimension of *ab*; $ldab \geq kd + 1$.

ldafb **INTEGER.** The first dimension of *afb*; $ldafb \geq kd + 1$.

<i>ldb</i>	INTEGER . The first dimension of <i>b</i> ; $ldb \geq \max(1, n)$.
<i>equed</i>	CHARACTER*1 . Must be 'N' or 'Y'. <i>equed</i> is an input argument if <i>fact</i> = 'F'. It specifies the form of equilibration that was done: If <i>equed</i> = 'N', no equilibration was done (always true if <i>fact</i> = 'N'); If <i>equed</i> = 'Y', equilibration was done and A has been replaced by $\text{diag}(s) * A * \text{diag}(s)$.
<i>s</i>	REAL for single precision flavors; DOUBLE PRECISION for double precision flavors. Array, DIMENSION (<i>n</i>). The array <i>s</i> contains the scale factors for A. This array is an input argument if <i>fact</i> = 'F' only; otherwise it is an output argument. If <i>equed</i> = 'N', <i>s</i> is not accessed. If <i>fact</i> = 'F' and <i>equed</i> = 'Y', each element of <i>s</i> must be positive.
<i>ldx</i>	INTEGER . The first dimension of the output array <i>x</i> ; $ldx \geq \max(1, n)$.
<i>iwork</i>	INTEGER . Workspace array, DIMENSION at least $\max(1, n)$; used in real flavors only.
<i>rwork</i>	REAL for <i>spbsvx</i> ; DOUBLE PRECISION for <i>zpbsvx</i> . Workspace array, DIMENSION at least $\max(1, n)$; used in complex flavors only.

Output Parameters

<i>x</i>	REAL for <i>spbsvx</i> DOUBLE PRECISION for <i>dpbsvx</i> COMPLEX for <i>cpbsvx</i> DOUBLE COMPLEX for <i>zpbsvx</i> . Array, DIMENSION (<i>ldx</i> , *).
----------	---

- If `info = 0` or `info = n+1`, the array `x` contains the solution matrix `X` to the *original* system of equations. Note that if `equed = 'Y'`, `A` and `B` are modified on exit, and the solution to the equilibrated system is $\text{diag}(s)^{-1} * X$.
- The second dimension of `x` must be at least $\max(1, nrhs)$.
- `ab` On exit, if `fact = 'E'` and `equed = 'Y'`, `A` is overwritten by $\text{diag}(s) * A * \text{diag}(s)$
- `afb` If `fact = 'N'` or `'E'`, then `afb` is an output argument and on exit returns the triangular factor `U` or `L` from the Cholesky factorization $A = U^H U$ or $A = LL^H$ of the original matrix `A` (if `fact = 'N'`), or of the equilibrated matrix `A` (if `fact = 'E'`). See the description of `ab` for the form of the equilibrated matrix.
- `b` Overwritten by $\text{diag}(s) * B$, if `equed = 'Y'`; not changed if `equed = 'N'`.
- `s` This array is an output argument if `fact ≠ 'F'`. See the description of `s` in *Input Arguments* section.
- `rcond` **REAL** for single precision flavors.
DOUBLE PRECISION for double precision flavors.
An estimate of the reciprocal condition number of the matrix `A` after equilibration (if done). If `rcond` is less than the machine precision (in particular, if `rcond = 0`), the matrix is singular to working precision. This condition is indicated by a return code of `info > 0`.
- `ferr, berr` **REAL** for single precision flavors.
DOUBLE PRECISION for double precision flavors.
Arrays, **DIMENSION** at least $\max(1, nrhs)$. Contain the component-wise forward and relative backward errors, respectively, for each solution vector.
- `equed` If `fact ≠ 'F'`, then `equed` is an output argument. It specifies the form of equilibration that was done (see the description of `equed` in *Input Arguments* section).

info INTEGER. If *info*=0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i*, and $i \leq n$, the leading minor of order *i* (and hence the matrix *A* itself) is not positive definite, so the factorization could not be completed, and the solution and error bounds could not be computed; *rcond* = 0 is returned.
If *info* = *i*, and $i = n + 1$, then *U* is nonsingular, but *rcond* is less than machine precision, meaning that the matrix is singular to working precision. Nevertheless, the solution and error bounds are computed because there are a number of situations where the computed solution can be more accurate than the value of *rcond* would suggest.

?ptsv

Computes the solution to the system of linear equations with a symmetric or Hermitian positive definite tridiagonal matrix A and multiple right-hand sides.

```
call sptsv (n, nrhs, d, e, b, ldb, info)
call dptsv (n, nrhs, d, e, b, ldb, info)
call cptsv (n, nrhs, d, e, b, ldb, info)
call zptsv (n, nrhs, d, e, b, ldb, info)
```

Discussion

This routine solves for *X* the real or complex system of linear equations $AX = B$, where *A* is an *n*-by-*n* symmetric/Hermitian positive definite tridiagonal matrix, the columns of matrix *B* are individual right-hand sides, and the columns of *X* are the corresponding solutions.

A is factored as $A = L D L^H$, and the factored form of *A* is then used to solve the system of equations $AX = B$.

Input Parameters

<i>n</i>	INTEGER. The order of matrix A ($n \geq 0$).
<i>nrhs</i>	INTEGER. The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).
<i>d</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. Array, dimension at least $\max(1, n)$. Contains the diagonal elements of the tridiagonal matrix A .
<i>e, b</i>	REAL for sptsv DOUBLE PRECISION for dptsv COMPLEX for cptsv DOUBLE COMPLEX for zptsv . Arrays: $e(n - 1)$, $b(ldb, *)$. The array e contains the $(n - 1)$ subdiagonal elements of A . The array b contains the matrix B whose columns are the right-hand sides for the systems of equations. The second dimension of b must be at least $\max(1, nrhs)$.
<i>ldb</i>	INTEGER. The first dimension of b ; $ldb \geq \max(1, n)$.

Output Parameters

<i>d</i>	Overwritten by the n diagonal elements of the diagonal matrix D from the LDL^H factorization of A .
<i>e</i>	Overwritten by the $(n - 1)$ subdiagonal elements of the unit bidiagonal factor L from the factorization of A .
<i>b</i>	Overwritten by the solution matrix X .
<i>info</i>	INTEGER. If $info=0$, the execution is successful. If $info = -i$, the i th parameter had an illegal value. If $info = i$, the leading minor of order i (and hence the matrix A itself) is not positive definite, and the solution has not been computed. The factorization has not been completed unless $i = n$.

?ptsvx

Uses the factorization $A=LDL^H$ to compute the solution to the system of linear equations with a symmetric (Hermitian) positive definite tridiagonal matrix A , and provides error bounds on the solution.

```
call sptsvx (fact, n, nrhs, d, e, df, ef, b, ldb, x, ldx, rcond,
            ferr, berr, work, info)
call dptsvx (fact, n, nrhs, d, e, df, ef, b, ldb, x, ldx, rcond,
            ferr, berr, work, info)
call cptsvx (fact, n, nrhs, d, e, df, ef, b, ldb, x, ldx, rcond,
            ferr, berr, work, rwork, info)
call zptsvx (fact, n, nrhs, d, e, df, ef, b, ldb, x, ldx, rcond,
            ferr, berr, work, rwork, info)
```

Discussion

This routine uses the Cholesky factorization $A=LDL^H$ to compute the solution to a real or complex system of linear equations $AX=B$, where A is a n -by- n symmetric or Hermitian positive definite tridiagonal matrix, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

Error bounds on the solution and a condition estimate are also provided.

The routine ?ptsvx performs the following steps:

1. If `fact = 'N'`, the matrix A is factored as $A = LD L^H$, where L is a unit lower bidiagonal matrix and D is diagonal. The factorization can also be regarded as having the form $A = U^H D U$.
2. If the leading i -by- i principal minor is not positive definite, then the routine returns with `info = i`. Otherwise, the factored form of A is used to estimate the condition number of the matrix A . If the reciprocal of the condition number is less than machine precision, `info = n + 1` is returned as a warning, but the routine still goes on to solve for X and compute error bounds as described below.

3. The system of equations is solved for X using the factored form of A .
4. Iterative refinement is applied to improve the computed solution matrix and calculate error bounds and backward error estimates for it.

Input Parameters

<i>fact</i>	<p>CHARACTER*1. Must be 'F' or 'N'.</p> <p>Specifies whether or not the factored form of the matrix A is supplied on entry.</p> <p>If <i>fact</i> = 'F': on entry, <i>df</i> and <i>ef</i> contain the factored form of A. Arrays <i>d</i>, <i>e</i>, <i>df</i>, and <i>ef</i> will not be modified.</p> <p>If <i>fact</i> = 'N', the matrix A will be copied to <i>df</i> and <i>ef</i> and factored.</p>
<i>n</i>	<p>INTEGER. The order of matrix A ($n \geq 0$).</p>
<i>nrhs</i>	<p>INTEGER. The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).</p>
<i>d,df,rwork</i>	<p>REAL for single precision flavors DOUBLE PRECISION for double precision flavors Arrays: <i>d</i>(n), <i>df</i>(n), <i>rwork</i>(n).</p> <p>The array <i>d</i> contains the n diagonal elements of the tridiagonal matrix A.</p> <p>The array <i>df</i> is an input argument if <i>fact</i> = 'F' and on entry contains the n diagonal elements of the diagonal matrix D from the LDL^H factorization of A.</p> <p>The array <i>rwork</i> is a workspace array used for complex flavors only.</p>
<i>e,ef,b,work</i>	<p>REAL for sptsvx DOUBLE PRECISION for dptsvx COMPLEX for cptsvx DOUBLE COMPLEX for zptsvx. Arrays: <i>e</i>($n - 1$), <i>ef</i>($n - 1$), <i>b</i>(<i>ldb</i>,*), <i>work</i>(*). The array <i>e</i> contains the ($n - 1$) subdiagonal elements of the tridiagonal matrix A.</p>

The array *ef* is an input argument if *fact* = 'F' and on entry contains the $(n - 1)$ subdiagonal elements of the unit bidiagonal factor L from the LDL^H factorization of A .

The array *b* contains the matrix B whose columns are the right-hand sides for the systems of equations.

The array *work* is a workspace array. The dimension of *work* must be at least $2 * n$ for real flavors, and at least n for complex flavors.

ldb INTEGER. The leading dimension of *b*; $ldb \geq \max(1, n)$.

ldx INTEGER. The leading dimension of *x*; $ldx \geq \max(1, n)$.

Output Parameters

x REAL for *sptsvx*
 DOUBLE PRECISION for *dptsvx*
 COMPLEX for *cptsvx*
 DOUBLE COMPLEX for *zptsvx*.
 Array, DIMENSION (*ldx*, *).

If *info* = 0 or *info* = $n + 1$, the array *x* contains the solution matrix X to the system of equations. The second dimension of *x* must be at least $\max(1, nrhs)$.

df, *ef* These arrays are output arguments if *fact* = 'N'. See the description of *df*, *ef* in *Input Arguments* section.

rcond REAL for single precision flavors.
 DOUBLE PRECISION for double precision flavors.
 An estimate of the reciprocal condition number of the matrix A after equilibration (if done). If *rcond* is less than the machine precision (in particular, if *rcond* = 0), the matrix is singular to working precision. This condition is indicated by a return code of *info* > 0.

ferr, *berr* REAL for single precision flavors.
 DOUBLE PRECISION for double precision flavors.
 Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and relative backward errors, respectively, for each solution vector.

info INTEGER. If *info*=0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.
 If *info* = *i*, and $i \leq n$, the leading minor of order *i* (and hence the matrix *A* itself) is not positive definite, so the factorization could not be completed, and the solution and error bounds could not be computed; *rcond* = 0 is returned.
 If *info* = *i*, and $i = n + 1$, then *U* is nonsingular, but *rcond* is less than machine precision, meaning that the matrix is singular to working precision. Nevertheless, the solution and error bounds are computed because there are a number of situations where the computed solution can be more accurate than the value of *rcond* would suggest.

?sysv

Computes the solution to the system of linear equations with a real or complex symmetric matrix A and multiple right-hand sides.

```
call ssysv (uplo, n, nrhs, a, lda, ipiv, b, ldb, work, lwork, info)
call dsysv (uplo, n, nrhs, a, lda, ipiv, b, ldb, work, lwork, info)
call csysv (uplo, n, nrhs, a, lda, ipiv, b, ldb, work, lwork, info)
call zsysv (uplo, n, nrhs, a, lda, ipiv, b, ldb, work, lwork, info)
```

Discussion

This routine solves for *X* the real or complex system of linear equations $AX = B$, where *A* is an *n*-by-*n* symmetric matrix, the columns of matrix *B* are individual right-hand sides, and the columns of *X* are the corresponding solutions.

The diagonal pivoting method is used to factor A as $A = U D U^T$ or $A = L D L^T$, where U (or L) is a product of permutation and unit upper (lower) triangular matrices, and D is symmetric and block diagonal with 1-by-1 and 2-by-2 diagonal blocks.

The factored form of A is then used to solve the system of equations $AX = B$.

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates whether the upper or lower triangular part of A is stored and how A is factored:
If *uplo* = 'U', the array *a* stores the upper triangular part of the matrix A , and A is factored as UDU^T .
If *uplo* = 'L', the array *a* stores the lower triangular part of the matrix A ; A is factored as LDL^T .

n INTEGER. The order of matrix A ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).

a, *b*, *work* REAL for *ssysv*
DOUBLE PRECISION for *dsysv*
COMPLEX for *csysv*
DOUBLE COMPLEX for *zsysv*.
Arrays: *a*(*lda*, *), *b*(*ldb*, *), *work*(*lwork*).
The array *a* contains either the upper or the lower triangular part of the symmetric matrix A (see *uplo*).
The second dimension of *a* must be at least $\max(1, n)$.
The array *b* contains the matrix B whose columns are the right-hand sides for the systems of equations.
The second dimension of *b* must be at least $\max(1, nrhs)$.
work(*lwork*) is a workspace array.

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, n)$.

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

lwork INTEGER. The size of the *work* array ($lwork \geq 1$)
See *Application notes* for the suggested value of *lwork*.

Output Parameters

<i>a</i>	If <i>info</i> = 0, <i>a</i> is overwritten by the block-diagonal matrix <i>D</i> and the multipliers used to obtain the factor <i>U</i> (or <i>L</i>) from the factorization of <i>A</i> as computed by ?sytrf .
<i>b</i>	If <i>info</i> = 0, <i>b</i> is overwritten by the solution matrix <i>X</i> .
<i>ipiv</i>	<p>INTEGER.</p> <p>Array, DIMENSION at least max(1,<i>n</i>).</p> <p>Contains details of the interchanges and the block structure of <i>D</i>, as determined by ?sytrf.</p> <p>If <i>ipiv</i>(<i>i</i>) = <i>k</i> > 0, then <i>d_{ii}</i> is a 1-by-1 diagonal block, and the <i>i</i>th row and column of <i>A</i> was interchanged with the <i>k</i>th row and column.</p> <p>If <i>uplo</i> = 'U' and <i>ipiv</i>(<i>i</i>) = <i>ipiv</i>(<i>i</i>-1) = -<i>m</i> < 0, then <i>D</i> has a 2-by-2 block in rows/columns <i>i</i> and <i>i</i>-1, and (<i>i</i>-1)th row and column of <i>A</i> was interchanged with the <i>m</i>th row and column.</p> <p>If <i>uplo</i> = 'L' and <i>ipiv</i>(<i>i</i>) = <i>ipiv</i>(<i>i</i>+1) = -<i>m</i> < 0, then <i>D</i> has a 2-by-2 block in rows/columns <i>i</i> and <i>i</i>+1, and (<i>i</i>+1)th row and column of <i>A</i> was interchanged with the <i>m</i>th row and column.</p>
<i>work</i> (1)	If <i>info</i> =0, on exit <i>work</i> (1) contains the minimum value of <i>lwork</i> required for optimum performance. Use this <i>lwork</i> for subsequent runs.
<i>info</i>	<p>INTEGER. If <i>info</i>=0, the execution is successful.</p> <p>If <i>info</i> = -<i>i</i>, the <i>i</i>th parameter had an illegal value.</p> <p>If <i>info</i> = <i>i</i>, <i>d_{ii}</i> is 0. The factorization has been completed, but <i>D</i> is exactly singular, so the solution could not be computed.</p>

Application Notes

For better performance, try using *lwork* = *n***blocksize*, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*.

If you are in doubt how much workspace to supply, use *lwork* = -1 for the first run. In this case, a workspace query is assumed; the routine only calculates the optimal size of the *work* array, returns this value as the first

entry `work(1)` of the `work` array, and no error message related to `lwork` is issued by XERBLA. On exit, examine `work(1)` and use this value for subsequent runs.

?sysvx

Uses the diagonal pivoting factorization to compute the solution to the system of linear equations with a real or complex symmetric matrix A, and provides error bounds on the solution.

```
call ssysvx (fact, uplo, n, nrhs, a, lda, af, ldaf, ipiv, b, ldb,  
            x, ldx, rcond, ferr, berr, work, lwork, iwork, info)  
call dsysvx (fact, uplo, n, nrhs, a, lda, af, ldaf, ipiv, b, ldb,  
            x, ldx, rcond, ferr, berr, work, lwork, iwork, info)  
call csysvx (fact, uplo, n, nrhs, a, lda, af, ldaf, ipiv, b, ldb,  
            x, ldx, rcond, ferr, berr, work, lwork, rwork, info)  
call zsysvx (fact, uplo, n, nrhs, a, lda, af, ldaf, ipiv, b, ldb,  
            x, ldx, rcond, ferr, berr, work, lwork, rwork, info)
```

Discussion

This routine uses the diagonal pivoting factorization to compute the solution to a real or complex system of linear equations $AX = B$, where A is a n -by- n symmetric matrix, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

Error bounds on the solution and a condition estimate are also provided.

The routine ?sysvx performs the following steps:

1. If `fact = 'N'`, the diagonal pivoting method is used to factor the matrix A. The form of the factorization is $A = U D U^T$ or $A = L D L^T$, where U (or L) is a product of permutation and unit upper (lower) triangular matrices, and D is symmetric and block diagonal with 1-by-1 and 2-by-2 diagonal blocks.

2. If some $d_{i,i} = 0$, so that D is exactly singular, then the routine returns with $info = i$. Otherwise, the factored form of A is used to estimate the condition number of the matrix A . If the reciprocal of the condition number is less than machine precision, $info = n + 1$ is returned as a warning, but the routine still goes on to solve for X and compute error bounds as described below.
3. The system of equations is solved for X using the factored form of A .
4. Iterative refinement is applied to improve the computed solution matrix and calculate error bounds and backward error estimates for it.

Input Parameters

<i>fact</i>	<p>CHARACTER*1. Must be 'F' or 'N'.</p> <p>Specifies whether or not the factored form of the matrix A has been supplied on entry.</p> <p>If <i>fact</i> = 'F': on entry, <i>af</i> and <i>ipiv</i> contain the factored form of A. Arrays <i>a</i>, <i>af</i>, and <i>ipiv</i> will not be modified.</p> <p>If <i>fact</i> = 'N', the matrix A will be copied to <i>af</i> and factored.</p>
<i>uplo</i>	<p>CHARACTER*1. Must be 'U' or 'L'.</p> <p>Indicates whether the upper or lower triangular part of A is stored and how A is factored:</p> <p>If <i>uplo</i> = 'U', the array <i>a</i> stores the upper triangular part of the symmetric matrix A, and A is factored as UDU^T.</p> <p>If <i>uplo</i> = 'L', the array <i>a</i> stores the lower triangular part of the symmetric matrix A; A is factored as LDL^T.</p>
<i>n</i>	INTEGER. The order of matrix A ($n \geq 0$).
<i>nrhs</i>	INTEGER. The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).
<i>a, af, b, work</i>	<p>REAL for <i>ssysvx</i></p> <p>DOUBLE PRECISION for <i>dsysvx</i></p> <p>COMPLEX for <i>csysvx</i></p>

DOUBLE COMPLEX for `zsysvx`.

Arrays: `a(lda,*)`, `af(ldaf,*)`, `b ldb,*)`,
`work(*)`.

The array `a` contains either the upper or the lower triangular part of the symmetric matrix A (see `uplo`). The second dimension of `a` must be at least $\max(1,n)$.

The array `af` is an input argument if `fact = 'F'`. It contains the block diagonal matrix D and the multipliers used to obtain the factor U or L from the factorization $A = U D U^T$ or $A = L D L^T$ as computed by `?sytrf`. The second dimension of `af` must be at least $\max(1,n)$.

The array `b` contains the matrix B whose columns are the right-hand sides for the systems of equations. The second dimension of `b` must be at least $\max(1,nrhs)$.

`work(*)` is a workspace array of dimension `(lwork)`.

`lda` INTEGER. The first dimension of `a`; $lda \geq \max(1, n)$.
`ldaf` INTEGER. The first dimension of `af`; $ldaf \geq \max(1, n)$.
`ldb` INTEGER. The first dimension of `b`; $ldb \geq \max(1, n)$.
`ipiv` INTEGER.

Array, DIMENSION at least $\max(1,n)$.

The array `ipiv` is an input argument if `fact = 'F'`. It contains details of the interchanges and the block structure of D , as determined by `?sytrf`.

If `ipiv(i) = k > 0`, then d_{ii} is a 1-by-1 diagonal block, and the i th row and column of A was interchanged with the k th row and column.

If `uplo = 'U'` and `ipiv(i) = ipiv(i-1) = -m < 0`, then D has a 2-by-2 block in rows/columns i and $i-1$, and $(i-1)$ th row and column of A was interchanged with the m th row and column.

If `uplo = 'L'` and `ipiv(i) = ipiv(i+1) = -m < 0`, then D has a 2-by-2 block in rows/columns i and $i+1$, and $(i+1)$ th row and column of A was interchanged with the m th row and column.

<i>ldx</i>	INTEGER . The leading dimension of the output array <i>x</i> ; $ldx \geq \max(1, n)$.
<i>lwork</i>	INTEGER . The size of the <i>work</i> array . See <i>Application notes</i> for the suggested value of <i>lwork</i> .
<i>iwork</i>	INTEGER . Workspace array, DIMENSION at least $\max(1, n)$; used in real flavors only.
<i>rwork</i>	REAL for <i>csysvx</i> ; DOUBLE PRECISION for <i>zsysvx</i> . Workspace array, DIMENSION at least $\max(1, n)$; used in complex flavors only.

Output Parameters

<i>x</i>	REAL for <i>ssysvx</i> DOUBLE PRECISION for <i>dsysvx</i> COMPLEX for <i>csysvx</i> DOUBLE COMPLEX for <i>zsysvx</i> . Array, DIMENSION (<i>ldx</i> , *). If <i>info</i> = 0 or <i>info</i> = <i>n</i> +1, the array <i>x</i> contains the solution matrix <i>X</i> to the system of equations. The second dimension of <i>x</i> must be at least $\max(1, nrhs)$.
<i>af</i> , <i>ipiv</i>	These arrays are output arguments if <i>fact</i> = 'N' . See the description of <i>af</i> , <i>ipiv</i> in <i>Input Arguments</i> section.
<i>rcond</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. An estimate of the reciprocal condition number of the matrix <i>A</i> . If <i>rcond</i> is less than the machine precision (in particular, if <i>rcond</i> = 0), the matrix is singular to working precision. This condition is indicated by a return code of <i>info</i> > 0.
<i>ferr</i> , <i>berr</i>	REAL for single precision flavors. DOUBLE PRECISION for double precision flavors. Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and relative backward errors, respectively, for each solution vector.

work(1) If *info*=0, on exit *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

info INTEGER. If *info*=0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i*, and $i \leq n$, then d_{ii} is exactly zero. The factorization has been completed, but the block diagonal matrix *D* is exactly singular, so the solution and error bounds could not be computed; *rcond* = 0 is returned.
If *info* = *i*, and $i = n + 1$, then *D* is nonsingular, but *rcond* is less than machine precision, meaning that the matrix is singular to working precision. Nevertheless, the solution and error bounds are computed because there are a number of situations where the computed solution can be more accurate than the value of *rcond* would suggest.

Application Notes

For real flavors, *lwork* must be at least $3*n$, and for complex flavors at least $2*n$. For better performance, try using $lwork = n*blocksize$, where *blocksize* is the optimal block size for *?sytrf*.

If you are in doubt how much workspace to supply, use *lwork* = -1 for the first run. In this case, a workspace query is assumed; the routine only calculates the optimal size of the *work* array, returns this value as the first entry *work(1)* of the *work* array, and no error message related to *lwork* is issued by XERBLA. On exit, examine *work(1)* and use this value for subsequent runs.

?hesvx

Uses the diagonal pivoting factorization to compute the solution to the complex system of linear equations with a Hermitian matrix A, and provides error bounds on the solution.

```
call chesvx (fact, uplo, n, nrhs, a, lda, af, ldaf, ipiv, b, ldb,  
            x, ldx, rcond, ferr, berr, work, lwork, rwork, info)  
call zhesvx (fact, uplo, n, nrhs, a, lda, af, ldaf, ipiv, b, ldb,  
            x, ldx, rcond, ferr, berr, work, lwork, rwork, info)
```

Discussion

This routine uses the diagonal pivoting factorization to compute the solution to a complex system of linear equations $AX = B$, where A is a n -by- n Hermitian matrix, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

Error bounds on the solution and a condition estimate are also provided.

The routine ?hesvx performs the following steps:

1. If $fact = 'N'$, the diagonal pivoting method is used to factor the matrix A. The form of the factorization is $A = UD U^H$ or $A = LD L^H$, where U (or L) is a product of permutation and unit upper (lower) triangular matrices, and D is Hermitian and block diagonal with 1-by-1 and 2-by-2 diagonal blocks.
2. If some $d_{i,i} = 0$, so that D is exactly singular, then the routine returns with $info = i$. Otherwise, the factored form of A is used to estimate the condition number of the matrix A. If the reciprocal of the condition number is less than machine precision, $info = n + 1$ is returned as a warning, but the routine still goes on to solve for X and compute error bounds as described below.
3. The system of equations is solved for X using the factored form of A.
4. Iterative refinement is applied to improve the computed solution matrix and calculate error bounds and backward error estimates for it.

Input Parameters

<i>fact</i>	<p>CHARACTER*1. Must be 'F' or 'N'.</p> <p>Specifies whether or not the factored form of the matrix A has been supplied on entry.</p> <p>If <i>fact</i> = 'F': on entry, <i>a</i>f and <i>ipiv</i> contain the factored form of A. Arrays <i>a</i>, <i>a</i>f, and <i>ipiv</i> will not be modified.</p> <p>If <i>fact</i> = 'N', the matrix A will be copied to <i>a</i>f and factored.</p>
<i>uplo</i>	<p>CHARACTER*1. Must be 'U' or 'L'.</p> <p>Indicates whether the upper or lower triangular part of A is stored and how A is factored:</p> <p>If <i>uplo</i> = 'U', the array <i>a</i> stores the upper triangular part of the Hermitian matrix A, and A is factored as UDU^H.</p> <p>If <i>uplo</i> = 'L', the array <i>a</i> stores the lower triangular part of the Hermitian matrix A; A is factored as LDL^H.</p>
<i>n</i>	<p>INTEGER. The order of matrix A ($n \geq 0$).</p>
<i>nrhs</i>	<p>INTEGER. The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).</p>
<i>a,af,b,work</i>	<p>COMPLEX for <i>chesvx</i> DOUBLE COMPLEX for <i>zhesvx</i>.</p> <p>Arrays: <i>a</i>(<i>lda</i>, *), <i>a</i>f(<i>ldaf</i>, *), <i>b</i>(<i>ldb</i>, *), <i>work</i>(*).</p> <p>The array <i>a</i> contains either the upper or the lower triangular part of the Hermitian matrix A (see <i>uplo</i>). The second dimension of <i>a</i> must be at least $\max(1,n)$.</p> <p>The array <i>a</i>f is an input argument if <i>fact</i> = 'F'. It contains the block diagonal matrix D and the multipliers used to obtain the factor U or L from the factorization $A = UD U^H$ or $A = LDL^H$ as computed by ?hetrf. The second dimension of <i>a</i>f must be at least $\max(1,n)$.</p>

The array *b* contains the matrix *B* whose columns are the right-hand sides for the systems of equations. The second dimension of *b* must be at least $\max(1, nrhs)$.

work()* is a workspace array of dimension (*lwork*).

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, n)$.

ldaf INTEGER. The first dimension of *af*; $ldaf \geq \max(1, n)$.

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

ipiv INTEGER.
Array, DIMENSION at least $\max(1, n)$.
The array *ipiv* is an input argument if *fact* = 'F'.
It contains details of the interchanges and the block structure of *D*, as determined by ?hetrf.
If $ipiv(i) = k > 0$, then d_{ii} is a 1-by-1 diagonal block, and the *i*th row and column of *A* was interchanged with the *k*th row and column.
If $uplo = 'U'$ and $ipiv(i) = ipiv(i-1) = -m < 0$, then *D* has a 2-by-2 block in rows/columns *i* and *i-1*, and (*i-1*)th row and column of *A* was interchanged with the *m*th row and column.
If $uplo = 'L'$ and $ipiv(i) = ipiv(i+1) = -m < 0$, then *D* has a 2-by-2 block in rows/columns *i* and *i+1*, and (*i+1*)th row and column of *A* was interchanged with the *m*th row and column.

ldx INTEGER. The leading dimension of the output array *x*; $ldx \geq \max(1, n)$.

lwork INTEGER. The size of the *work* array .
See *Application notes* for the suggested value of *lwork*.

rwork REAL for chesvx;
DOUBLE PRECISION for zhesvx.
Workspace array, DIMENSION at least $\max(1, n)$.

Output Parameters

x COMPLEX for chesvx
DOUBLE COMPLEX for zhesvx.
Array, DIMENSION (*ldx*, *).

	If $info = 0$ or $info = n+1$, the array x contains the solution matrix X to the system of equations. The second dimension of x must be at least $\max(1, nrhs)$.
$af, ipiv$	These arrays are output arguments if $fact = 'N'$. See the description of $af, ipiv$ in <i>Input Arguments</i> section.
$rcond$	REAL for $chesvx$; DOUBLE PRECISION for $zhesvx$. An estimate of the reciprocal condition number of the matrix A . If $rcond$ is less than the machine precision (in particular, if $rcond = 0$), the matrix is singular to working precision. This condition is indicated by a return code of $info > 0$.
$ferr, berr$	REAL for $chesvx$; DOUBLE PRECISION for $zhesvx$. Arrays, DIMENSION at least $\max(1, nrhs)$. Contain the component-wise forward and relative backward errors, respectively, for each solution vector.
$work(1)$	If $info=0$, on exit $work(1)$ contains the minimum value of $lwork$ required for optimum performance. Use this $lwork$ for subsequent runs.
$info$	INTEGER. If $info=0$, the execution is successful. If $info = -i$, the i th parameter had an illegal value. If $info = i$, and $i \leq n$, then d_{ii} is exactly zero. The factorization has been completed, but the block diagonal matrix D is exactly singular, so the solution and error bounds could not be computed; $rcond = 0$ is returned. If $info = i$, and $i = n + 1$, then D is nonsingular, but $rcond$ is less than machine precision, meaning that the matrix is singular to working precision. Nevertheless, the solution and error bounds are computed because there are a number of situations where the computed solution can be more accurate than the value of $rcond$ would suggest.

Application Notes

The value of `lwork` must be at least $2*n$. For better performance, try using `lwork = n*blocksize`, where `blocksize` is the optimal block size for `?hetrf`.

If you are in doubt how much workspace to supply, use `lwork = -1` for the first run. In this case, a workspace query is assumed; the routine only calculates the optimal size of the `work` array, returns this value as the first entry `work(1)` of the `work` array, and no error message related to `lwork` is issued by XERBLA. On exit, examine `work(1)` and use this value for subsequent runs.

?hesv

Computes the solution to the system of linear equations with a Hermitian matrix A and multiple right-hand sides.

```
call chesv (uplo, n, nrhs, a, lda, ipiv, b, ldb, work, lwork, info)
call zhesv (uplo, n, nrhs, a, lda, ipiv, b, ldb, work, lwork, info)
```

Discussion

This routine solves for X the real or complex system of linear equations $AX = B$, where A is an n -by- n symmetric matrix, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

The diagonal pivoting method is used to factor A as $A = U D U^H$ or $A = L D L^H$, where U (or L) is a product of permutation and unit upper (lower) triangular matrices, and D is Hermitian and block diagonal with 1-by-1 and 2-by-2 diagonal blocks.

The factored form of A is then used to solve the system of equations $AX = B$.

Input Parameters

`uplo` CHARACTER*1. Must be 'U' or 'L'.

Indicates whether the upper or lower triangular part of A is stored and how A is factored:
 If `uplo = 'U'`, the array `a` stores the upper triangular part of the matrix A , and A is factored as UDU^H .
 If `uplo = 'L'`, the array `a` stores the lower triangular part of the matrix A ; A is factored as LDL^H .

`n` **INTEGER**. The order of matrix A ($n \geq 0$).

`nrhs` **INTEGER**. The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).

`a, b, work` **COMPLEX** for `chesv`
DOUBLE COMPLEX for `zhesv`.
 Arrays: `a(lda,*)`, `b ldb,*)`, `work(lwork)`.
 The array `a` contains either the upper or the lower triangular part of the Hermitian matrix A (see `uplo`).
 The second dimension of `a` must be at least $\max(1, n)$.
 The array `b` contains the matrix B whose columns are the right-hand sides for the systems of equations.
 The second dimension of `b` must be at least $\max(1, nrhs)$.
`work(lwork)` is a workspace array.

`lda` **INTEGER**. The first dimension of `a`; $lda \geq \max(1, n)$.

`ldb` **INTEGER**. The first dimension of `b`; $ldb \geq \max(1, n)$.

`lwork` **INTEGER**. The size of the `work` array ($lwork \geq 1$)
 See *Application notes* for the suggested value of `lwork`.

Output Parameters

`a` If `info = 0`, `a` is overwritten by the block-diagonal matrix D and the multipliers used to obtain the factor U (or L) from the factorization of A as computed by `?hetrf`.

`b` If `info = 0`, `b` is overwritten by the solution matrix X .

`ipiv` **INTEGER**.
 Array, **DIMENSION** at least $\max(1, n)$.
 Contains details of the interchanges and the block structure of D , as determined by `?hetrf`.

If $ipiv(i) = k > 0$, then d_{ii} is a 1-by-1 diagonal block, and the i th row and column of A was interchanged with the k th row and column.

If $uplo = 'U'$ and $ipiv(i) = ipiv(i-1) = -m < 0$, then D has a 2-by-2 block in rows/columns i and $i-1$, and $(i-1)$ th row and column of A was interchanged with the m th row and column.

If $uplo = 'L'$ and $ipiv(i) = ipiv(i+1) = -m < 0$, then D has a 2-by-2 block in rows/columns i and $i+1$, and $(i+1)$ th row and column of A was interchanged with the m th row and column.

work(1) If $info=0$, on exit **work(1)** contains the minimum value of **lwork** required for optimum performance. Use this **lwork** for subsequent runs.

info **INTEGER**. If $info=0$, the execution is successful. If $info = -i$, the i th parameter had an illegal value. If $info = i$, d_{ii} is 0. The factorization has been completed, but D is exactly singular, so the solution could not be computed.

Application Notes

For better performance, try using $lwork = n * blocksize$, where $blocksize$ is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*.

If you are in doubt how much workspace to supply, use $lwork = -1$ for the first run. In this case, a workspace query is assumed; the routine only calculates the optimal size of the **work** array, returns this value as the first entry **work(1)** of the **work** array, and no error message related to **lwork** is issued by XERBLA. On exit, examine **work(1)** and use this value for subsequent runs.

?spsv

Computes the solution to the system of linear equations with a real or complex symmetric matrix A stored in packed format, and multiple right-hand sides.

```
call sspsv (uplo, n, nrhs, ap, ipiv, b, ldb, info)
call dspsv (uplo, n, nrhs, ap, ipiv, b, ldb, info)
call cspsv (uplo, n, nrhs, ap, ipiv, b, ldb, info)
call zspsv (uplo, n, nrhs, ap, ipiv, b, ldb, info)
```

Discussion

This routine solves for X the real or complex system of linear equations $AX = B$, where A is an n -by- n symmetric matrix stored in packed format, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

The diagonal pivoting method is used to factor A as $A = U D U^T$ or $A = L D L^T$, where U (or L) is a product of permutation and unit upper (lower) triangular matrices, and D is symmetric and block diagonal with 1-by-1 and 2-by-2 diagonal blocks.

The factored form of A is then used to solve the system of equations $AX = B$.

Input Parameters

<code>uplo</code>	CHARACTER*1 . Must be 'U' or 'L'. Indicates whether the upper or lower triangular part of A is stored and how A is factored: If <code>uplo</code> = 'U', the array <code>ap</code> stores the upper triangular part of the matrix A , and A is factored as UDU^T . If <code>uplo</code> = 'L', the array <code>ap</code> stores the lower triangular part of the matrix A ; A is factored as LDL^T .
<code>n</code>	INTEGER . The order of matrix A ($n \geq 0$).
<code>nrhs</code>	INTEGER . The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).

ap, *b* REAL for *sspsv*
 DOUBLE PRECISION for *dspsv*
 COMPLEX for *cspsv*
 DOUBLE COMPLEX for *zspsv*.
 Arrays: *ap*(*), *b*(*ldb*,*)
 The dimension of *ap* must be at least $\max(1, n(n+1)/2)$.
 The array *ap* contains the factor *U* or *L*, as specified by *uplo*, in *packed storage* (see [Matrix Storage Schemes](#)).
 The array *b* contains the matrix *B* whose columns are the right-hand sides for the systems of equations.
 The second dimension of *b* must be at least $\max(1, nrhs)$.

ldb INTEGER. The first dimension of *b*; $ldb \geq \max(1, n)$.

Output Parameters

ap The block-diagonal matrix *D* and the multipliers used to obtain the factor *U* (or *L*) from the factorization of *A* as computed by [?sptf](#), stored as a packed triangular matrix in the same storage format as *A*.

b If *info* = 0, *b* is overwritten by the solution matrix *X*.

ipiv INTEGER.
 Array, DIMENSION at least $\max(1, n)$.
 Contains details of the interchanges and the block structure of *D*, as determined by [?sptf](#).
 If *ipiv*(*i*) = *k* > 0, then *d_{ii}* is a 1-by-1 block, and the *i*th row and column of *A* was interchanged with the *k*th row and column.
 If *uplo* = 'U' and *ipiv*(*i*) = *ipiv*(*i*-1) = -*m* < 0, then *D* has a 2-by-2 block in rows/columns *i* and *i*-1, and (*i*-1)th row and column of *A* was interchanged with the *m*th row and column.
 If *uplo* = 'L' and *ipiv*(*i*) = *ipiv*(*i*+1) = -*m* < 0, then *D* has a 2-by-2 block in rows/columns *i* and *i*+1, and (*i*+1)th row and column of *A* was interchanged with the *m*th row and column.

info INTEGER. If *info*=0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i*, d_{ii} is 0. The factorization has been completed, but D is exactly singular, so the solution could not be computed.

?spsvx

Uses the diagonal pivoting factorization to compute the solution to the system of linear equations with a real or complex symmetric matrix A stored in packed format, and provides error bounds on the solution.

```
call sspsvx (fact, uplo, n, nrhs, ap, AFP, ipiv, b, ldb, x, ldx,
             rcond, ferr, berr, work, iwork, info)
call dspsvx (fact, uplo, n, nrhs, ap, AFP, ipiv, b, ldb, x, ldx,
             rcond, ferr, berr, work, iwork, info)
call cspsvx (fact, uplo, n, nrhs, ap, AFP, ipiv, b, ldb, x, ldx,
             rcond, ferr, berr, work, rwork, info)
call zspsvx (fact, uplo, n, nrhs, ap, AFP, ipiv, b, ldb, x, ldx,
             rcond, ferr, berr, work, rwork, info)
```

Discussion

This routine uses the diagonal pivoting factorization to compute the solution to a real or complex system of linear equations $AX = B$, where A is a n -by- n symmetric matrix stored in packed format, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

Error bounds on the solution and a condition estimate are also provided.

The routine ?spsvx performs the following steps:

1. If $fact = 'N'$, the diagonal pivoting method is used to factor the matrix A . The form of the factorization is $A = UD U^T$ or $A = LDL^T$, where U (or L) is a product of permutation and unit upper (lower) triangular matrices, and D is symmetric and block diagonal with 1-by-1 and 2-by-2 diagonal blocks.
2. If some $d_{i,i} = 0$, so that D is exactly singular, then the routine returns with $info = i$. Otherwise, the factored form of A is used to estimate the condition number of the matrix A . If the reciprocal of the condition number

is less than machine precision, $info = n + 1$ is returned as a warning, but the routine still goes on to solve for X and compute error bounds as described below.

3. The system of equations is solved for X using the factored form of A .
4. Iterative refinement is applied to improve the computed solution matrix and calculate error bounds and backward error estimates for it.

Input Parameters

<i>fact</i>	<p>CHARACTER*1. Must be 'F' or 'N'.</p> <p>Specifies whether or not the factored form of the matrix A has been supplied on entry.</p> <p>If <i>fact</i> = 'F': on entry, <i>afp</i> and <i>ipiv</i> contain the factored form of A. Arrays <i>ap</i>, <i>afp</i>, and <i>ipiv</i> will not be modified.</p> <p>If <i>fact</i> = 'N', the matrix A will be copied to <i>afp</i> and factored.</p>
<i>uplo</i>	<p>CHARACTER*1. Must be 'U' or 'L'.</p> <p>Indicates whether the upper or lower triangular part of A is stored and how A is factored:</p> <p>If <i>uplo</i> = 'U', the array <i>ap</i> stores the upper triangular part of the symmetric matrix A, and A is factored as UDU^T.</p> <p>If <i>uplo</i> = 'L', the array <i>ap</i> stores the lower triangular part of the symmetric matrix A; A is factored as LDL^T.</p>
<i>n</i>	INTEGER. The order of matrix A ($n \geq 0$).
<i>nrhs</i>	INTEGER. The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).
<i>ap,afp,b,work</i>	<p>REAL for <i>sspsvx</i></p> <p>DOUBLE PRECISION for <i>dspsvx</i></p> <p>COMPLEX for <i>cspsvx</i></p> <p>DOUBLE COMPLEX for <i>zpsvx</i>.</p> <p>Arrays: <i>ap</i>(*), <i>afp</i>(*), <i>b</i>(<i>ldb</i>,*), <i>work</i>(*).</p>

The array *ap* contains the upper or lower triangle of the symmetric matrix *A* in *packed storage* (see [Matrix Storage Schemes](#)).

The array *afp* is an input argument if *fact* = 'F'. It contains the block diagonal matrix *D* and the multipliers used to obtain the factor *U* or *L* from the factorization $A = U D U^T$ or $A = L D L^T$ as computed by [?sptf](#), in the same storage format as *A*.

The array *b* contains the matrix *B* whose columns are the right-hand sides for the systems of equations.

work ()* is a workspace array.

The dimension of arrays *ap* and *afp* must be at least $\max(1, n(n+1)/2)$; the second dimension of *b* must be at least $\max(1, nrhs)$; the dimension of *work* must be at least $\max(1, 3*n)$ for real flavors and $\max(1, 2*n)$ for complex flavors.

ldb **INTEGER.** The first dimension of *b*; $ldb \geq \max(1, n)$.

ipiv **INTEGER.**

Array, **DIMENSION** at least $\max(1, n)$.

The array *ipiv* is an input argument if *fact* = 'F'. It contains details of the interchanges and the block structure of *D*, as determined by [?sptf](#).

If $ipiv(i) = k > 0$, then d_{ii} is a 1-by-1 diagonal block, and the *i*th row and column of *A* was interchanged with the *k*th row and column.

If $uplo = 'U'$ and $ipiv(i) = ipiv(i-1) = -m < 0$, then *D* has a 2-by-2 block in rows/columns *i* and *i-1*, and (*i-1*)th row and column of *A* was interchanged with the *m*th row and column.

If $uplo = 'L'$ and $ipiv(i) = ipiv(i+1) = -m < 0$, then *D* has a 2-by-2 block in rows/columns *i* and *i+1*, and (*i+1*)th row and column of *A* was interchanged with the *m*th row and column.

ldx **INTEGER.** The leading dimension of the output array *x*; $ldx \geq \max(1, n)$.

iwork **INTEGER**.
Workspace array, **DIMENSION** at least $\max(1, n)$; used in real flavors only.

rwork **REAL** for **cspsvx**;
DOUBLE PRECISION for **zspsvx**.
Workspace array, **DIMENSION** at least $\max(1, n)$; used in complex flavors only.

Output Parameters

x **REAL** for **sspsvx**
DOUBLE PRECISION for **dspsvx**
COMPLEX for **cspsvx**
DOUBLE COMPLEX for **zspsvx**.
Array, **DIMENSION** (*ldx*, *).
If *info* = 0 or *info* = *n*+1, the array *x* contains the solution matrix *X* to the system of equations. The second dimension of *x* must be at least $\max(1, nrhs)$.

afp, *ipiv* These arrays are output arguments if *fact* = 'N'. See the description of *afp*, *ipiv* in *Input Arguments* section.

rcond **REAL** for single precision flavors.
DOUBLE PRECISION for double precision flavors.
An estimate of the reciprocal condition number of the matrix *A*. If *rcond* is less than the machine precision (in particular, if *rcond* = 0), the matrix is singular to working precision. This condition is indicated by a return code of *info* > 0.

ferr, *berr* **REAL** for single precision flavors.
DOUBLE PRECISION for double precision flavors.
Arrays, **DIMENSION** at least $\max(1, nrhs)$. Contain the component-wise forward and relative backward errors, respectively, for each solution vector.

info **INTEGER**. If *info*=0, the execution is successful. If *info* = -*i*, the *i*th parameter had an illegal value. If *info* = *i*, and *i* ≤ *n*, then d_{ii} is exactly zero. The factorization has been completed, but the block diagonal

matrix D is exactly singular, so the solution and error bounds could not be computed; $rcond = 0$ is returned. If $info = i$, and $i = n + 1$, then D is nonsingular, but $rcond$ is less than machine precision, meaning that the matrix is singular to working precision. Nevertheless, the solution and error bounds are computed because there are a number of situations where the computed solution can be more accurate than the value of $rcond$ would suggest.

?hpsvx

Uses the diagonal pivoting factorization to compute the solution to the system of linear equations with a Hermitian matrix A stored in packed format, and provides error bounds on the solution.

```
call chpsvx (fact, uplo, n, nrhs, ap, afp, ipiv, b, ldb, x, ldx,  
            rcond, ferr, berr, work, rwork, info)  
call zhpsvx (fact, uplo, n, nrhs, ap, afp, ipiv, b, ldb, x, ldx,  
            rcond, ferr, berr, work, rwork, info)
```

Discussion

This routine uses the diagonal pivoting factorization to compute the solution to a complex system of linear equations $AX = B$, where A is a n -by- n Hermitian matrix stored in packed format, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

Error bounds on the solution and a condition estimate are also provided.

The routine ?hpsvx performs the following steps:

1. If `fact = 'N'`, the diagonal pivoting method is used to factor the matrix A . The form of the factorization is $A = UD U^H$ or $A = LD L^H$, where U (or L) is a product of permutation and unit upper (lower) triangular matrices, and D is Hermitian and block diagonal with 1-by-1 and 2-by-2 diagonal blocks.
2. If some $d_{i,i} = 0$, so that D is exactly singular, then the routine returns with `info = i`. Otherwise, the factored form of A is used to estimate the condition number of the matrix A . If the reciprocal of the condition number is less than machine precision, `info = n + 1` is returned as a warning, but the routine still goes on to solve for X and compute error bounds as described below.
3. The system of equations is solved for X using the factored form of A .
4. Iterative refinement is applied to improve the computed solution matrix and calculate error bounds and backward error estimates for it.

Input Parameters

<i>fact</i>	<p>CHARACTER*1. Must be 'F' or 'N'.</p> <p>Specifies whether or not the factored form of the matrix <i>A</i> has been supplied on entry.</p> <p>If <i>fact</i> = 'F': on entry, <i>afp</i> and <i>ipiv</i> contain the factored form of <i>A</i>. Arrays <i>ap</i>, <i>afp</i>, and <i>ipiv</i> will not be modified.</p> <p>If <i>fact</i> = 'N', the matrix <i>A</i> will be copied to <i>afp</i> and factored.</p>
<i>uplo</i>	<p>CHARACTER*1. Must be 'U' or 'L'.</p> <p>Indicates whether the upper or lower triangular part of <i>A</i> is stored and how <i>A</i> is factored:</p> <p>If <i>uplo</i> = 'U', the array <i>ap</i> stores the upper triangular part of the Hermitian matrix <i>A</i>, and <i>A</i> is factored as UDU^H.</p> <p>If <i>uplo</i> = 'L', the array <i>ap</i> stores the lower triangular part of the Hermitian matrix <i>A</i>; <i>A</i> is factored as LDL^H.</p>
<i>n</i>	<p>INTEGER. The order of matrix <i>A</i> ($n \geq 0$).</p>
<i>nrhs</i>	<p>INTEGER. The number of right-hand sides; the number of columns in <i>B</i> ($nrhs \geq 0$).</p>
<i>ap,afp,b,work</i>	<p>COMPLEX for <i>chpsvx</i> DOUBLE COMPLEX for <i>zhpsvx</i>.</p> <p>Arrays: <i>ap</i>(*), <i>afp</i>(*), <i>b</i>(<i>ldb</i>,*), <i>work</i>(*).</p> <p>The array <i>ap</i> contains the upper or lower triangle of the Hermitian matrix <i>A</i> in <i>packed storage</i> (see Matrix Storage Schemes).</p> <p>The array <i>afp</i> is an input argument if <i>fact</i> = 'F'. It contains the block diagonal matrix <i>D</i> and the multipliers used to obtain the factor <i>U</i> or <i>L</i> from the factorization $A = UD U^H$ or $A = LDL^H$ as computed by ?hptrf, in the same storage format as <i>A</i>.</p> <p>The array <i>b</i> contains the matrix <i>B</i> whose columns are the right-hand sides for the systems of equations.</p> <p><i>work</i>(*) is a workspace array.</p>

The dimension of arrays *ap* and *afp* must be at least $\max(1, n(n+1)/2)$; the second dimension of *b* must be at least $\max(1, nrhs)$; the dimension of *work* must be at least $\max(1, 2 * n)$.

ldb **INTEGER**. The first dimension of *b*; $ldb \geq \max(1, n)$.

ipiv **INTEGER**.
 Array, **DIMENSION** at least $\max(1, n)$.
 The array *ipiv* is an input argument if *fact* = 'F'.
 It contains details of the interchanges and the block structure of *D*, as determined by [?hptrf](#).
 If $ipiv(i) = k > 0$, then d_{ii} is a 1-by-1 diagonal block, and the *i*th row and column of *A* was interchanged with the *k*th row and column.
 If $uplo = 'U'$ and $ipiv(i) = ipiv(i-1) = -m < 0$, then *D* has a 2-by-2 block in rows/columns *i* and *i-1*, and (*i-1*)th row and column of *A* was interchanged with the *m*th row and column.
 If $uplo = 'L'$ and $ipiv(i) = ipiv(i+1) = -m < 0$, then *D* has a 2-by-2 block in rows/columns *i* and *i+1*, and (*i+1*)th row and column of *A* was interchanged with the *m*th row and column.

ldx **INTEGER**. The leading dimension of the output array *x*; $ldx \geq \max(1, n)$.

rwork **REAL** for *chpsvx*;
DOUBLE PRECISION for *zhpsvx*.
 Workspace array, **DIMENSION** at least $\max(1, n)$.

Output Parameters

x **COMPLEX** for *chpsvx*
DOUBLE COMPLEX for *zhpsvx*.
 Array, **DIMENSION** (*ldx*, *).
 If *info* = 0 or *info* = *n*+1, the array *x* contains the solution matrix *X* to the system of equations. The second dimension of *x* must be at least $\max(1, nrhs)$.

<i>afp, ipiv</i>	These arrays are output arguments if <i>fact</i> = 'N'. See the description of <i>afp, ipiv</i> in <i>Input Arguments</i> section.
<i>rcond</i>	REAL for <i>chpsvx</i> ; DOUBLE PRECISION for <i>zhpsvx</i> . An estimate of the reciprocal condition number of the matrix <i>A</i> . If <i>rcond</i> is less than the machine precision (in particular, if <i>rcond</i> = 0), the matrix is singular to working precision. This condition is indicated by a return code of <i>info</i> > 0.
<i>ferr, berr</i>	REAL for <i>chpsvx</i> ; DOUBLE PRECISION for <i>zhpsvx</i> . Arrays, DIMENSION at least max(1, <i>nrhs</i>). Contain the component-wise forward and relative backward errors, respectively, for each solution vector.
<i>info</i>	INTEGER. If <i>info</i> =0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = <i>i</i> , and <i>i</i> ≤ <i>n</i> , then <i>d_{ii}</i> is exactly zero. The factorization has been completed, but the block diagonal matrix <i>D</i> is exactly singular, so the solution and error bounds could not be computed; <i>rcond</i> = 0 is returned. If <i>info</i> = <i>i</i> , and <i>i</i> = <i>n</i> + 1, then <i>D</i> is nonsingular, but <i>rcond</i> is less than machine precision, meaning that the matrix is singular to working precision. Nevertheless, the solution and error bounds are computed because there are a number of situations where the computed solution can be more accurate than the value of <i>rcond</i> would suggest.

?hpsv

Computes the solution to the system of linear equations with a Hermitian matrix A stored in packed format, and multiple right-hand sides.

```
call chpsv (uplo, n, nrhs, ap, ipiv, b, ldb, info)
call zhpsv (uplo, n, nrhs, ap, ipiv, b, ldb, info)
```

Discussion

This routine solves for X the system of linear equations $AX = B$, where A is an n -by- n Hermitian matrix stored in packed format, the columns of matrix B are individual right-hand sides, and the columns of X are the corresponding solutions.

The diagonal pivoting method is used to factor A as $A = U D U^H$ or $A = L D L^H$, where U (or L) is a product of permutation and unit upper (lower) triangular matrices, and D is Hermitian and block diagonal with 1-by-1 and 2-by-2 diagonal blocks.

The factored form of A is then used to solve the system of equations $AX = B$.

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Indicates whether the upper or lower triangular part of A is stored and how A is factored:
If **uplo** = 'U', the array **ap** stores the upper triangular part of the matrix A , and A is factored as UDU^H .
If **uplo** = 'L', the array **ap** stores the lower triangular part of the matrix A ; A is factored as LDL^H .

n INTEGER. The order of matrix A ($n \geq 0$).

nrhs INTEGER. The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).

ap, *b* **COMPLEX** for `chpsv`
DOUBLE COMPLEX for `zhpsv`.
 Arrays: *ap*(*), *b*(*ldb*,*)
 The dimension of *ap* must be at least $\max(1, n(n+1)/2)$.
 The array *ap* contains the factor *U* or *L*, as specified by *uplo*, in *packed storage* (see [Matrix Storage Schemes](#)).
 The array *b* contains the matrix *B* whose columns are the right-hand sides for the systems of equations.
 The second dimension of *b* must be at least $\max(1, nrhs)$.

ldb **INTEGER**. The first dimension of *b*; $ldb \geq \max(1, n)$.

Output Parameters

ap The block-diagonal matrix *D* and the multipliers used to obtain the factor *U* (or *L*) from the factorization of *A* as computed by [?hptrf](#), stored as a packed triangular matrix in the same storage format as *A*.

b If *info* = 0, *b* is overwritten by the solution matrix *X*.

ipiv **INTEGER**.
 Array, **DIMENSION** at least $\max(1, n)$.
 Contains details of the interchanges and the block structure of *D*, as determined by [?hptrf](#).
 If *ipiv*(*i*) = *k* > 0, then *d_{ii}* is a 1-by-1 block, and the *i*th row and column of *A* was interchanged with the *k*th row and column.
 If *uplo* = 'U' and *ipiv*(*i*) = *ipiv*(*i*-1) = -*m* < 0, then *D* has a 2-by-2 block in rows/columns *i* and *i*-1, and (*i*-1)th row and column of *A* was interchanged with the *m*th row and column.
 If *uplo* = 'L' and *ipiv*(*i*) = *ipiv*(*i*+1) = -*m* < 0, then *D* has a 2-by-2 block in rows/columns *i* and *i*+1, and (*i*+1)th row and column of *A* was interchanged with the *m*th row and column.

info **INTEGER**. If *info*=0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i*, d_{ii} is 0. The factorization has been completed, but D is exactly singular, so the solution could not be computed.

LAPACK Routines: Least Squares and Eigenvalue Problems

5

This chapter describes the Intel[®] Math Kernel Library implementation of routines from the LAPACK package that are used for solving linear least-squares problems, eigenvalue and singular value problems, as well as performing a number of related computational tasks.

Sections in this chapter include descriptions of LAPACK [computational routines](#) and [driver routines](#).

For full reference on LAPACK routines and related information see [\[LUG\]](#).

Least-Squares Problems. A typical *least-squares problem* is as follows: given a matrix A and a vector b , find the vector x that minimizes the sum of squares $\sum_i ((Ax)_i - b_i)^2$ or, equivalently, find the vector x that minimizes the 2-norm $\|Ax - b\|_2$.

In the most usual case, A is an m by n matrix with $m \geq n$ and $\text{rank}(A) = n$. This problem is also referred to as finding the *least-squares solution* to an *overdetermined* system of linear equations (here we have more equations than unknowns). To solve this problem, you can use the *QR* factorization of the matrix A (see *QR Factorization* on [page 5-6](#)).

If $m < n$ and $\text{rank}(A) = m$, there exist an infinite number of solutions x which exactly satisfy $Ax = b$, and thus minimize the norm $\|Ax - b\|_2$. In this case it is often useful to find the unique solution that minimizes $\|x\|_2$. This problem is referred to as finding the *minimum-norm solution* to an *underdetermined* system of linear equations (here we have more unknowns than equations).

To solve this problem, you can use the *LQ* factorization of the matrix A (see *LQ Factorization* on [page 5-7](#)).

In the general case you may have a *rank-deficient least-squares problem*, with $\text{rank}(A) < \min(m, n)$: find the *minimum-norm least-squares solution* that minimizes both $\|x\|_2$ and $\|Ax - b\|_2$. In this case (or when the rank of A is in doubt) you can use the *QR factorization with pivoting* or *singular value decomposition* (see [page 5-74](#)).

Eigenvalue Problems (from German *eigen* “own”) are stated as follows: given a matrix A , find the *eigenvalues* λ and the corresponding *eigenvectors* z that satisfy the equation

$$Az = \lambda z \text{ (right eigenvectors } z)$$

or the equation

$$z^H A = \lambda z^H \text{ (left eigenvectors } z).$$

If A is a real symmetric or complex Hermitian matrix, the above two equations are equivalent, and the problem is called a *symmetric eigenvalue problem*. Routines for solving this type of problems are described in the section *Symmetric Eigenvalue Problems* (see [page 5-101](#)).

Routines for solving eigenvalue problems with nonsymmetric or non-Hermitian matrices are described in the section *Nonsymmetric Eigenvalue Problems* (see [page 5-174](#)).

The library also includes routines that handle *generalized symmetric-definite eigenvalue problems*: find the eigenvalues λ and the corresponding eigenvectors x that satisfy one of the following equations:

$$Az = \lambda Bz, \quad ABz = \lambda z, \quad \text{or} \quad BAz = \lambda z$$

where A is symmetric or Hermitian, and B is symmetric positive-definite or Hermitian positive-definite. Routines for reducing these problems to standard symmetric eigenvalue problems are described in the section *Generalized Symmetric-Definite Eigenvalue Problems* (see [page 5-157](#)).

* * *

To solve a particular problem, you usually call several computational routines. Sometimes you need to combine the routines of this chapter with other LAPACK routines described in Chapter 4 as well as with BLAS routines (Chapter 2).

For example, to solve a set of least-squares problems minimizing $\|Ax - b\|_2$ for all columns b of a given matrix B (where A and B are real matrices), you can call `?gexrf` to form the factorization $A = QR$, then call `?ormqr` to compute $C = Q^H B$, and finally call the BLAS routine `?trsm` to solve for X the system of equations $RX = C$.

Another way is to call an appropriate driver routine that performs several tasks in one call. For example, to solve the least-squares problem the driver routine `?gels` can be used.

Routine Naming Conventions

For each routine in this chapter, you can use the LAPACK name.

LAPACK names have the structure `xyyzzz`, which is described below.

The initial letter `x` indicates the data type:

<code>s</code>	real, single precision	<code>c</code>	complex, single precision
<code>d</code>	real, double precision	<code>z</code>	complex, double precision

The second and third letters `yy` indicate the matrix type and storage scheme:

<code>bd</code>	bidiagonal matrix
<code>ge</code>	general matrix
<code>gb</code>	general band matrix
<code>hs</code>	upper Hessenberg matrix
<code>or</code>	(real) orthogonal matrix
<code>op</code>	(real) orthogonal matrix (packed storage)
<code>un</code>	(complex) unitary matrix
<code>up</code>	(complex) unitary matrix (packed storage)
<code>pt</code>	symmetric or Hermitian positive-definite tridiagonal matrix
<code>sy</code>	symmetric matrix
<code>sp</code>	symmetric matrix (packed storage)
<code>sb</code>	(real) symmetric band matrix
<code>st</code>	(real) symmetric tridiagonal matrix
<code>he</code>	Hermitian matrix
<code>hp</code>	Hermitian matrix (packed storage)
<code>hb</code>	(complex) Hermitian band matrix
<code>tr</code>	triangular or quasi-triangular matrix.

The last three letters `zzz` indicate the computation performed, for example:

<code>qrf</code>	form the QR factorization
<code>lqf</code>	form the LQ factorization.

Thus, the routine `sgeqrf` forms the QR factorization of general real matrices in single precision; the corresponding routine for complex matrices is `cgeqrf`.

Matrix Storage Schemes

LAPACK routines use the following matrix storage schemes:

- *Full storage*: a matrix A is stored in a two-dimensional array a , with the matrix element a_{ij} stored in the array element $a(i, j)$.
- *Packed storage* scheme allows you to store symmetric, Hermitian, or triangular matrices more compactly: the upper or lower triangle of the matrix is packed by columns in a one-dimensional array.
- *Band storage*: an m by n band matrix with kl sub-diagonals and ku super-diagonals is stored compactly in a two-dimensional array ab with $kl+ku+1$ rows and n columns. Columns of the matrix are stored in the corresponding columns of the array, and *diagonals* of the matrix are stored in rows of the array.

In Chapters 4 and 5, arrays that hold matrices in packed storage have names ending in p ; arrays with matrices in band storage have names ending in b .

For more information on matrix storage schemes, see [Matrix Arguments](#) in Appendix A.

Mathematical Notation

In addition to the mathematical notation used in previous chapters, descriptions of routines in this chapter use the following notation:

λ_i	<i>Eigenvalues</i> of the matrix A (for the definition of eigenvalues, see Eigenvalue Problems on page 5-2).
σ_i	<i>Singular values</i> of the matrix A . They are equal to square roots of the eigenvalues of $A^H A$. (For more information, see Singular Value Decomposition).
$\ x\ _2$	The <i>2-norm</i> of the vector x : $\ x\ _2 = (\sum_i x_i ^2)^{1/2} = \ x\ _E$.
$\ A\ _2$	The <i>2-norm</i> (or <i>spectral norm</i>) of the matrix A . $\ A\ _2 = \max_i \sigma_i$, $\ A\ _2^2 = \max_{\ x\ =1} (Ax \cdot Ax)$.
$\ A\ _E$	The <i>Euclidean norm</i> of the matrix A : $\ A\ _E^2 = \sum_i \sum_j a_{ij} ^2$ (for vectors, the Euclidean norm and the 2-norm are equal: $\ x\ _E = \ x\ _2$).
$\theta(x, y)$	The <i>acute angle between vectors</i> x and y : $\cos \theta(x, y) = x \cdot y / (\ x\ _2 \ y\ _2)$.

Computational Routines

In the sections that follow, the descriptions of LAPACK computational routines are given. These routines perform distinct computational tasks that can be used for:

[Orthogonal Factorizations](#)
[Singular Value Decomposition](#)
[Symmetric Eigenvalue Problems](#)
[Generalized Symmetric-Definite Eigenvalue Problems](#)
[Nonsymmetric Eigenvalue Problems](#)
[Generalized Nonsymmetric Eigenvalue Problems](#)
[Generalized Singular Value Decomposition](#)

See also the respective [driver routines](#).

Orthogonal Factorizations

This section describes the LAPACK routines for the QR (RQ) and LQ (QL) factorization of matrices. Routines for the RZ factorization as well as for generalized QR and RQ factorizations are also included.

QR Factorization. Assume that A is an m by n matrix to be factored.

If $m \geq n$, the QR factorization is given by

$$A = Q \begin{pmatrix} R \\ 0 \end{pmatrix} = (Q_1, Q_2) \begin{pmatrix} R \\ 0 \end{pmatrix}$$

where R is an n by n upper triangular matrix with real diagonal elements, and Q is an m by m orthogonal (or unitary) matrix.

You can use the QR factorization for solving the following least-squares problem: minimize $\|Ax - b\|_2$ where A is a full-rank m by n matrix ($m \geq n$). After factoring the matrix, compute the solution x by solving $Rx = (Q_1)^T b$.

If $m < n$, the QR factorization is given by

$$A = QR = Q(R_1 R_2)$$

where R is trapezoidal, R_1 is upper triangular and R_2 is rectangular.

The LAPACK routines do not form the matrix Q explicitly. Instead, Q is represented as a product of $\min(m, n)$ *elementary reflectors*. Routines are provided to work with Q in this representation.

LQ Factorization of an m by n matrix A is as follows. If $m \leq n$,

$$A = (L, 0)Q = (L, 0) \begin{pmatrix} Q_1 \\ Q_2 \end{pmatrix} = LQ_1$$

where L is an m by m lower triangular matrix with real diagonal elements, and Q is an n by n orthogonal (or unitary) matrix.

If $m > n$, the LQ factorization is

$$A = \begin{pmatrix} L_1 \\ L_2 \end{pmatrix} Q$$

where L_1 is an n by n lower triangular matrix, L_2 is rectangular, and Q is an n by n orthogonal (or unitary) matrix.

You can use the LQ factorization to find the minimum-norm solution of an underdetermined system of linear equations $Ax = b$ where A is an m by n matrix of rank m ($m < n$). After factoring the matrix, compute the solution vector x as follows: solve $Ly = b$ for y , and then compute $x = (Q_1)^H y$.

Table 5-1 lists LAPACK routines that perform orthogonal factorization of matrices.

Table 5-1 Computational Routines for Orthogonal Factorization

Matrix type, factorization	Factorize without pivoting	Factorize with pivoting	Generate matrix Q	Apply matrix Q
general matrices, QR factorization	?geqrf	?geqpf ?geqp3	?orgqr ?ungqr	?ormqr ?unmqr
general matrices, RQ factorization	?gerqf		?orgrq ?ungrq	?ormrq ?unmrq
general matrices, LQ factorization	?gelqf		?orglq ?unglq	?ormlq ?unmlq
general matrices, QL factorization	?geqlf		?orgql ?ungql	?ormql ?unmql
trapezoidal matrices, RZ factorization	?tzzrf			?ormrz ?unmrz
pair of matrices, generalized QR factorization	?ggqrf			
pair of matrices, generalized RQ factorization	?ggrqf			

?geqrf

Computes the *QR* factorization of a general m by n matrix.

```
call sgeqrf ( m, n, a, lda, tau, work, lwork, info )
call dgeqrf ( m, n, a, lda, tau, work, lwork, info )
call cgeqrf ( m, n, a, lda, tau, work, lwork, info )
call zgeqrf ( m, n, a, lda, tau, work, lwork, info )
```

Discussion

The routine forms the *QR* factorization of a general m by n matrix A (see *Orthogonal Factorizations* on [page 5-6](#)). No pivoting is performed.

The routine does not form the matrix Q explicitly. Instead, Q is represented as a product of $\min(m, n)$ elementary reflectors. Routines are provided to work with Q in this representation.

Input Parameters

m **INTEGER**. The number of rows in the matrix A ($m \geq 0$).

n **INTEGER**. The number of columns in A ($n \geq 0$).

$a, work$ **REAL** for `sgeqrf`
DOUBLE PRECISION for `dgeqrf`
COMPLEX for `cgeqrf`
DOUBLE COMPLEX for `zgeqrf`.

Arrays:
 $a(lda, *)$ contains the matrix A .
The second dimension of a must be at least $\max(1, n)$.
 $work(lwork)$ is a workspace array.

lda **INTEGER**. The first dimension of a ; at least $\max(1, m)$.

$lwork$ **INTEGER**. The size of the $work$ array ($lwork \geq n$)
See [Application notes](#) for the suggested value of $lwork$.

Output Parameters

- a*** Overwritten by the factorization data as follows:
If $m \geq n$, the elements below the diagonal are overwritten by the details of the unitary matrix Q , and the upper triangle is overwritten by the corresponding elements of the upper triangular matrix R .
If $m < n$, the strictly lower triangular part is overwritten by the details of the unitary matrix Q , and the remaining elements are overwritten by the corresponding elements of the m by n upper trapezoidal matrix R .
- tau*** REAL for `sgeqrf`
DOUBLE PRECISION for `dgeqrf`
COMPLEX for `cgeqrf`
DOUBLE COMPLEX for `zgeqrf`.
Array, DIMENSION at least $\max(1, \min(m, n))$.
Contains additional information on the matrix Q .
- work(1)*** If $info = 0$, on exit `work(1)` contains the minimum value of `lwork` required for optimum performance. Use this `lwork` for subsequent runs.
- info*** INTEGER.
If $info = 0$, the execution is successful.
If $info = -i$, the i th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = n * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*.

If you are in doubt how much workspace to supply, use a generous value of `lwork` for the first run. On exit, examine `work(1)` and use this value for subsequent runs.

The computed factorization is the exact factorization of a matrix $A + E$, where $\|E\|_2 = O(\epsilon) \|A\|_2$.

The approximate number of floating-point operations for real flavors is

$$(4/3)n^3 \quad \text{if } m = n,$$

$$(2/3)n^2(3m-n) \quad \text{if } m > n,$$

$$(2/3)m^2(3n-m) \quad \text{if } m < n.$$

The number of operations for complex flavors is 4 times greater.

To solve a set of least-squares problems minimizing $\|Ax - b\|_2$ for all columns b of a given matrix B , you can call the following:

`?geqrf` (this routine) to factorize $A = QR$;

`?ormqr` to compute $C = Q^TB$ (for real matrices);

`?unmqr` to compute $C = Q^HB$ (for complex matrices);

`?trsm` (a BLAS routine) to solve $RX = C$.

(The columns of the computed X are the least-squares solution vectors x .)

To compute the elements of Q explicitly, call

`?orgqr` (for real matrices)

`?ungqr` (for complex matrices).

?geqpf

Computes the *QR* factorization of a general m by n matrix with pivoting.

```
call sgeqpf ( m, n, a, lda, jpvt, tau, work, info )
call dgeqpf ( m, n, a, lda, jpvt, tau, work, info )
call cgeqpf ( m, n, a, lda, jpvt, tau, work, rwork, info )
call zgeqpf ( m, n, a, lda, jpvt, tau, work, rwork, info )
```

Discussion

This routine is deprecated and has been replaced by routine [?geqp3](#).

The routine `?geqpf` forms the *QR* factorization of a general m by n matrix A with column pivoting: $AP = QR$ (see *Orthogonal Factorizations* on [page 5-6](#)). Here P denotes an n by n permutation matrix.

The routine does not form the matrix Q explicitly. Instead, Q is represented as a product of $\min(m, n)$ elementary reflectors. Routines are provided to work with Q in this representation.

Input Parameters

m **INTEGER**. The number of rows in the matrix A ($m \geq 0$).

n **INTEGER**. The number of columns in A ($n \geq 0$).

$a, work$ **REAL** for `sgeqpf`
DOUBLE PRECISION for `dgeqpf`
COMPLEX for `cgeqpf`
DOUBLE COMPLEX for `zgeqpf`.

Arrays:
 $a(lda, *)$ contains the matrix A .
The second dimension of a must be at least $\max(1, n)$.
 $work(lwork)$ is a workspace array.

lda **INTEGER**. The first dimension of a ; at least $\max(1, m)$.

$lwork$ **INTEGER**. The size of the $work$ array; must be at least $\max(1, 3*n)$.

<i>jpvt</i>	<p>INTEGER. Array, DIMENSION at least $\max(1, n)$.</p> <p>On entry, if $jpvt(i) > 0$, the ith column of A is moved to the beginning of AP before the computation, and fixed in place during the computation.</p> <p>If $jpvt(i) = 0$, the ith column of A is a free column (that is, it may be interchanged during the computation with any other free column).</p>
<i>rwork</i>	<p>REAL for cgeqpf</p> <p>DOUBLE PRECISION for zgeqpf.</p> <p>A workspace array, DIMENSION at least $\max(1, 2*n)$.</p>

Output Parameters

<i>a</i>	<p>Overwritten by the factorization data as follows:</p> <p>If $m \geq n$, the elements below the diagonal are overwritten by the details of the unitary (orthogonal) matrix Q, and the upper triangle is overwritten by the corresponding elements of the upper triangular matrix R.</p> <p>If $m < n$, the strictly lower triangular part is overwritten by the details of the matrix Q, and the remaining elements are overwritten by the corresponding elements of the m by n upper trapezoidal matrix R.</p>
<i>tau</i>	<p>REAL for sgeqpf</p> <p>DOUBLE PRECISION for dgeqpf</p> <p>COMPLEX for cgeqpf</p> <p>DOUBLE COMPLEX for zgeqpf.</p> <p>Array, DIMENSION at least $\max(1, \min(m, n))$.</p> <p>Contains additional information on the matrix Q.</p>
<i>jpvt</i>	<p>Overwritten by details of the permutation matrix P in the factorization $AP = QR$. More precisely, the columns of AP are the columns of A in the following order:</p> <p>$jpvt(1), jpvt(2), \dots, jpvt(n)$.</p>
<i>info</i>	<p>INTEGER.</p> <p>If $info = 0$, the execution is successful.</p> <p>If $info = -i$, the ith parameter had an illegal value.</p>

Application Notes

The computed factorization is the exact factorization of a matrix $A + E$, where $\|E\|_2 = O(\epsilon) \|A\|_2$.

The approximate number of floating-point operations for real flavors is

$$(4/3)n^3 \quad \text{if } m = n,$$

$$(2/3)n^2(3m-n) \quad \text{if } m > n,$$

$$(2/3)m^2(3n-m) \quad \text{if } m < n.$$

The number of operations for complex flavors is 4 times greater.

To solve a set of least-squares problems minimizing $\|Ax - b\|_2$ for all columns b of a given matrix B , you can call the following:

`?geqpf` (this routine) to factorize $AP = QR$;

`?ormqr` to compute $C = Q^T B$ (for real matrices);

`?unmqr` to compute $C = Q^H B$ (for complex matrices);

`?trsm` (a BLAS routine) to solve $RX = C$.

(The columns of the computed X are the permuted least-squares solution vectors x ; the output array `jpvt` specifies the permutation order.)

To compute the elements of Q explicitly, call

`?orgqr` (for real matrices)

`?ungqr` (for complex matrices).

?geqp3

Computes the *QR* factorization of a general *m* by *n* matrix with column pivoting using Level 3 BLAS.

```
call sgeqp3 ( m, n, a, lda, jpvt, tau, work, lwork, info )
call dgeqp3 ( m, n, a, lda, jpvt, tau, work, lwork, info )
call cgeqp3 ( m, n, a, lda, jpvt, tau, work, lwork, rwork, info )
call zgeqp3 ( m, n, a, lda, jpvt, tau, work, lwork, rwork, info )
```

Discussion

The routine forms the *QR* factorization of a general *m* by *n* matrix *A* with column pivoting: $AP = QR$ (see *Orthogonal Factorizations* on [page 5-6](#)) using Level 3 BLAS. Here *P* denotes an *n* by *n* permutation matrix.

Use this routine instead of `?geqpf` for better performance.

The routine does not form the matrix *Q* explicitly. Instead, *Q* is represented as a product of $\min(m, n)$ elementary reflectors. Routines are provided to work with *Q* in this representation.

Input Parameters

<i>m</i>	INTEGER. The number of rows in the matrix <i>A</i> ($m \geq 0$).
<i>n</i>	INTEGER. The number of columns in <i>A</i> ($n \geq 0$).
<i>a</i> , <i>work</i>	REAL for <code>sgeqp3</code> DOUBLE PRECISION for <code>dgeqp3</code> COMPLEX for <code>cgeqp3</code> DOUBLE COMPLEX for <code>zgeqp3</code> . Arrays: <i>a</i> (<i>lda</i> , *) contains the matrix <i>A</i> . The second dimension of <i>a</i> must be at least $\max(1, n)$. <i>work</i> (<i>lwork</i>) is a workspace array.
<i>lda</i>	INTEGER. The first dimension of <i>a</i> ; at least $\max(1, m)$.

- lwork* **INTEGER**. The size of the *work* array; must be at least $\max(1, 3*n+1)$ for real flavors, and at least $\max(1, n+1)$ for complex flavors.
- jpvt* **INTEGER**. Array, **DIMENSION** at least $\max(1, n)$.
On entry, if $jpvt(i) \neq 0$, the i th column of A is moved to the beginning of AP before the computation, and fixed in place during the computation.
If $jpvt(i) = 0$, the i th column of A is a free column (that is, it may be interchanged during the computation with any other free column).
- rwork* **REAL** for *cgeqp3*
DOUBLE PRECISION for *zgeqp3*.
A workspace array, **DIMENSION** at least $\max(1, 2*n)$.
Used in complex flavors only.

Output Parameters

- a* Overwritten by the factorization data as follows:
If $m \geq n$, the elements below the diagonal are overwritten by the details of the unitary (orthogonal) matrix Q , and the upper triangle is overwritten by the corresponding elements of the upper triangular matrix R .
If $m < n$, the strictly lower triangular part is overwritten by the details of the matrix Q , and the remaining elements are overwritten by the corresponding elements of the m by n upper trapezoidal matrix R .
- tau* **REAL** for *sgeqp3*
DOUBLE PRECISION for *dgeqp3*
COMPLEX for *cgeqp3*
DOUBLE COMPLEX for *zgeqp3*.
Array, **DIMENSION** at least $\max(1, \min(m, n))$.
Contains scalar factors of the elementary reflectors for the matrix Q .

jpvt Overwritten by details of the permutation matrix P in the factorization $AP = QR$. More precisely, the columns of AP are the columns of A in the following order:
jpvt(1), *jpvt*(2), . . . , *jpvt*(n).

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = $-i$, the i th parameter had an illegal value.

Application Notes

To solve a set of least-squares problems minimizing $\|Ax - b\|_2$ for all columns b of a given matrix B , you can call the following:

[?geqp3](#) (this routine) to factorize $AP = QR$;
[?ormqr](#) to compute $C = Q^T B$ (for real matrices);
[?unmqr](#) to compute $C = Q^H B$ (for complex matrices);
[?trsm](#) (a BLAS routine) to solve $RX = C$.

(The columns of the computed X are the permuted least-squares solution vectors x ; the output array *jpvt* specifies the permutation order.)

To compute the elements of Q explicitly, call

[?orgqr](#) (for real matrices)
[?ungqr](#) (for complex matrices).

?orgqr

Generates the real orthogonal matrix Q of the QR factorization formed by [?geqrf](#).

```
call sorgqr ( m, n, k, a, lda, tau, work, lwork, info )
call dorgqr ( m, n, k, a, lda, tau, work, lwork, info )
```

Discussion

The routine generates the whole or part of m by m orthogonal matrix Q of the QR factorization formed by the routines [sgeqrf/dgeqrf](#) (see [page 5-8](#)) or [sgeqpf/dgeqpf](#) (see [page 5-11](#)). Use this routine after a call to [sgeqrf/dgeqrf](#) or [sgeqpf/dgeqpf](#).

Usually Q is determined from the QR factorization of an m by p matrix A with $m \geq p$. To compute the whole matrix Q , use:

```
call ?orgqr ( m, m, p, a, lda, tau, work, lwork, info )
```

To compute the leading p columns of Q (which form an orthonormal basis in the space spanned by the columns of A):

```
call ?orgqr ( m, p, p, a, lda, tau, work, lwork, info )
```

To compute the matrix Q^k of the QR factorization of A 's leading k columns:

```
call ?orgqr ( m, m, k, a, lda, tau, work, lwork, info )
```

To compute the leading k columns of Q^k (which form an orthonormal basis in the space spanned by A 's leading k columns):

```
call ?orgqr ( m, k, k, a, lda, tau, work, lwork, info )
```

Input Parameters

- m **INTEGER**. The order of the orthogonal matrix Q ($m \geq 0$).
- n **INTEGER**. The number of columns of Q to be computed ($0 \leq n \leq m$).
- k **INTEGER**. The number of elementary reflectors whose product defines the matrix Q ($0 \leq k \leq n$).

a, *tau*, *work* REAL for `sorgqr`
DOUBLE PRECISION for `dorgqr`
Arrays:
a(*lda*,*) and *tau*(*) are the arrays returned by
`sgeqrf / dgeqrf` or `sgeqpf / dgeqpf`.
The second dimension of *a* must be at least $\max(1, n)$.
The dimension of *tau* must be at least $\max(1, k)$.
work(*lwork*) is a workspace array.

lda INTEGER. The first dimension of *a*; at least $\max(1, m)$.

lwork INTEGER. The size of the *work* array ($lwork \geq n$)
See *Application notes* for the suggested value of *lwork*.

Output Parameters

a Overwritten by *n* leading columns of the *m* by *m*
orthogonal matrix *Q*.

work(1) If *info* = 0, on exit *work*(1) contains the minimum
value of *lwork* required for optimum performance. Use
this *lwork* for subsequent runs.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = n * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work*(1) and use this value for subsequent runs.

The computed *Q* differs from an exactly orthogonal matrix by a matrix *E* such that $\|E\|_2 = O(\epsilon) \|A\|_2$ where ϵ is the machine precision.

The total number of floating-point operations is approximately

$$4 * m * n * k - 2 * (m + n) * k^2 + (4/3) * k^3.$$

If $n = k$, the number is approximately $(2/3) * n^2 * (3m - n)$.

The complex counterpart of this routine is [?ungqr](#).

?ormqr

Multiplies a real matrix by the orthogonal matrix Q of the QR factorization formed by `?geqrf` or `?geqpf`.

```
call sormqr ( side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info )
call dormqr ( side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info )
```

Discussion

The routine multiplies a real matrix C by Q or Q^T , where Q is the orthogonal matrix Q of the QR factorization formed by the routines `sgeqrf/dgeqrf` (see [page 5-8](#)) or `sgeqpf/dgeqpf` (see [page 5-11](#)).

Depending on the parameters `side` and `trans`, the routine can form one of the matrix products QC , Q^TC , CQ , or CQ^T (overwriting the result on C).

Input Parameters

`side` CHARACTER*1. Must be either 'L' or 'R'.
If `side = 'L'`, Q or Q^T is applied to C from the left.
If `side = 'R'`, Q or Q^T is applied to C from the right.

`trans` CHARACTER*1. Must be either 'N' or 'T'.
If `trans = 'N'`, the routine multiplies C by Q .
If `trans = 'T'`, the routine multiplies C by Q^T .

`m` INTEGER. The number of rows in the matrix C ($m \geq 0$).

`n` INTEGER. The number of columns in C ($n \geq 0$).

`k` INTEGER. The number of elementary reflectors whose product defines the matrix Q . Constraints:
 $0 \leq k \leq m$ if `side = 'L'`;
 $0 \leq k \leq n$ if `side = 'R'`.

`a,work,tau,c` REAL for `sgeqrf`
DOUBLE PRECISION for `dgeqrf`.
Arrays:
`a(lda,*)` and `tau(*)` are the arrays returned by

`sgeqrf / dgeqrf` or `sgeqpf / dgeqpf`.

The second dimension of `a` must be at least $\max(1, k)$.

The dimension of `tau` must be at least $\max(1, k)$.

`c(ldc,*)` contains the matrix `C`.

The second dimension of `c` must be at least $\max(1, n)$

`work(lwork)` is a workspace array.

<code>lda</code>	INTEGER. The first dimension of <code>a</code> . Constraints: <code>lda</code> \geq $\max(1, m)$ if <code>side = 'L'</code> ; <code>lda</code> \geq $\max(1, n)$ if <code>side = 'R'</code> .
<code>ldc</code>	INTEGER. The first dimension of <code>c</code> . Constraint: <code>ldc</code> \geq $\max(1, m)$.
<code>lwork</code>	INTEGER. The size of the <code>work</code> array. Constraints: <code>lwork</code> \geq $\max(1, n)$ if <code>side = 'L'</code> ; <code>lwork</code> \geq $\max(1, m)$ if <code>side = 'R'</code> . See <i>Application notes</i> for the suggested value of <code>lwork</code> .

Output Parameters

<code>c</code>	Overwritten by the product QC , Q^TC , CQ , or CQ^T (as specified by <code>side</code> and <code>trans</code>).
<code>work(1)</code>	If <code>info = 0</code> , on exit <code>work(1)</code> contains the minimum value of <code>lwork</code> required for optimum performance. Use this <code>lwork</code> for subsequent runs.
<code>info</code>	INTEGER. If <code>info = 0</code> , the execution is successful. If <code>info = -i</code> , the <code>i</code> th parameter had an illegal value.

Application Notes

For better performance, try using `lwork = n*blocksize` (if `side = 'L'`) or `lwork = m*blocksize` (if `side = 'R'`) where `blocksize` is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of `lwork` for the first run. On exit, examine `work(1)` and use this value for subsequent runs.

The complex counterpart of this routine is [?unmqr](#).

?ungqr

Generates the complex unitary matrix Q of the QR factorization formed by [?geqrf](#).

```
call cungqr ( m, n, k, a, lda, tau, work, lwork, info )
call zungqr ( m, n, k, a, lda, tau, work, lwork, info )
```

Discussion

The routine generates the whole or part of m by m unitary matrix Q of the QR factorization formed by the routines [cgeqrf/zgeqrf](#) (see [page 5-8](#)) or [cgeqpf/zgeqpf](#) (see [page 5-11](#)). Use this routine after a call to [cgeqrf/zgeqrf](#) or [cgeqpf/zgeqpf](#).

Usually Q is determined from the QR factorization of an m by p matrix A with $m \geq p$. To compute the whole matrix Q , use:

```
call ?ungqr ( m, m, p, a, lda, tau, work, lwork, info )
```

To compute the leading p columns of Q (which form an orthonormal basis in the space spanned by the columns of A):

```
call ?ungqr ( m, p, p, a, lda, tau, work, lwork, info )
```

To compute the matrix Q^k of the QR factorization of A 's leading k columns:

```
call ?ungqr ( m, m, k, a, lda, tau, work, lwork, info )
```

To compute the leading k columns of Q^k (which form an orthonormal basis in the space spanned by A 's leading k columns):

```
call ?ungqr ( m, k, k, a, lda, tau, work, lwork, info )
```

Input Parameters

- m **INTEGER**. The order of the unitary matrix Q ($m \geq 0$).
- n **INTEGER**. The number of columns of Q to be computed ($0 \leq n \leq m$).
- k **INTEGER**. The number of elementary reflectors whose product defines the matrix Q ($0 \leq k \leq n$).

a, *tau*, *work* COMPLEX for *cungqr*
 DOUBLE COMPLEX for *zungqr*
 Arrays:
a(lda,)* and *tau(*)* are the arrays returned by
cgeqrf/zgeqrf or *cgeqpf/zgeqpf*.
 The second dimension of *a* must be at least $\max(1, n)$.
 The dimension of *tau* must be at least $\max(1, k)$.
work(lwork) is a workspace array.

lda INTEGER. The first dimension of *a*; at least $\max(1, m)$.

lwork INTEGER. The size of the *work* array ($lwork \geq n$).
 See *Application notes* for the suggested value of *lwork*.

Output Parameters

a Overwritten by *n* leading columns of the *m* by *m* unitary matrix *Q*.

work(1) If *info* = 0, on exit *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = n * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

The computed *Q* differs from an exactly unitary matrix by a matrix *E* such that $\|E\|_2 = O(\epsilon) \|A\|_2$ where ϵ is the machine precision.

The total number of floating-point operations is approximately

$$16 * m * n * k - 8 * (m + n) * k^2 + (16/3) * k^3.$$

If $n = k$, the number is approximately $(8/3) * n^2 * (3m - n)$.

The real counterpart of this routine is [?orgqr](#).

?unmqr

Multiplies a complex matrix by the unitary matrix Q of the QR factorization formed by ?geqrf.

```
call cunmqr ( side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info )
call zunmqr ( side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info )
```

Discussion

The routine multiplies a rectangular complex matrix C by Q or Q^H , where Q is the unitary matrix Q of the QR factorization formed by the routines `cgeqrf/zgeqrf` (see [page 5-8](#)) or `cgeqpf/zgeqpf` (see [page 5-11](#)).

Depending on the parameters `side` and `trans`, the routine can form one of the matrix products QC , Q^HC , CQ , or CQ^H (overwriting the result on C).

Input Parameters

`side` CHARACTER*1. Must be either 'L' or 'R'.
If `side = 'L'`, Q or Q^H is applied to C from the left.
If `side = 'R'`, Q or Q^H is applied to C from the right.

`trans` CHARACTER*1. Must be either 'N' or 'C'.
If `trans = 'N'`, the routine multiplies C by Q .
If `trans = 'C'`, the routine multiplies C by Q^H .

`m` INTEGER. The number of rows in the matrix C ($m \geq 0$).

`n` INTEGER. The number of columns in C ($n \geq 0$).

`k` INTEGER. The number of elementary reflectors whose product defines the matrix Q . Constraints:
 $0 \leq k \leq m$ if `side = 'L'`;
 $0 \leq k \leq n$ if `side = 'R'`.

`a,work,tau,c` COMPLEX for `cgeqrf`
DOUBLE COMPLEX for `zgeqrf`.

Arrays:
`a(lda,*)` and `tau(*)` are the arrays returned by

`cgeqrf / zgeqrf` or `cgeqpf / zgeqpf`.

The second dimension of `a` must be at least $\max(1, k)$.

The dimension of `tau` must be at least $\max(1, k)$.

`c(ldc,*)` contains the matrix `C`.

The second dimension of `c` must be at least $\max(1, n)$

`work(lwork)` is a workspace array.

<code>lda</code>	INTEGER. The first dimension of <code>a</code> . Constraints: <code>lda</code> \geq $\max(1, m)$ if <code>side = 'L'</code> ; <code>lda</code> \geq $\max(1, n)$ if <code>side = 'R'</code> .
<code>ldc</code>	INTEGER. The first dimension of <code>c</code> . Constraint: <code>ldc</code> \geq $\max(1, m)$.
<code>lwork</code>	INTEGER. The size of the <code>work</code> array. Constraints: <code>lwork</code> \geq $\max(1, n)$ if <code>side = 'L'</code> ; <code>lwork</code> \geq $\max(1, m)$ if <code>side = 'R'</code> . See <i>Application notes</i> for the suggested value of <code>lwork</code> .

Output Parameters

<code>c</code>	Overwritten by the product QC , $Q^H C$, CQ , or CQ^H (as specified by <code>side</code> and <code>trans</code>).
<code>work(1)</code>	If <code>info = 0</code> , on exit <code>work(1)</code> contains the minimum value of <code>lwork</code> required for optimum performance. Use this <code>lwork</code> for subsequent runs.
<code>info</code>	INTEGER. If <code>info = 0</code> , the execution is successful. If <code>info = -i</code> , the <code>i</code> th parameter had an illegal value.

Application Notes

For better performance, try using `lwork = n*blocksize` (if `side = 'L'`) or `lwork = m*blocksize` (if `side = 'R'`) where `blocksize` is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of `lwork` for the first run. On exit, examine `work(1)` and use this value for subsequent runs.

The real counterpart of this routine is [?ormqr](#).

?gelqf

Computes the LQ factorization of a general m by n matrix.

```
call sgelqf ( m, n, a, lda, tau, work, lwork, info )
call dgelqf ( m, n, a, lda, tau, work, lwork, info )
call cgelqf ( m, n, a, lda, tau, work, lwork, info )
call zgelqf ( m, n, a, lda, tau, work, lwork, info )
```

Discussion

The routine forms the LQ factorization of a general m by n matrix A (see *Orthogonal Factorizations* on [page 5-6](#)). No pivoting is performed.

The routine does not form the matrix Q explicitly. Instead, Q is represented as a product of $\min(m, n)$ elementary reflectors. Routines are provided to work with Q in this representation.

Input Parameters

m **INTEGER**. The number of rows in the matrix A ($m \geq 0$).

n **INTEGER**. The number of columns in A ($n \geq 0$).

$a, work$ **REAL** for `sgelqf`
DOUBLE PRECISION for `dgelqf`
COMPLEX for `cgelqf`
DOUBLE COMPLEX for `zgelqf`.

Arrays:

$a(lda, *)$ contains the matrix A .

The second dimension of a must be at least $\max(1, n)$.

$work(lwork)$ is a workspace array.

lda **INTEGER**. The first dimension of a ; at least $\max(1, m)$.

$lwork$ **INTEGER**. The size of the $work$ array; at least $\max(1, m)$.
 See *Application notes* for the suggested value of $lwork$.

Output Parameters

<i>a</i>	Overwritten by the factorization data as follows: If $m \leq n$, the elements above the diagonal are overwritten by the details of the unitary (orthogonal) matrix Q , and the lower triangle is overwritten by the corresponding elements of the lower triangular matrix L . If $m > n$, the strictly upper triangular part is overwritten by the details of the matrix Q , and the remaining elements are overwritten by the corresponding elements of the m by n lower trapezoidal matrix L .
<i>tau</i>	REAL for <code>sgelqf</code> DOUBLE PRECISION for <code>dgelqf</code> COMPLEX for <code>cgelqf</code> DOUBLE COMPLEX for <code>zgelqf</code> . Array, DIMENSION at least $\max(1, \min(m, n))$. Contains additional information on the matrix Q .
<i>work(1)</i>	If <i>info</i> = 0, on exit <i>work(1)</i> contains the minimum value of <i>lwork</i> required for optimum performance. Use this <i>lwork</i> for subsequent runs.
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = m * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*.

If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

The computed factorization is the exact factorization of a matrix $A + E$, where $\|E\|_2 = O(\epsilon) \|A\|_2$.

The approximate number of floating-point operations for real flavors is

$$(4/3)n^3 \quad \text{if } m = n,$$

$$(2/3)n^2(3m-n) \quad \text{if } m > n,$$

$$(2/3)m^2(3n-m) \quad \text{if } m < n.$$

The number of operations for complex flavors is 4 times greater.

To find the minimum-norm solution of an underdetermined least-squares problem minimizing $\|Ax - b\|_2$ for all columns b of a given matrix B , you can call the following:

`?gelsf` (this routine) to factorize $A = LQ$;

`?trsm` (a BLAS routine) to solve $LY = B$ for Y ;

`?ormlq` to compute $X = (Q_1)^T Y$ (for real matrices);

`?unmlq` to compute $X = (Q_1)^H Y$ (for complex matrices).

(The columns of the computed X are the minimum-norm solution vectors x . Here A is an m by n matrix with $m < n$; Q_1 denotes the first m columns of Q).

To compute the elements of Q explicitly, call

`?orglq` (for real matrices)

`?unglq` (for complex matrices).

?orglq

Generates the real orthogonal matrix Q of the LQ factorization formed by [?gelqf](#).

```
call sorglq ( m, n, k, a, lda, tau, work, lwork, info )
call dorglq ( m, n, k, a, lda, tau, work, lwork, info )
```

Discussion

The routine generates the whole or part of n by n orthogonal matrix Q of the LQ factorization formed by the routines [sgelqf/dgelqf](#) (see [page 5-25](#)). Use this routine after a call to [sgelqf/dgelqf](#).

Usually Q is determined from the LQ factorization of an p by n matrix A with $n \geq p$. To compute the whole matrix Q , use:

```
call ?orglq ( n, n, p, a, lda, tau, work, lwork, info )
```

To compute the leading p rows of Q (which form an orthonormal basis in the space spanned by the rows of A):

```
call ?orglq ( p, n, p, a, lda, tau, work, lwork, info )
```

To compute the matrix Q^k of the LQ factorization of A 's leading k rows:

```
call ?orglq ( n, n, k, a, lda, tau, work, lwork, info )
```

To compute the leading k rows of Q^k (which form an orthonormal basis in the space spanned by A 's leading k rows):

```
call ?orgqr ( k, n, k, a, lda, tau, work, lwork, info )
```

Input Parameters

- m **INTEGER.** The number of rows of Q to be computed ($0 \leq m \leq n$).
- n **INTEGER.** The order of the orthogonal matrix Q ($n \geq m$).
- k **INTEGER.** The number of elementary reflectors whose product defines the matrix Q ($0 \leq k \leq m$).

`a`, `tau`, `work` REAL for `sorglq`
 DOUBLE PRECISION for `dorglq`
 Arrays:
`a(lda,*)` and `tau(*)` are the arrays returned by
`sgelqf/dgelqf`.
 The second dimension of `a` must be at least $\max(1, n)$.
 The dimension of `tau` must be at least $\max(1, k)$.
`work(lwork)` is a workspace array.

`lda` INTEGER. The first dimension of `a`; at least $\max(1, m)$.

`lwork` INTEGER. The size of the `work` array; at least $\max(1, m)$.
 See *Application notes* for the suggested value of `lwork`.

Output Parameters

`a` Overwritten by `m` leading rows of the `n` by `n` orthogonal matrix Q .

`work(1)` If `info` = 0, on exit `work(1)` contains the minimum value of `lwork` required for optimum performance. Use this `lwork` for subsequent runs.

`info` INTEGER.
 If `info` = 0, the execution is successful.
 If `info` = $-i$, the i th parameter had an illegal value.

Application Notes

For better performance, try using `lwork` = $m \cdot \text{blocksize}$, where `blocksize` is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of `lwork` for the first run. On exit, examine `work(1)` and use this value for subsequent runs.

The computed Q differs from an exactly orthogonal matrix by a matrix E such that $\|E\|_2 = O(\epsilon) \|A\|_2$ where ϵ is the machine precision.

The total number of floating-point operations is approximately

$$4 \cdot m \cdot n \cdot k - 2 \cdot (m + n) \cdot k^2 + (4/3) \cdot k^3.$$

If $m = k$, the number is approximately $(2/3) \cdot m^2 \cdot (3n - m)$.

The complex counterpart of this routine is [?unglq](#).

?ormlq

Multiplies a real matrix by the orthogonal matrix Q of the LQ factorization formed by ?gelqf.

```
call sormlq ( side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info )
call dormlq ( side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info )
```

Discussion

The routine multiplies a real m -by- n matrix C by Q or Q^T , where Q is the orthogonal matrix Q of the LQ factorization formed by the routine [sgelqf/dgelqf](#) (see [page 5-25](#)).

Depending on the parameters *side* and *trans*, the routine can form one of the matrix products QC , Q^TC , CQ , or CQ^T (overwriting the result on C).

Input Parameters

side CHARACTER*1. Must be either 'L' or 'R'.
If *side* = 'L', Q or Q^T is applied to C from the left.
If *side* = 'R', Q or Q^T is applied to C from the right.

trans CHARACTER*1. Must be either 'N' or 'T'.
If *trans* = 'N', the routine multiplies C by Q .
If *trans* = 'T', the routine multiplies C by Q^T .

m INTEGER. The number of rows in the matrix C ($m \geq 0$).

n INTEGER. The number of columns in C ($n \geq 0$).

k INTEGER. The number of elementary reflectors whose product defines the matrix Q . Constraints:
 $0 \leq k \leq m$ if *side* = 'L';
 $0 \leq k \leq n$ if *side* = 'R'.

a,work,tau,c REAL for sormlq
DOUBLE PRECISION for dormlq.
Arrays:
a(lda,)* and *tau(*)* are arrays returned by ?gelqf.

The second dimension of a must be:
 at least $\max(1, m)$ if $side = 'L'$;
 at least $\max(1, n)$ if $side = 'R'$.
 The dimension of τ must be at least $\max(1, k)$.

$c(ldc, *)$ contains the matrix C .

The second dimension of c must be at least $\max(1, n)$

$work(lwork)$ is a workspace array.

lda INTEGER. The first dimension of a ; $lda \geq \max(1, k)$.

ldc INTEGER. The first dimension of c ; $ldc \geq \max(1, m)$.

$lwork$ INTEGER. The size of the $work$ array. Constraints:

$lwork \geq \max(1, n)$ if $side = 'L'$;

$lwork \geq \max(1, m)$ if $side = 'R'$.

See *Application notes* for the suggested value of $lwork$.

Output Parameters

c Overwritten by the product QC , Q^TC , CQ , or CQ^T
 (as specified by $side$ and $trans$).

$work(1)$ If $info = 0$, on exit $work(1)$ contains the minimum
 value of $lwork$ required for optimum performance. Use
 this $lwork$ for subsequent runs.

$info$ INTEGER.
 If $info = 0$, the execution is successful.
 If $info = -i$, the i th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = n * blocksize$ (if $side = 'L'$) or
 $lwork = m * blocksize$ (if $side = 'R'$) where $blocksize$ is a
 machine-dependent value (typically, 16 to 64) required for optimum
 performance of the *blocked algorithm*. If you are in doubt how much
 workspace to supply, use a generous value of $lwork$ for the first run. On
 exit, examine $work(1)$ and use this value for subsequent runs.

The complex counterpart of this routine is [?unmlq](#).

?unglq

Generates the complex unitary matrix Q of the LQ factorization formed by ?gelqf.

```
call cunglq ( m, n, k, a, lda, tau, work, lwork, info )
call zunglq ( m, n, k, a, lda, tau, work, lwork, info )
```

Discussion

The routine generates the whole or part of n by n unitary matrix Q of the LQ factorization formed by the routines `cgelqf/zgelqf` (see [page 5-25](#)). Use this routine after a call to `cgelqf/zgelqf`.

Usually Q is determined from the LQ factorization of an p by n matrix A with $n \geq p$. To compute the whole matrix Q , use:

```
call ?unglq ( n, n, p, a, lda, tau, work, lwork, info )
```

To compute the leading p rows of Q (which form an orthonormal basis in the space spanned by the rows of A):

```
call ?unglq ( p, n, p, a, lda, tau, work, lwork, info )
```

To compute the matrix Q^k of the LQ factorization of A 's leading k rows:

```
call ?unglq ( n, n, k, a, lda, tau, work, lwork, info )
```

To compute the leading k rows of Q^k (which form an orthonormal basis in the space spanned by A 's leading k rows):

```
call ?ungqr ( k, n, k, a, lda, tau, work, lwork, info )
```

Input Parameters

- m **INTEGER.** The number of rows of Q to be computed ($0 \leq m \leq n$).
- n **INTEGER.** The order of the unitary matrix Q ($n \geq m$).
- k **INTEGER.** The number of elementary reflectors whose product defines the matrix Q ($0 \leq k \leq m$).

a, *tau*, *work* **COMPLEX** for *cunglq*
 DOUBLE COMPLEX for *zunglq*
 Arrays:
a(*lda*,*) and *tau*(*) are the arrays returned by
sgelqf/dgelqf.
 The second dimension of *a* must be at least $\max(1, n)$.
 The dimension of *tau* must be at least $\max(1, k)$.
work (*lwork*) is a workspace array.

lda **INTEGER**. The first dimension of *a*; at least $\max(1, m)$.

lwork **INTEGER**. The size of the *work* array; at least $\max(1, m)$.
 See *Application notes* for the suggested value of *lwork*.

Output Parameters

a Overwritten by *m* leading rows of the *n* by *n* unitary
 matrix *Q*.

work(1) If *info* = 0, on exit *work*(1) contains the minimum
 value of *lwork* required for optimum performance. Use
 this *lwork* for subsequent runs.

info **INTEGER**.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using *lwork* = *m***blocksize*, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work*(1) and use this value for subsequent runs.

The computed *Q* differs from an exactly unitary matrix by a matrix *E* such that $\|E\|_2 = O(\epsilon) \|A\|_2$ where ϵ is the machine precision.

The total number of floating-point operations is approximately

$$16*m*n*k - 8*(m+n)*k^2 + (16/3)*k^3.$$

If *m* = *k*, the number is approximately $(8/3)*m^2*(3n - m)$.

The real counterpart of this routine is [?orglq](#).

?unmlq

Multiplies a complex matrix by the unitary matrix Q of the LQ factorization formed by ?gelqf.

```
call cunmlq ( side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info )
call zunmlq ( side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info )
```

Discussion

The routine multiplies a real m -by- n matrix C by Q or Q^H , where Q is the unitary matrix Q of the LQ factorization formed by the routine [cgelqf/zgelqf](#) (see [page 5-25](#)).

Depending on the parameters *side* and *trans*, the routine can form one of the matrix products QC , Q^HC , CQ , or CQ^H (overwriting the result on C).

Input Parameters

<i>side</i>	CHARACTER*1. Must be either 'L' or 'R'. If <i>side</i> = 'L', Q or Q^H is applied to C from the left. If <i>side</i> = 'R', Q or Q^H is applied to C from the right.
<i>trans</i>	CHARACTER*1. Must be either 'N' or 'C'. If <i>trans</i> = 'N', the routine multiplies C by Q . If <i>trans</i> = 'C', the routine multiplies C by Q^H .
<i>m</i>	INTEGER. The number of rows in the matrix C ($m \geq 0$).
<i>n</i>	INTEGER. The number of columns in C ($n \geq 0$).
<i>k</i>	INTEGER. The number of elementary reflectors whose product defines the matrix Q . Constraints: $0 \leq k \leq m$ if <i>side</i> = 'L'; $0 \leq k \leq n$ if <i>side</i> = 'R'.
<i>a,work,tau,c</i>	COMPLEX for cunmlq DOUBLE COMPLEX for zunmlq. Arrays: <i>a(lda,*)</i> and <i>tau(*)</i> are arrays returned by ?gelqf.

The second dimension of a must be:
 at least $\max(1, m)$ if $side = 'L'$;
 at least $\max(1, n)$ if $side = 'R'$.
 The dimension of τ must be at least $\max(1, k)$.

$c(ldc, *)$ contains the matrix C .

The second dimension of c must be at least $\max(1, n)$

$work(lwork)$ is a workspace array.

lda INTEGER. The first dimension of a ; $lda \geq \max(1, k)$.

ldc INTEGER. The first dimension of c ; $ldc \geq \max(1, m)$.

$lwork$ INTEGER. The size of the $work$ array. Constraints:

$lwork \geq \max(1, n)$ if $side = 'L'$;

$lwork \geq \max(1, m)$ if $side = 'R'$.

See *Application notes* for the suggested value of $lwork$.

Output Parameters

c Overwritten by the product QC , $Q^H C$, CQ , or CQ^H
 (as specified by $side$ and $trans$).

$work(1)$ If $info = 0$, on exit $work(1)$ contains the minimum
 value of $lwork$ required for optimum performance. Use
 this $lwork$ for subsequent runs.

$info$ INTEGER.
 If $info = 0$, the execution is successful.
 If $info = -i$, the i th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = n * blocksize$ (if $side = 'L'$) or
 $lwork = m * blocksize$ (if $side = 'R'$) where $blocksize$ is a
 machine-dependent value (typically, 16 to 64) required for optimum
 performance of the *blocked algorithm*. If you are in doubt how much
 workspace to supply, use a generous value of $lwork$ for the first run. On
 exit, examine $work(1)$ and use this value for subsequent runs.

The real counterpart of this routine is [?ormlq](#).

?geqlf

Computes the QL factorization of a general m by n matrix.

```
call sgeqlf ( m, n, a, lda, tau, work, lwork, info )
call dgeqlf ( m, n, a, lda, tau, work, lwork, info )
call cgeqlf ( m, n, a, lda, tau, work, lwork, info )
call zgeqlf ( m, n, a, lda, tau, work, lwork, info )
```

Discussion

The routine forms the QL factorization of a general m -by- n matrix A . No pivoting is performed.

The routine does not form the matrix Q explicitly. Instead, Q is represented as a product of $\min(m, n)$ elementary reflectors. Routines are provided to work with Q in this representation.

Input Parameters

m **INTEGER**. The number of rows in the matrix A ($m \geq 0$).

n **INTEGER**. The number of columns in A ($n \geq 0$).

a, work **REAL** for `sgeqlf`
DOUBLE PRECISION for `dgeqlf`
COMPLEX for `cgeqlf`
DOUBLE COMPLEX for `zgeqlf`.

Arrays:
a(*lda*,*) contains the matrix A .
The second dimension of *a* must be at least $\max(1, n)$.
work(*lwork*) is a workspace array.

lda **INTEGER**. The first dimension of *a*; at least $\max(1, m)$.

lwork **INTEGER**. The size of the *work* array; at least $\max(1, n)$.
See *Application notes* for the suggested value of *lwork*.

Output Parameters

- a* Overwritten on exit by the factorization data as follows:
 if $m \geq n$, the lower triangle of the subarray
 $a(m-n+1:m, 1:n)$ contains the n -by- n lower triangular
 matrix L ;
 if $m \leq n$, the elements on and below the $(n-m)$ th
 superdiagonal contain the m -by- n lower trapezoidal
 matrix L ;
 in both cases, the remaining elements, with the array
 τ , represent the orthogonal/unitary matrix Q as a
 product of elementary reflectors.
- tau* REAL for `sgeqlf`
 DOUBLE PRECISION for `dgeqlf`
 COMPLEX for `cgeqlf`
 DOUBLE COMPLEX for `zgeqlf`.
 Array, DIMENSION at least $\max(1, \min(m, n))$.
 Contains scalar factors of the elementary reflectors for
 the matrix Q .
- work(1)* If $info = 0$, on exit *work(1)* contains the minimum
 value of *lwork* required for optimum performance.
- info* INTEGER.
 If $info = 0$, the execution is successful.
 If $info = -i$, the i th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = n * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

Related routines include:

- [?orgql](#) to generate matrix Q (for real matrices);
- [?ungql](#) to generate matrix Q (for complex matrices);
- [?ormql](#) to apply matrix Q (for real matrices);
- [?unmq1](#) to apply matrix Q (for complex matrices).

?orgql

Generates the real matrix Q of the QL factorization formed by ?geqlf.

```
call sorgql ( m, n, k, a, lda, tau, work, lwork, info )
call dorgql ( m, n, k, a, lda, tau, work, lwork, info )
```

Discussion

The routine generates an m -by- n real matrix Q with orthonormal columns, which is defined as the last n columns of a product of k elementary reflectors H_i of order m : $Q = H_k \cdot \cdot \cdot H_2 H_1$ as returned by the routines [sgeqlf/dgeqlf](#). Use this routine after a call to [sgeqlf/dgeqlf](#).

Input Parameters

m **INTEGER.** The number of rows of the matrix Q ($m \geq 0$).

n **INTEGER.** The number of columns of the matrix Q ($m \geq n \geq 0$).

k **INTEGER.** The number of elementary reflectors whose product defines the matrix Q ($n \geq k \geq 0$).

a , τ , $work$ **REAL** for `sorgql`
DOUBLE PRECISION for `dorgql`
 Arrays: $a(lda, *)$, $\tau(*)$, $work(lwork)$.

On entry, the $(n - k + i)$ th column of a must contain the vector which defines the elementary reflector H_i , for $i = 1, 2, \dots, k$, as returned by [sgeqlf/dgeqlf](#) in the last k columns of its array argument a ;

$\tau(i)$ must contain the scalar factor of the elementary reflector H_i , as returned by [sgeqlf/dgeqlf](#);

The second dimension of a must be at least $\max(1, n)$.

The dimension of τ must be at least $\max(1, k)$.

$work(lwork)$ is a workspace array.

lda **INTEGER**. The first dimension of *a*; at least $\max(1, m)$.
lwork **INTEGER**. The size of the *work* array; at least $\max(1, n)$.
See *Application notes* for the suggested value of *lwork*.

Output Parameters

a Overwritten by the *m*-by-*n* matrix *Q*.
work(1) If *info* = 0, on exit *work(1)* contains the minimum
value of *lwork* required for optimum performance. Use
this *lwork* for subsequent runs.
info **INTEGER**.
If *info* = 0, the execution is successful.
If *info* = *-i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = n * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

The complex counterpart of this routine is [?ungql](#).

?ungql

Generates the complex matrix Q of the QL factorization formed by ?geqlf.

```
call cungql ( m, n, k, a, lda, tau, work, lwork, info )
call zungql ( m, n, k, a, lda, tau, work, lwork, info )
```

Discussion

The routine generates an m -by- n complex matrix Q with orthonormal columns, which is defined as the last n columns of a product of k elementary reflectors H_i of order m : $Q = H_k \cdot \cdot \cdot H_2 H_1$ as returned by the routines [cgeqlf/zgeqlf](#). Use this routine after a call to [cgeqlf/zgeqlf](#).

Input Parameters

m **INTEGER.** The number of rows of the matrix Q ($m \geq 0$).

n **INTEGER.** The number of columns of the matrix Q ($m \geq n \geq 0$).

k **INTEGER.** The number of elementary reflectors whose product defines the matrix Q ($n \geq k \geq 0$).

a , tau , $work$ **COMPLEX** for [cungql](#)
DOUBLE COMPLEX for [zungql](#)
 Arrays: $a(lda, *)$, $tau(*)$, $work(lwork)$.

On entry, the $(n - k + i)$ th column of a must contain the vector which defines the elementary reflector H_i , for $i = 1, 2, \dots, k$, as returned by [cgeqlf/zgeqlf](#) in the last k columns of its array argument a ;

$tau(i)$ must contain the scalar factor of the elementary reflector H_i , as returned by [cgeqlf/zgeqlf](#);

The second dimension of a must be at least $\max(1, n)$.

The dimension of tau must be at least $\max(1, k)$.

$work(lwork)$ is a workspace array.

lda **INTEGER**. The first dimension of *a*; at least $\max(1, m)$.
lwork **INTEGER**. The size of the *work* array; at least $\max(1, n)$.
See *Application notes* for the suggested value of *lwork*.

Output Parameters

a Overwritten by the *m*-by-*n* matrix *Q*.
work(1) If *info* = 0, on exit *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.
info **INTEGER**.
If *info* = 0, the execution is successful.
If *info* = *-i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = n * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

The real counterpart of this routine is [?orgql](#).

?ormql

Multiplies a real matrix by the orthogonal matrix Q of the QL factorization formed by [?geqlf](#).

```
call sormql ( side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info )
call dormql ( side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info )
```

Discussion

This routine multiplies a real m -by- n matrix C by Q or Q^T , where Q is the orthogonal matrix Q of the QL factorization formed by the routine [sgeqlf/dgeqlf](#).

Depending on the parameters *side* and *trans*, the routine [?ormql](#) can form one of the matrix products QC , Q^TC , CQ , or CQ^T (overwriting the result over C).

Input Parameters

<i>side</i>	CHARACTER*1. Must be either 'L' or 'R'. If <i>side</i> = 'L', Q or Q^T is applied to C from the left. If <i>side</i> = 'R', Q or Q^T is applied to C from the right.
<i>trans</i>	CHARACTER*1. Must be either 'N' or 'T'. If <i>trans</i> = 'N', the routine multiplies C by Q . If <i>trans</i> = 'T', the routine multiplies C by Q^T .
<i>m</i>	INTEGER. The number of rows in the matrix C ($m \geq 0$).
<i>n</i>	INTEGER. The number of columns in C ($n \geq 0$).
<i>k</i>	INTEGER. The number of elementary reflectors whose product defines the matrix Q . Constraints: $0 \leq k \leq m$ if <i>side</i> = 'L'; $0 \leq k \leq n$ if <i>side</i> = 'R'.

a, *tau*, *c*, *work* REAL for *sormql*
DOUBLE PRECISION for *dormql*.
Arrays: *a*(*lda*, *), *tau*(*), *c*(*ldc*, *),
work(*lwork*).

On entry, the *i*th column of *a* must contain the vector which defines the elementary reflector H_i , for $i = 1, 2, \dots, k$, as returned by *sgeqlf*/*dgeqlf* in the last *k* columns of its array argument *a*.
The second dimension of *a* must be at least $\max(1, k)$.
tau(*i*) must contain the scalar factor of the elementary reflector H_i , as returned by *sgeqlf*/*dgeqlf*.
The dimension of *tau* must be at least $\max(1, k)$.
c(*ldc*, *) contains the *m*-by-*n* matrix *C*.
The second dimension of *c* must be at least $\max(1, n)$.
work(*lwork*) is a workspace array.

lda INTEGER. The first dimension of *a*;
if *side* = 'L', $lda \geq \max(1, m)$;
if *side* = 'R', $lda \geq \max(1, n)$.

ldc INTEGER. The first dimension of *c*; $ldc \geq \max(1, m)$.

lwork INTEGER. The size of the *work* array. Constraints:
lwork $\geq \max(1, n)$ if *side* = 'L';
lwork $\geq \max(1, m)$ if *side* = 'R'.
See *Application notes* for the suggested value of *lwork*.

Output Parameters

c Overwritten by the product QC , Q^TC , CQ , or CQ^T (as specified by *side* and *trans*).

work(1) If *info* = 0, on exit *work*(1) contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = n * blocksize$ (if $side = 'L'$) or $lwork = m * blocksize$ (if $side = 'R'$) where $blocksize$ is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of $lwork$ for the first run. On exit, examine $work(1)$ and use this value for subsequent runs.

The complex counterpart of this routine is [?unmql](#).

?unmql

Multiplies a complex matrix by the unitary matrix Q of the QL factorization formed by `?geqlf`.

```
call cunmql ( side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info )
call zunmql ( side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info )
```

Discussion

The routine multiplies a complex m -by- n matrix C by Q or Q^H , where Q is the unitary matrix Q of the QL factorization formed by the routine `cgeqlf/zgeqlf`.

Depending on the parameters `side` and `trans`, the routine `?unmql` can form one of the matrix products QC , $Q^H C$, CQ , or CQ^H (overwriting the result over C).

Input Parameters

`side` CHARACTER*1. Must be either 'L' or 'R'.
 If `side` = 'L', Q or Q^H is applied to C from the left.
 If `side` = 'R', Q or Q^H is applied to C from the right.

`trans` CHARACTER*1. Must be either 'N' or 'C'.
 If `trans` = 'N', the routine multiplies C by Q .
 If `trans` = 'C', the routine multiplies C by Q^H .

`m` INTEGER. The number of rows in the matrix C ($m \geq 0$).

`n` INTEGER. The number of columns in C ($n \geq 0$).

`k` INTEGER. The number of elementary reflectors whose product defines the matrix Q . Constraints:
 $0 \leq k \leq m$ if `side` = 'L';
 $0 \leq k \leq n$ if `side` = 'R'.

a, tau, c, work COMPLEX for `cunmql`
DOUBLE COMPLEX for `zunmql`.
Arrays: *a(lda, *)*, *tau(*)*, *c(ldc, *)*,
work(lwork).

On entry, the *i*th column of *a* must contain the vector which defines the elementary reflector H_i , for $i = 1, 2, \dots, k$, as returned by `cgeqlf/zgeqlf` in the last *k* columns of its array argument *a*.
The second dimension of *a* must be at least $\max(1, k)$.
tau(i) must contain the scalar factor of the elementary reflector H_i , as returned by `cgeqlf/zgeqlf`.
The dimension of *tau* must be at least $\max(1, k)$.
*c(ldc, *)* contains the *m*-by-*n* matrix *C*.
The second dimension of *c* must be at least $\max(1, n)$.
work(lwork) is a workspace array.

lda INTEGER. The first dimension of *a*;
if *side* = 'L', $lda \geq \max(1, m)$;
if *side* = 'R', $lda \geq \max(1, n)$.

ldc INTEGER. The first dimension of *c*; $ldc \geq \max(1, m)$.

lwork INTEGER. The size of the *work* array. Constraints:
 $lwork \geq \max(1, n)$ if *side* = 'L';
 $lwork \geq \max(1, m)$ if *side* = 'R'.
See *Application notes* for the suggested value of *lwork*.

Output Parameters

c Overwritten by the product QC , $Q^H C$, CQ , or CQ^H (as specified by *side* and *trans*).

work(1) If *info* = 0, on exit *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = n * blocksize$ (if $side = 'L'$) or $lwork = m * blocksize$ (if $side = 'R'$) where $blocksize$ is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of $lwork$ for the first run. On exit, examine $work(1)$ and use this value for subsequent runs.

The real counterpart of this routine is [?ormql](#).

?gerqf

Computes the RQ factorization of a general m by n matrix.

```
call sgerqf ( m, n, a, lda, tau, work, lwork, info )
call dgerqf ( m, n, a, lda, tau, work, lwork, info )
call cgerqf ( m, n, a, lda, tau, work, lwork, info )
call zgerqf ( m, n, a, lda, tau, work, lwork, info )
```

Discussion

The routine forms the RQ factorization of a general m -by- n matrix A . No pivoting is performed.

The routine does not form the matrix Q explicitly. Instead, Q is represented as a product of $\min(m, n)$ elementary reflectors. Routines are provided to work with Q in this representation.

Input Parameters

m **INTEGER**. The number of rows in the matrix A ($m \geq 0$).

n **INTEGER**. The number of columns in A ($n \geq 0$).

a, work **REAL** for `sgerqf`
DOUBLE PRECISION for `dgerqf`
COMPLEX for `cgerqf`
DOUBLE COMPLEX for `zgerqf`.

Arrays:
a(*lda*,*) contains the m -by- n matrix A .
The second dimension of *a* must be at least $\max(1, n)$.
work(*lwork*) is a workspace array.

lda **INTEGER**. The first dimension of *a*; at least $\max(1, m)$.

lwork **INTEGER**. The size of the *work* array;
lwork $\geq \max(1, m)$.
See [Application notes](#) for the suggested value of *lwork*.

Output Parameters

- a* Overwritten on exit by the factorization data as follows:
 if $m \leq n$, the upper triangle of the subarray
 $a(1:m, n-m+1:n)$ contains the m -by- m upper triangular
 matrix R ;
 if $m \geq n$, the elements on and above the $(m-n)$ th
 subdiagonal contain the m -by- n upper trapezoidal
 matrix R ;
 in both cases, the remaining elements, with the array
tau, represent the orthogonal/unitary matrix Q as a
 product of $\min(m,n)$ elementary reflectors.
- tau* REAL for *sgerqf*
 DOUBLE PRECISION for *dgerqf*
 COMPLEX for *cgerqf*
 DOUBLE COMPLEX for *zgerqf*.
 Array, DIMENSION at least $\max(1, \min(m, n))$.
 Contains scalar factors of the elementary reflectors for
 the matrix Q .
- work(1)* If *info* = 0, on exit *work(1)* contains the minimum
 value of *lwork* required for optimum performance.
- info* INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = $-i$, the i th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = m * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

Related routines include:

- [?orgqr](#) to generate matrix Q (for real matrices);
[?ungqr](#) to generate matrix Q (for complex matrices);
[?ormqr](#) to apply matrix Q (for real matrices);
[?unmqr](#) to apply matrix Q (for complex matrices).

?orgrq

Generates the real matrix Q of the RQ factorization formed by ?gerqf.

```
call sorgrq ( m, n, k, a, lda, tau, work, lwork, info )
call dorgrq ( m, n, k, a, lda, tau, work, lwork, info )
```

Discussion

The routine generates an m -by- n real matrix Q with orthonormal rows, which is defined as the last m rows of a product of k elementary reflectors H_i of order n : $Q = H_1 H_2 \cdots H_k$ as returned by the routines [sgerqf/dgerqf](#). Use this routine after a call to [sgerqf/dgerqf](#).

Input Parameters

m **INTEGER.** The number of rows of the matrix Q ($m \geq 0$).

n **INTEGER.** The number of columns of the matrix Q ($n \geq m$).

k **INTEGER.** The number of elementary reflectors whose product defines the matrix Q ($m \geq k \geq 0$).

a , tau , $work$ **REAL** for [sorgrq](#)
DOUBLE PRECISION for [dorgrq](#)
 Arrays: $a(lda, *)$, $tau(*)$, $work(lwork)$.

On entry, the $(m - k + i)$ th row of a must contain the vector which defines the elementary reflector H_i , for $i = 1, 2, \dots, k$, as returned by [sgerqf/dgerqf](#) in the last k rows of its array argument a ;

$tau(i)$ must contain the scalar factor of the elementary reflector H_i , as returned by [sgerqf/dgerqf](#);

The second dimension of a must be at least $\max(1, n)$.

The dimension of tau must be at least $\max(1, k)$.

$work(lwork)$ is a workspace array.

lda **INTEGER**. The first dimension of *a*; at least $\max(1, m)$.
lwork **INTEGER**. The size of the *work* array; at least $\max(1, m)$.
See *Application notes* for the suggested value of *lwork*.

Output Parameters

a Overwritten by the *m*-by-*n* matrix *Q*.
work(1) If *info* = 0, on exit *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.
info **INTEGER**.
If *info* = 0, the execution is successful.
If *info* = *-i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = m * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

The complex counterpart of this routine is [?ungqr](#).

?ungrq

Generates the complex matrix Q of the RQ factorization formed by `?gerqf`.

```
call cungrq ( m, n, k, a, lda, tau, work, lwork, info )
call zungrq ( m, n, k, a, lda, tau, work, lwork, info )
```

Discussion

The routine generates an m -by- n complex matrix Q with orthonormal rows, which is defined as the last m rows of a product of k elementary reflectors H_i of order n : $Q = H_1^H H_2^H \cdot \cdot \cdot H_k^H$ as returned by the routines `sgerqf/dgerqf`. Use this routine after a call to `sgerqf/dgerqf`.

Input Parameters

m **INTEGER.** The number of rows of the matrix Q ($m \geq 0$).

n **INTEGER.** The number of columns of the matrix Q ($n \geq m$).

k **INTEGER.** The number of elementary reflectors whose product defines the matrix Q ($m \geq k \geq 0$).

a , tau , $work$ **REAL** for `cungrq`
DOUBLE PRECISION for `zungrq`
 Arrays: `a(lda,*)`, `tau(*)`, `work(lwork)`.

On entry, the $(m - k + i)$ th row of a must contain the vector which defines the elementary reflector H_i , for $i = 1, 2, \dots, k$, as returned by `sgerqf/dgerqf` in the last k rows of its array argument a ;
 $tau(i)$ must contain the scalar factor of the elementary reflector H_i , as returned by `sgerqf/dgerqf`;

The second dimension of a must be at least $\max(1, n)$.
 The dimension of tau must be at least $\max(1, k)$.
 $work(lwork)$ is a workspace array.

lda INTEGER. The first dimension of *a*; at least $\max(1, m)$.
lwork INTEGER. The size of the *work* array; at least $\max(1, m)$.
See *Application notes* for the suggested value of *lwork*.

Output Parameters

a Overwritten by the *m*-by-*n* matrix *Q*.
work(1) If *info* = 0, on exit *work(1)* contains the minimum
value of *lwork* required for optimum performance. Use
this *lwork* for subsequent runs.
info INTEGER.
If *info* = 0, the execution is successful.
If *info* = *-i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = m * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

The real counterpart of this routine is [?orgrrq](#).

?ormrq

Multiplies a real matrix by the orthogonal matrix Q of the RQ factorization formed by [?gerqf](#).

```
call sormrq ( side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info )
call dormrq ( side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info )
```

Discussion

The routine multiplies a real m -by- n matrix C by Q or Q^T , where Q is the real orthogonal matrix defined as a product of k elementary reflectors H_i : $Q = H_1 H_2 \cdot \cdot \cdot H_k$ as returned by the RQ factorization routine [sgerqf/dgerqf](#).

Depending on the parameters *side* and *trans*, the routine can form one of the matrix products QC , Q^TC , CQ , or CQ^T (overwriting the result over C).

Input Parameters

<i>side</i>	CHARACTER*1. Must be either 'L' or 'R'. If <i>side</i> = 'L', Q or Q^T is applied to C from the left. If <i>side</i> = 'R', Q or Q^T is applied to C from the right.
<i>trans</i>	CHARACTER*1. Must be either 'N' or 'T'. If <i>trans</i> = 'N', the routine multiplies C by Q . If <i>trans</i> = 'T', the routine multiplies C by Q^T .
<i>m</i>	INTEGER. The number of rows in the matrix C ($m \geq 0$).
<i>n</i>	INTEGER. The number of columns in C ($n \geq 0$).
<i>k</i>	INTEGER. The number of elementary reflectors whose product defines the matrix Q . Constraints: $0 \leq k \leq m$, if <i>side</i> = 'L'; $0 \leq k \leq n$, if <i>side</i> = 'R'.

a, tau, c, work REAL for *sormrq*
 DOUBLE PRECISION for *dormrq*.
 Arrays: *a(lda, *)*, *tau(*)*, *c(ldc, *)*,
work(lwork).

On entry, the *i*th row of *a* must contain the vector which defines the elementary reflector H_i , for $i = 1, 2, \dots, k$, as returned by *sgerqf/dgerqf* in the last *k* rows of its array argument *a*.
 The second dimension of *a* must be at least $\max(1, m)$ if *side* = 'L', and at least $\max(1, n)$ if *side* = 'R'.
tau(i) must contain the scalar factor of the elementary reflector H_i , as returned by *sgerqf/dgerqf*.
 The dimension of *tau* must be at least $\max(1, k)$.
*c(ldc, *)* contains the *m*-by-*n* matrix *C*.
 The second dimension of *c* must be at least $\max(1, n)$.
work(lwork) is a workspace array.

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, k)$.

ldc INTEGER. The first dimension of *c*; $ldc \geq \max(1, m)$.

lwork INTEGER. The size of the *work* array. Constraints:
 $lwork \geq \max(1, n)$ if *side* = 'L';
 $lwork \geq \max(1, m)$ if *side* = 'R'.
 See *Application notes* for the suggested value of *lwork*.

Output Parameters

c Overwritten by the product QC , $Q^T C$, CQ , or CQ^T (as specified by *side* and *trans*).

work(1) If *info* = 0, on exit *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = n * blocksize$ (if $side = 'L'$) or $lwork = m * blocksize$ (if $side = 'R'$) where $blocksize$ is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of $lwork$ for the first run. On exit, examine $work(1)$ and use this value for subsequent runs.

The complex counterpart of this routine is [?unmrq](#).

?unmrq

Multiplies a complex matrix by the unitary matrix Q of the RQ factorization formed by [?gerqf](#).

```
call cunmrq ( side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info )
call zunmrq ( side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info )
```

Discussion

The routine multiplies a complex m -by- n matrix C by Q or Q^H , where Q is the complex unitary matrix defined as a product of k elementary reflectors H_i : $Q = H_1^H H_2^H \cdot \cdot \cdot H_k^H$ as returned by the RQ factorization routine [cgerqf/zgerqf](#).

Depending on the parameters *side* and *trans*, the routine can form one of the matrix products QC , $Q^H C$, CQ , or CQ^H (overwriting the result over C).

Input Parameters

side CHARACTER*1. Must be either 'L' or 'R'.
If *side* = 'L', Q or Q^H is applied to C from the left.
If *side* = 'R', Q or Q^H is applied to C from the right.

trans CHARACTER*1. Must be either 'N' or 'C'.
If *trans* = 'N', the routine multiplies C by Q .
If *trans* = 'C', the routine multiplies C by Q^H .

m INTEGER. The number of rows in the matrix C ($m \geq 0$).

n INTEGER. The number of columns in C ($n \geq 0$).

k INTEGER. The number of elementary reflectors whose product defines the matrix Q . Constraints:
 $0 \leq k \leq m$, if *side* = 'L';
 $0 \leq k \leq n$, if *side* = 'R'.

a, tau, c, work COMPLEX for `cunmrq`
 DOUBLE COMPLEX for `zunmrq`.
 Arrays: *a(lda, *)*, *tau(*)*, *c(ldc, *)*,
work(lwork).

On entry, the *i*th row of *a* must contain the vector which defines the elementary reflector H_i , for $i = 1, 2, \dots, k$, as returned by `cgerqf/zgerqf` in the last *k* rows of its array argument *a*.
 The second dimension of *a* must be at least $\max(1, m)$ if *side* = 'L', and at least $\max(1, n)$ if *side* = 'R'.
tau(i) must contain the scalar factor of the elementary reflector H_i , as returned by `cgerqf/zgerqf`.
 The dimension of *tau* must be at least $\max(1, k)$.
*c(ldc, *)* contains the *m*-by-*n* matrix *C*.
 The second dimension of *c* must be at least $\max(1, n)$.
work(lwork) is a workspace array.

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, k)$.

ldc INTEGER. The first dimension of *c*; $ldc \geq \max(1, m)$.

lwork INTEGER. The size of the *work* array. Constraints:
 $lwork \geq \max(1, n)$ if *side* = 'L';
 $lwork \geq \max(1, m)$ if *side* = 'R'.
 See *Application notes* for the suggested value of *lwork*.

Output Parameters

c Overwritten by the product QC , $Q^H C$, CQ , or CQ^H (as specified by *side* and *trans*).

work(1) If *info* = 0, on exit *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = n * blocksize$ (if $side = 'L'$) or $lwork = m * blocksize$ (if $side = 'R'$) where $blocksize$ is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of $lwork$ for the first run. On exit, examine $work(1)$ and use this value for subsequent runs.

The real counterpart of this routine is [?ormrq](#).

?tzzrf

Reduces the upper trapezoidal matrix A to upper triangular form.

```
call stzzrf ( m, n, a, lda, tau, work, lwork, info )
call dtzzrf ( m, n, a, lda, tau, work, lwork, info )
call ctzzrf ( m, n, a, lda, tau, work, lwork, info )
call ztzzrf ( m, n, a, lda, tau, work, lwork, info )
```

Discussion

This routine reduces the m -by- n ($m \leq n$) real/complex upper trapezoidal matrix A to upper triangular form by means of orthogonal/unitary transformations. The upper trapezoidal matrix A is factored as

$$A = (R \ 0) * Z,$$

where Z is an n -by- n orthogonal/unitary matrix and R is an m -by- m upper triangular matrix.

Input Parameters

m	INTEGER. The number of rows in the matrix A ($m \geq 0$).
n	INTEGER. The number of columns in A ($n \geq m$).
$a, work$	REAL for stzzrf DOUBLE PRECISION for dtzzrf COMPLEX for ctzzrf DOUBLE COMPLEX for ztzzrf. Arrays: $a(lda, *)$, $work(lwork)$. The leading m -by- n upper trapezoidal part of the array a contains the matrix A to be factorized. The second dimension of a must be at least $\max(1, n)$. $work$ is a workspace array.
lda	INTEGER. The first dimension of a ; at least $\max(1, m)$.
$lwork$	INTEGER. The size of the $work$ array;

$lwork \geq \max(1, m)$.

See *Application notes* for the suggested value of $lwork$.

Output Parameters

- a Overwritten on exit by the factorization data as follows:
the leading m -by- m upper triangular part of a contains the upper triangular matrix R , and elements $m + 1$ to n of the first m rows of a , with the array tau , represent the orthogonal matrix Z as a product of m elementary reflectors.
- tau REAL for `stzrzf`
DOUBLE PRECISION for `dtzrzf`
COMPLEX for `ctzrzf`
DOUBLE COMPLEX for `ztzrzf`.
Array, DIMENSION at least $\max(1, m)$.
Contains scalar factors of the elementary reflectors for the matrix Z .
- $work(1)$ If $info = 0$, on exit $work(1)$ contains the minimum value of $lwork$ required for optimum performance. Use this $lwork$ for subsequent runs.
- $info$ INTEGER.
If $info = 0$, the execution is successful.
If $info = -i$, the i th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = m * blocksize$, where $blocksize$ is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*.

If you are in doubt how much workspace to supply, use a generous value of $lwork$ for the first run. On exit, examine $work(1)$ and use this value for subsequent runs.

Related routines include:

- [?ormrz](#) to apply matrix Q (for real matrices);
[?unmrz](#) to apply matrix Q (for complex matrices).

?ormrz

Multiplies a real matrix by the orthogonal matrix defined from the factorization formed by [?tzzrf](#).

```
call sormrz ( side,trans,m,n,k,l,a,lda,tau,c,ldc,work,lwork,info )
call dormrz ( side,trans,m,n,k,l,a,lda,tau,c,ldc,work,lwork,info )
```

Discussion

The routine multiplies a real m -by- n matrix C by Q or Q^T , where Q is the real orthogonal matrix defined as a product of k elementary reflectors H_i : $Q = H_1 H_2 \cdot \cdot \cdot H_k$ as returned by the factorization routine [stzzrf/dtzzrf](#).

Depending on the parameters *side* and *trans*, the routine can form one of the matrix products QC , $Q^T C$, CQ , or CQ^T (overwriting the result over C).

The matrix Q is of order m if *side* = 'L' and of order n if *side* = 'R'.

Input Parameters

<i>side</i>	CHARACTER*1. Must be either 'L' or 'R'. If <i>side</i> = 'L', Q or Q^T is applied to C from the left. If <i>side</i> = 'R', Q or Q^T is applied to C from the right.
<i>trans</i>	CHARACTER*1. Must be either 'N' or 'T'. If <i>trans</i> = 'N', the routine multiplies C by Q . If <i>trans</i> = 'T', the routine multiplies C by Q^T .
<i>m</i>	INTEGER. The number of rows in the matrix C ($m \geq 0$).
<i>n</i>	INTEGER. The number of columns in C ($n \geq 0$).
<i>k</i>	INTEGER. The number of elementary reflectors whose product defines the matrix Q . Constraints: $0 \leq k \leq m$, if <i>side</i> = 'L'; $0 \leq k \leq n$, if <i>side</i> = 'R'.
<i>l</i>	INTEGER.

The number of columns of the matrix A containing the meaningful part of the Householder reflectors.

Constraints:

$0 \leq l \leq m$, if *side* = 'L';

$0 \leq l \leq n$, if *side* = 'R'.

a, tau, c, work REAL for *sormrz*
DOUBLE PRECISION for *dormrz*.
Arrays: *a(lda,*)*, *tau(*)*, *c(ldc,*)*,
work(lwork).

On entry, the i th row of *a* must contain the vector which defines the elementary reflector H_i , for $i = 1, 2, \dots, k$, as returned by *stzrzf/dtzrzf* in the last k rows of its array argument *a*.

The second dimension of *a* must be at least $\max(1, m)$ if *side* = 'L', and at least $\max(1, n)$ if *side* = 'R'.

tau(i) must contain the scalar factor of the elementary reflector H_i , as returned by *stzrzf/dtzrzf*.

The dimension of *tau* must be at least $\max(1, k)$.

c(ldc,)* contains the m -by- n matrix C .

The second dimension of *c* must be at least $\max(1, n)$

work(lwork) is a workspace array.

lda INTEGER. The first dimension of *a*; $lda \geq \max(1, k)$.

ldc INTEGER. The first dimension of *c*; $ldc \geq \max(1, m)$.

lwork INTEGER. The size of the *work* array. Constraints:
 $lwork \geq \max(1, n)$ if *side* = 'L';
 $lwork \geq \max(1, m)$ if *side* = 'R'.

See *Application notes* for the suggested value of *lwork*.

Output Parameters

c Overwritten by the product QC , Q^TC , CQ , or CQ^T
(as specified by *side* and *trans*).

work(1) If *info* = 0, on exit *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using *lwork* = *n***blocksize* (if *side* = 'L') or *lwork* = *m***blocksize* (if *side* = 'R') where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

The complex counterpart of this routine is [?unmrz](#).

?unmrz

Multiplies a complex matrix by the unitary matrix defined from the factorization formed by `?tzrzf`.

```
call cummrz ( side,trans,m,n,k,l,a,lda,tau,c,ldc,work,lwork,info )
call zummrz ( side,trans,m,n,k,l,a,lda,tau,c,ldc,work,lwork,info )
```

Discussion

The routine multiplies a complex m -by- n matrix C by Q or Q^H , where Q is the unitary matrix defined as a product of k elementary reflectors H_i :

$Q = H_1^H H_2^H \cdots H_k^H$ as returned by the factorization routine [ctzrzf/ztzrzf](#).

Depending on the parameters `side` and `trans`, the routine can form one of the matrix products QC , $Q^H C$, CQ , or CQ^H (overwriting the result over C).

The matrix Q is of order m if `side = 'L'` and of order n if `side = 'R'`.

Input Parameters

<code>side</code>	CHARACTER*1. Must be either 'L' or 'R'. If <code>side = 'L'</code> , Q or Q^H is applied to C from the left. If <code>side = 'R'</code> , Q or Q^H is applied to C from the right.
<code>trans</code>	CHARACTER*1. Must be either 'N' or 'C'. If <code>trans = 'N'</code> , the routine multiplies C by Q . If <code>trans = 'C'</code> , the routine multiplies C by Q^H .
<code>m</code>	INTEGER. The number of rows in the matrix C ($m \geq 0$).
<code>n</code>	INTEGER. The number of columns in C ($n \geq 0$).
<code>k</code>	INTEGER. The number of elementary reflectors whose product defines the matrix Q . Constraints: $0 \leq k \leq m$, if <code>side = 'L'</code> ; $0 \leq k \leq n$, if <code>side = 'R'</code> .
<code>l</code>	INTEGER.

The number of columns of the matrix A containing the meaningful part of the Householder reflectors.

Constraints:

$0 \leq l \leq m$, if *side* = 'L';

$0 \leq l \leq n$, if *side* = 'R'.

a, tau, c, work **COMPLEX** for `cunmrz`
 DOUBLE COMPLEX for `zunmrz`.
 Arrays: *a(lda,*)*, *tau(*)*, *c(ldc,*)*,
work(lwork).

On entry, the *i*th row of *a* must contain the vector which defines the elementary reflector H_i , for $i = 1, 2, \dots, k$, as returned by `ctzrzf/ztzrzf` in the last *k* rows of its array argument *a*.

The second dimension of *a* must be at least $\max(1, m)$ if *side* = 'L', and at least $\max(1, n)$ if *side* = 'R'.

tau(*i*) must contain the scalar factor of the elementary reflector H_i , as returned by `ctzrzf/ztzrzf`.

The dimension of *tau* must be at least $\max(1, k)$.

c(ldc,)* contains the *m*-by-*n* matrix *C*.

The second dimension of *c* must be at least $\max(1, n)$

work(lwork) is a workspace array.

lda **INTEGER**. The first dimension of *a*; $lda \geq \max(1, k)$.

ldc **INTEGER**. The first dimension of *c*; $ldc \geq \max(1, m)$.

lwork **INTEGER**. The size of the *work* array. Constraints:

$lwork \geq \max(1, n)$ if *side* = 'L';

$lwork \geq \max(1, m)$ if *side* = 'R'.

See *Application notes* for the suggested value of *lwork*.

Output Parameters

c Overwritten by the product QC , $Q^H C$, CQ , or CQ^H (as specified by *side* and *trans*).

work(1) If *info* = 0, on exit *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using *lwork* = *n***blocksize* (if *side* = 'L') or *lwork* = *m***blocksize* (if *side* = 'R') where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work*(1) and use this value for subsequent runs.

The real counterpart of this routine is [?ormrz](#).

?ggqrf

Computes the generalized QR factorization of two matrices.

```
call sggqrf (n, m, p, a, lda, taua, b, ldb, taub, work, lwork, info)
call dggqrf (n, m, p, a, lda, taua, b, ldb, taub, work, lwork, info)
call cggqrf (n, m, p, a, lda, taua, b, ldb, taub, work, lwork, info)
call zggqrf (n, m, p, a, lda, taua, b, ldb, taub, work, lwork, info)
```

Discussion

The routine forms the generalized QR factorization of an n -by- m matrix A and an n -by- p matrix B as $A = QR$, $B = QTZ$, where Q is an n -by- n orthogonal/unitary matrix, Z is a p -by- p orthogonal/unitary matrix, and R and T assume one of the forms:

$$R = \begin{matrix} & & m \\ & m & \\ n - m & & \end{matrix} \begin{pmatrix} R_{11} \\ 0 \end{pmatrix}, \text{ if } n \geq m$$

or

$$R = \begin{matrix} & n & m - n \\ n & (R_{11} & R_{12}) \end{matrix}, \text{ if } n < m,$$

where R_{11} is upper triangular, and

$$T = \begin{matrix} & p - n & n \\ n & (0 & T_{12}) \end{matrix}, \text{ if } n \leq p, \text{ or}$$

$$T = \begin{matrix} & & p \\ n - p & & \\ & p & \end{matrix} \begin{pmatrix} T_{11} \\ T_{21} \end{pmatrix}, \text{ if } n > p$$

where T_{12} or T_{21} is a p -by- p upper triangular matrix.

In particular, if B is square and nonsingular, the GQR factorization of A and B implicitly gives the QR factorization of $B^{-1}A$ as:

$$B^{-1}A = Z^H(T^{-1}R)$$

Input Parameters

- n **INTEGER**. The number of rows of the matrices A and B ($n \geq 0$).
- m **INTEGER**. The number of columns in A ($m \geq 0$).
- p **INTEGER**. The number of columns in B ($p \geq 0$).
- $a, b, work$ **REAL** for `sggqrf`
DOUBLE PRECISION for `dggqrf`
COMPLEX for `cggqrf`
DOUBLE COMPLEX for `zggqrf`.
- Arrays:
 $a(lda, *)$ contains the matrix A .
 The second dimension of a must be at least $\max(1, m)$.
 $b(l db, *)$ contains the matrix B .
 The second dimension of b must be at least $\max(1, p)$.
 $work(lwork)$ is a workspace array.
- lda **INTEGER**. The first dimension of a ; at least $\max(1, n)$.
- ldb **INTEGER**. The first dimension of b ; at least $\max(1, n)$.
- $lwork$ **INTEGER**. The size of the $work$ array; must be at least $\max(1, n, m, p)$
 See *Application notes* for the suggested value of $lwork$.

Output Parameters

- a, b Overwritten by the factorization data as follows:
 on exit, the elements on and above the diagonal of the array a contain the $\min(n, m)$ -by- m upper trapezoidal matrix R (R is upper triangular if $n \geq m$); the elements below the diagonal, with the array $taua$, represent the orthogonal/unitary matrix Q as a product of $\min(n, m)$ elementary reflectors ;

- if $n \leq p$, the upper triangle of the subarray $b(1:n, p-n+1:p)$ contains the n -by- n upper triangular matrix T ;
- if $n > p$, the elements on and above the $(n-p)$ th subdiagonal contain the n -by- p upper trapezoidal matrix T ; the remaining elements, with the array $taub$, represent the orthogonal/unitary matrix Z as a product of elementary reflectors.
- taua*, *taub* REAL for `sggqrf`
DOUBLE PRECISION for `dggqrf`
COMPLEX for `cggqrf`
DOUBLE COMPLEX for `zggqrf`.
Arrays, DIMENSION at least $\max(1, \min(n, m))$ for *taua* and at least $\max(1, \min(n, p))$ for *taub*.
The array *taua* contains the scalar factors of the elementary reflectors which represent the orthogonal/unitary matrix Q .
The array *taub* contains the scalar factors of the elementary reflectors which represent the orthogonal/unitary matrix Z .
- work(1)* If *info* = 0, on exit *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.
- info* INTEGER.
If *info* = 0, the execution is successful.
If *info* = $-i$, the i th parameter had an illegal value.

Application Notes

For better performance, try using

$$lwork \geq \max(n, m, p) * \max(nb1, nb2, nb3),$$

where *nb1* is the optimal blocksize for the QR factorization of an n -by- m matrix, *nb2* is the optimal blocksize for the RQ factorization of an n -by- p matrix, and *nb3* is the optimal blocksize for a call of `?ormqr`/`?unmqr`.

?ggrqf

Computes the generalized RQ factorization of two matrices.

```
call sggrqf (m, p, n, a, lda, taua, b, ldb, taub, work, lwork, info)
call dggrqf (m, p, n, a, lda, taua, b, ldb, taub, work, lwork, info)
call cggrqf (m, p, n, a, lda, taua, b, ldb, taub, work, lwork, info)
call zggrqf (m, p, n, a, lda, taua, b, ldb, taub, work, lwork, info)
```

Discussion

The routine forms the generalized RQ factorization of an m -by- n matrix A and an p -by- n matrix B as $A = R Q$, $B = Z T Q$, where Q is an n -by- n orthogonal/unitary matrix, Z is a p -by- p orthogonal/unitary matrix, and R and T assume one of the forms:

$$R = \begin{matrix} n-m & m \\ m & \end{matrix} \begin{pmatrix} 0 & R_{12} \end{pmatrix}, \text{ if } m \leq n,$$

or

$$R = \begin{matrix} m-n & n \\ m-n & \end{matrix} \begin{pmatrix} R_{11} \\ R_{21} \end{pmatrix}, \text{ if } m > n$$

where R_{11} or R_{21} is upper triangular, and

$$T = \begin{matrix} n \\ p-n & \end{matrix} \begin{pmatrix} T_{11} \\ 0 \end{pmatrix}, \text{ if } p \geq n$$

or

$$T = \begin{matrix} p & n-p \\ p & \end{matrix} \begin{pmatrix} T_{11} & T_{12} \end{pmatrix}, \text{ if } p < n,$$

where T_{11} is upper triangular.

In particular, if B is square and nonsingular, the GRQ factorization of A and B implicitly gives the RQ factorization of AB^{-1} as:

$$AB^{-1} = (R T^{-1}) Z^H$$

Input Parameters

<i>m</i>	INTEGER. The number of rows of the matrix A ($m \geq 0$).
<i>p</i>	INTEGER. The number of rows in B ($p \geq 0$).
<i>n</i>	INTEGER. The number of columns of the matrices A and B ($n \geq 0$).
<i>a</i> , <i>b</i> , <i>work</i>	REAL for <code>sggrqf</code> DOUBLE PRECISION for <code>dggrqf</code> COMPLEX for <code>cggrqf</code> DOUBLE COMPLEX for <code>zggrqf</code> . Arrays: <i>a</i> (<i>lda</i> ,*) contains the m -by- n matrix A . The second dimension of <i>a</i> must be at least $\max(1, n)$. <i>b</i> (<i>ldb</i> ,*) contains the p -by- n matrix B . The second dimension of <i>b</i> must be at least $\max(1, n)$. <i>work</i> (<i>lwork</i>) is a workspace array.
<i>lda</i>	INTEGER. The first dimension of <i>a</i> ; at least $\max(1, m)$.
<i>ldb</i>	INTEGER. The first dimension of <i>b</i> ; at least $\max(1, p)$.
<i>lwork</i>	INTEGER. The size of the <i>work</i> array; must be at least $\max(1, n, m, p)$ See <i>Application notes</i> for the suggested value of <i>lwork</i> .

Output Parameters

<i>a</i> , <i>b</i>	Overwritten by the factorization data as follows: on exit, if $m \leq n$, the upper triangle of the subarray <i>a</i> (1: <i>m</i> , <i>n</i> - <i>m</i> +1: <i>n</i>) contains the m -by- m upper triangular matrix R ; if $m > n$, the elements on and above the (m - n)th subdiagonal contain the m -by- n upper trapezoidal
---------------------	---

matrix R ; the remaining elements, with the array `taua`, represent the orthogonal/unitary matrix Q as a product of elementary reflectors; the elements on and above the diagonal of the array `b` contain the $\min(p,n)$ -by- n upper trapezoidal matrix T (T is upper triangular if $p \geq n$); the elements below the diagonal, with the array `taub`, represent the orthogonal/unitary matrix Z as a product of elementary reflectors.

`taua, taub` REAL for `sggrqf`
 DOUBLE PRECISION for `dggrqf`
 COMPLEX for `cggrqf`
 DOUBLE COMPLEX for `zggrqf`.
 Arrays, DIMENSION at least $\max(1, \min(m, n))$ for `taua` and at least $\max(1, \min(p, n))$ for `taub`.
 The array `taua` contains the scalar factors of the elementary reflectors which represent the orthogonal/unitary matrix Q .
 The array `taub` contains the scalar factors of the elementary reflectors which represent the orthogonal/unitary matrix Z .

`work(1)` If `info` = 0, on exit `work(1)` contains the minimum value of `lwork` required for optimum performance. Use this `lwork` for subsequent runs.

`info` INTEGER.
 If `info` = 0, the execution is successful.
 If `info` = $-i$, the i th parameter had an illegal value.

Application Notes

For better performance, try using

$$lwork \geq \max(n, m, p) * \max(nb1, nb2, nb3),$$

where $nb1$ is the optimal blocksize for the RQ factorization of an m -by- n matrix, $nb2$ is the optimal blocksize for the QR factorization of an p -by- n matrix, and $nb3$ is the optimal blocksize for a call of `?ormrq/?unmrq`.

If you are in doubt how much workspace to supply, use a generous value of `lwork` for the first run. On exit, examine `work(1)` and use this value for subsequent runs.

Singular Value Decomposition

This section describes LAPACK routines for computing the *singular value decomposition* (SVD) of a general m by n matrix A :

$$A = U\Sigma V^H.$$

In this decomposition, U and V are unitary (for complex A) or orthogonal (for real A); Σ is an m by n diagonal matrix with real diagonal elements σ_i :

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{\min(m, n)} \geq 0.$$

The diagonal elements σ_i are *singular values* of A . The first $\min(m, n)$ columns of the matrices U and V are, respectively, *left* and *right singular vectors* of A . The singular values and singular vectors satisfy

$$Av_i = \sigma_i u_i \quad \text{and} \quad A^H u_i = \sigma_i v_i$$

where u_i and v_i are the i th columns of U and V , respectively.

To find the SVD of a general matrix A , call the LAPACK routine `?gebrd` or `?gbbbrd` for reducing A to a bidiagonal matrix B by a unitary (orthogonal) transformation: $A = QBP^H$. Then call `?bdsqr`, which forms the SVD of a bidiagonal matrix: $B = U_1 \Sigma V_1^H$.

Thus, the sought-for SVD of A is given by $A = U\Sigma V^H = (QU_1) \Sigma (V_1^H P^H)$.

Table 5-2 Computational Routines for Singular Value Decomposition (SVD)

Operation	Real matrices	Complex matrices
Reduce A to a bidiagonal matrix B : $A = QBP^H$ (full storage)	?gebrd	?gebrd
Reduce A to a bidiagonal matrix B : $A = QBP^H$ (band storage)	?gbbbrd	?gbbbrd
Generate the orthogonal (unitary) matrix Q or P	?orgbr	?ungbr
Apply the orthogonal (unitary) matrix Q or P	?ormbr	?unmbr
Form singular value decomposition of the bidiagonal matrix B : $B = U\Sigma V^H$?bdsqr ?bdsdc	?bdsqr

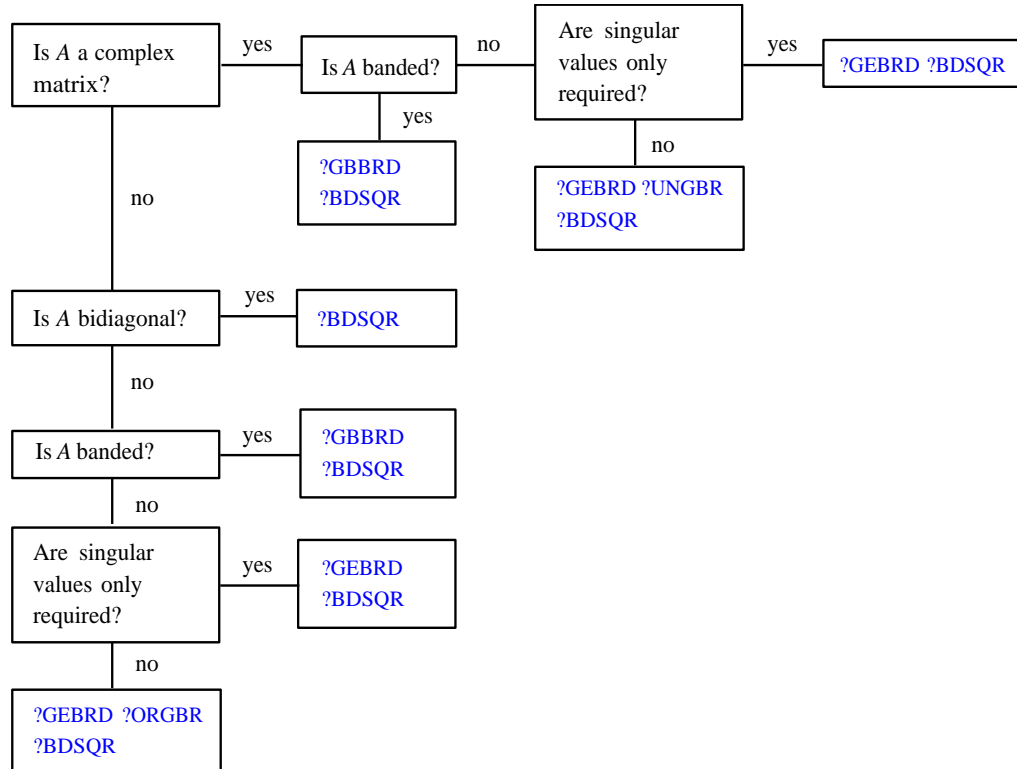
Figure 5-1 Decision Tree: Singular Value Decomposition

Figure 5-1 presents a decision tree that helps you choose the right sequence of routines for SVD, depending on whether you need singular values only or singular vectors as well, whether A is real or complex, and so on.

You can use the SVD to find a minimum-norm solution to a (possibly) rank-deficient least-squares problem of minimizing $\|Ax - b\|_2$. The effective rank k of the matrix A can be determined as the number of singular values which exceed a suitable threshold. The minimum-norm solution is

$$x = V_k(\Sigma_k)^{-1}c$$

where Σ_k is the leading k by k submatrix of Σ , the matrix V_k consists of the first k columns of $V = PV_1$, and the vector c consists of the first k elements of $U^H b = U_1^H Q^H b$.

?gebrd

Reduces a general matrix to bidiagonal form.

```
call sgebrd ( m, n, a, lda, d, e, tauq, taup, work, lwork, info )
call dgebrd ( m, n, a, lda, d, e, tauq, taup, work, lwork, info )
call cgebrd ( m, n, a, lda, d, e, tauq, taup, work, lwork, info )
call zgebrd ( m, n, a, lda, d, e, tauq, taup, work, lwork, info )
```

Discussion

The routine reduces a general m by n matrix A to a bidiagonal matrix B by an orthogonal (unitary) transformation.

If $m \geq n$, the reduction is given by

$$A = QBP^H = Q \begin{pmatrix} B_1 \\ 0 \end{pmatrix} P^H = Q_1 B_1 P^H,$$

where B_1 is an n by n upper diagonal matrix, Q and P are orthogonal or, for a complex A , unitary matrices; Q_1 consists of the first n columns of Q .

If $m < n$, the reduction is given by

$$A = QBP^H = Q(B_1 0)P^H = Q_1 B_1 P_1^H,$$

where B_1 is an m by m lower diagonal matrix, Q and P are orthogonal or, for a complex A , unitary matrices; P_1 consists of the first m rows of P .

The routine does not form the matrices Q and P explicitly, but represents them as products of elementary reflectors. Routines are provided to work with the matrices Q and P in this representation:

If the matrix A is real,

- to compute Q and P explicitly, call [?orgbr](#).
- to multiply a general matrix by Q or P , call [?ormbr](#).

If the matrix A is complex,

- to compute Q and P explicitly, call [?ungbr](#).
- to multiply a general matrix by Q or P , call [?unmbr](#).

Input Parameters

- m* **INTEGER**. The number of rows in the matrix *A* ($m \geq 0$).
- n* **INTEGER**. The number of columns in *A* ($n \geq 0$).
- a*, *work* **REAL** for **sgebrd**
DOUBLE PRECISION for **dgebrd**
COMPLEX for **cgebrd**
DOUBLE COMPLEX for **zgebrd**.
- Arrays:
a(*lda*,*) contains the matrix *A*.
The second dimension of *a* must be at least $\max(1, n)$.
work(*lwork*) is a workspace array.
- lda* **INTEGER**. The first dimension of *a*; at least $\max(1, m)$.
- lwork* **INTEGER**. The dimension of *work*; at least $\max(1, m, n)$.
See *Application notes* for the suggested value of *lwork*.

Output Parameters

- a* If $m \geq n$, the diagonal and first super-diagonal of *a* are overwritten by the upper bidiagonal matrix *B*. Elements below the diagonal are overwritten by details of *Q*, and the remaining elements are overwritten by details of *P*.
If $m < n$, the diagonal and first sub-diagonal of *a* are overwritten by the lower bidiagonal matrix *B*. Elements above the diagonal are overwritten by details of *P*, and the remaining elements are overwritten by details of *Q*.
- d* **REAL** for single-precision flavors
DOUBLE PRECISION for double-precision flavors.
Array, **DIMENSION** at least $\max(1, \min(m, n))$.
Contains the diagonal elements of *B*.
- e* **REAL** for single-precision flavors
DOUBLE PRECISION for double-precision flavors.
Array, **DIMENSION** at least $\max(1, \min(m, n) - 1)$.
Contains the off-diagonal elements of *B*.

<i>tauq, taup</i>	REAL for <i>sgebrd</i> DOUBLE PRECISION for <i>dgebrd</i> COMPLEX for <i>cgebrd</i> DOUBLE COMPLEX for <i>zgebrd</i> . Arrays, DIMENSION at least $\max(1, \min(m, n))$. Contain further details of the matrices <i>Q</i> and <i>P</i> .
<i>work(1)</i>	If <i>info</i> = 0, on exit <i>work(1)</i> contains the minimum value of <i>lwork</i> required for optimum performance. Use this <i>lwork</i> for subsequent runs.
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = (m + n) * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*.

If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

The computed matrices *Q*, *B*, and *P* satisfy $QBP^H = A + E$, where $\|E\|_2 = c(n)\epsilon \|A\|_2$, *c(n)* is a modestly increasing function of *n*, and ϵ is the machine precision.

The approximate number of floating-point operations for real flavors is
 $(4/3) * n^2 * (3 * m - n)$ for $m \geq n$,
 $(4/3) * m^2 * (3 * n - m)$ for $m < n$.

The number of operations for complex flavors is four times greater.

If *n* is much less than *m*, it can be more efficient to first form the *QR* factorization of *A* by calling [?geqrf](#) and then reduce the factor *R* to bidiagonal form. This requires approximately $2 * n^2 * (m + n)$ floating-point operations.

If *m* is much less than *n*, it can be more efficient to first form the *LQ* factorization of *A* by calling [?gelqf](#) and then reduce the factor *L* to bidiagonal form. This requires approximately $2 * m^2 * (m + n)$ floating-point operations.

?gbbbrd

Reduces a general band matrix to
bidiagonal form.

```
call sgbbrd ( vect, m, n, ncc, kl, ku, ab, ldab, d, e, q, ldq, pt,
             ldpt, c, ldc, work, info )
call dgbbrd ( vect, m, n, ncc, kl, ku, ab, ldab, d, e, q, ldq, pt,
             ldpt, c, ldc, work, info )
call cgbbrd ( vect, m, n, ncc, kl, ku, ab, ldab, d, e, q, ldq, pt,
             ldpt, c, ldc, work, rwork, info )
call zgbbrd ( vect, m, n, ncc, kl, ku, ab, ldab, d, e, q, ldq, pt,
             ldpt, c, ldc, work, rwork, info )
```

Discussion

This routine reduces an m by n band matrix A to upper bidiagonal matrix B : $A = QBPH$. Here the matrices Q and P are orthogonal (for real A) or unitary (for complex A). They are determined as products of Givens rotation matrices, and may be formed explicitly by the routine if required. The routine can also update a matrix C as follows: $C = Q^H C$.

Input Parameters

vect CHARACTER*1. Must be 'N' or 'Q' or 'P' or 'B'.
If **vect** = 'N', neither Q nor P^H is generated.
If **vect** = 'Q', the routine generates the matrix Q .
If **vect** = 'P', the routine generates the matrix P^H .
If **vect** = 'B', the routine generates both Q and P^H .

m INTEGER. The number of rows in the matrix A ($m \geq 0$).

n INTEGER. The number of columns in A ($n \geq 0$).

ncc INTEGER. The number of columns in C ($ncc \geq 0$).

kl INTEGER. The number of sub-diagonals within the band of A ($kl \geq 0$).

ku INTEGER. The number of super-diagonals within the band of A ($ku \geq 0$).

<i>ab, c, work</i>	<p>REAL for <code>sgbbrd</code> DOUBLE PRECISION for <code>dgbbrd</code> COMPLEX for <code>cgbbrd</code> DOUBLE COMPLEX for <code>zgbbrd</code>.</p> <p>Arrays: <i>ab</i>(<i>ldab</i>, *) contains the matrix <i>A</i> in band storage (see Matrix Storage Schemes). The second dimension of <i>a</i> must be at least $\max(1, n)$. <i>c</i>(<i>ldc</i>, *) contains an <i>m</i> by <i>ncc</i> matrix <i>C</i>. If <i>ncc</i> = 0, the array <i>c</i> is not referenced. The second dimension of <i>c</i> must be at least $\max(1, ncc)$. <i>work</i>(*) is a workspace array. The dimension of <i>work</i> must be at least $2 * \max(m, n)$ for real flavors, or $\max(m, n)$ for complex flavors.</p>
<i>ldab</i>	<p>INTEGER. The first dimension of the array <i>ab</i> (<i>ldab</i> $\geq k_l + k_u + 1$).</p>
<i>ldq</i>	<p>INTEGER. The first dimension of the output array <i>q</i>. <i>ldq</i> $\geq \max(1, m)$ if <i>vect</i> = 'Q' or 'B', <i>ldq</i> ≥ 1 otherwise.</p>
<i>ldpt</i>	<p>INTEGER. The first dimension of the output array <i>pt</i>. <i>ldpt</i> $\geq \max(1, n)$ if <i>vect</i> = 'P' or 'B', <i>ldpt</i> ≥ 1 otherwise.</p>
<i>ldc</i>	<p>INTEGER. The first dimension of the array <i>c</i>. <i>ldc</i> $\geq \max(1, m)$ if <i>ncc</i> > 0; <i>ldc</i> ≥ 1 if <i>ncc</i> = 0.</p>
<i>rwork</i>	<p>REAL for <code>cgbbrd</code> DOUBLE PRECISION for <code>zgbbrd</code>. A workspace array, DIMENSION at least $\max(m, n)$.</p>

Output Parameters

<i>ab</i>	Overwritten by values generated during the reduction.
<i>d</i>	<p>REAL for single-precision flavors DOUBLE PRECISION for double-precision flavors. Array, DIMENSION at least $\max(1, \min(m, n))$. Contains the diagonal elements of the matrix <i>B</i>.</p>

e REAL for single-precision flavors
 DOUBLE PRECISION for double-precision flavors.
 Array, DIMENSION at least $\max(1, \min(m, n) - 1)$.
 Contains the off-diagonal elements of *B*.

q, *pt* REAL for *sgebrd*
 DOUBLE PRECISION for *dgebrd*
 COMPLEX for *cgebrd*
 DOUBLE COMPLEX for *zgebrd*.
 Arrays:
q(*ldq*, *) contains the output *m* by *m* matrix *Q*.
 The second dimension of *q* must be at least $\max(1, m)$.
p(*ldpt*, *) contains the output *n* by *n* matrix *P^H*.
 The second dimension of *pt* must be at least $\max(1, n)$.

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The computed matrices *Q*, *B*, and *P* satisfy $QBP^H = A + E$, where $\|E\|_2 = c(n)\epsilon \|A\|_2$, *c*(*n*) is a modestly increasing function of *n*, and ϵ is the machine precision.

If *m* = *n*, the total number of floating-point operations for real flavors is approximately the sum of:

$$\begin{aligned}
 &6 * n^2 * (kl + ku) && \text{if } vect = 'N' \text{ and } ncc = 0, \\
 &3 * n^2 * ncc * (kl + ku - 1) / (kl + ku) && \text{if } C \text{ is updated, and} \\
 &3 * n^3 * (kl + ku - 1) / (kl + ku) && \text{if either } Q \text{ or } P^H \text{ is generated} \\
 &&& \text{(double this if both).}
 \end{aligned}$$

To estimate the number of operations for complex flavors, use the same formulas with the coefficients 20 and 10 (instead of 6 and 3).

?orgbr

Generates the real orthogonal matrix Q
or P^T determined by ?gebrd.

```
call sorgbr ( vect, m, n, k, a, lda, tau, work, lwork, info )
call dorgbr ( vect, m, n, k, a, lda, tau, work, lwork, info )
```

Discussion

The routine generates the whole or part of the orthogonal matrices Q and P^T formed by the routines `sgebrd/dgebrd` (see [page 5-76](#)). Use this routine after a call to `sgebrd/dgebrd`. All valid combinations of arguments are described in *Input parameters*. In most cases you'll need the following:

To compute the whole m by m matrix Q :

```
call ?orgbr ( 'Q', m, m, n, a ... )
```

(note that the array a must have at least m columns).

To form the n leading columns of Q if $m > n$:

```
call ?orgbr ( 'Q', m, n, n, a ... )
```

To compute the whole n by n matrix P^T :

```
call ?orgbr ( 'P', n, n, m, a ... )
```

(note that the array a must have at least n rows).

To form the m leading rows of P^T if $m < n$:

```
call ?orgbr ( 'P', m, n, m, a ... )
```

Input Parameters

<code>vect</code>	CHARACTER*1. Must be 'Q' or 'P'. If <code>vect = 'Q'</code> , the routine generates the matrix Q . If <code>vect = 'P'</code> , the routine generates the matrix P^T .
<code>m</code>	INTEGER. The number of required rows of Q or P^T .
<code>n</code>	INTEGER. The number of required columns of Q or P^T .
<code>k</code>	INTEGER. One of the dimensions of A in ?gebrd: If <code>vect = 'Q'</code> , the number of columns in A ; If <code>vect = 'P'</code> , the number of rows in A .

Constraints: $m \geq 0$, $n \geq 0$, $k \geq 0$.

For $vect = 'Q'$: $k \leq n \leq m$ if $m > k$, or $m = n$ if $m \leq k$.

For $vect = 'P'$: $k \leq m \leq n$ if $n > k$, or $m = n$ if $n \leq k$.

$a, work$ REAL for `sorgbr`
DOUBLE PRECISION for `dorgbr`.
Arrays:
 $a(lda, *)$ is the array a as returned by `?gebrd`.
The second dimension of a must be at least $\max(1, n)$.
 $work(lwork)$ is a workspace array.

lda INTEGER. The first dimension of a ; at least $\max(1, m)$.

tau REAL for `sorgbr`
DOUBLE PRECISION for `dorgbr`.
For $vect = 'Q'$, the array $tauq$ as returned by `?gebrd`.
For $vect = 'P'$, the array $taup$ as returned by `?gebrd`.
The dimension of tau must be at least $\max(1, \min(m, k))$
for $vect = 'Q'$, or $\max(1, \min(m, k))$ for $vect = 'P'$.

$lwork$ INTEGER. The size of the $work$ array.
See *Application notes* for the suggested value of $lwork$.

Output Parameters

a Overwritten by the orthogonal matrix Q or P^T (or the leading rows or columns thereof) as specified by $vect$, m , and n .

$work(1)$ If $info = 0$, on exit $work(1)$ contains the minimum value of $lwork$ required for optimum performance. Use this $lwork$ for subsequent runs.

$info$ INTEGER.
If $info = 0$, the execution is successful.
If $info = -i$, the i th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = \min(m, n) * blocksize$, where $blocksize$ is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*.

If you are in doubt how much workspace to supply, use a generous value of `lwork` for the first run. On exit, examine `work(1)` and use this value for subsequent runs.

The computed matrix Q differs from an exactly orthogonal matrix by a matrix E such that $\|E\|_2 = O(\epsilon)$.

The approximate numbers of floating-point operations for the cases listed in *Discussion* are as follows:

To form the whole of Q :

$$(4/3)n(3m^2 - 3m*n + n^2) \quad \text{if } m > n;$$

$$(4/3)m^3 \quad \text{if } m \leq n.$$

To form the n leading columns of Q when $m > n$:

$$(2/3)n^2(3m - n^2) \quad \text{if } m > n.$$

To form the whole of P^T :

$$(4/3)n^3 \quad \text{if } m \geq n;$$

$$(4/3)m(3n^2 - 3m*n + m^2) \quad \text{if } m < n.$$

To form the m leading columns of P^T when $m < n$:

$$(2/3)n^2(3m - n^2) \quad \text{if } m > n.$$

The complex counterpart of this routine is [?ungbr](#).

?ormbr

Multiplies an arbitrary real matrix by
the real orthogonal matrix Q or P^T
determined by ?gebrd.

```
call sormbr (vect,side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info)
call dormbr (vect,side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info)
```

Discussion

Given an arbitrary real matrix C , this routine forms one of the matrix products QC , $Q^T C$, CQ , CQ^T , PC , $P^T C$, CP , or CP^T , where Q and P are orthogonal matrices computed by a call to [sgebrd/dgebrd](#) (see [page 5-76](#)). The routine overwrites the product on C .

Input Parameters

In the descriptions below, r denotes the order of Q or P^T :

If $side = 'L'$, $r = m$; if $side = 'R'$, $r = n$.

vect CHARACTER*1. Must be 'Q' or 'P'.
If $vect = 'Q'$, then Q or Q^T is applied to C .
If $vect = 'P'$, then P or P^T is applied to C .

side CHARACTER*1. Must be 'L' or 'R'.
If $side = 'L'$, multipliers are applied to C from the left.
If $side = 'R'$, they are applied to C from the right.

trans CHARACTER*1. Must be 'N' or 'T'.
If $trans = 'N'$, then Q or P is applied to C .
If $trans = 'T'$, then Q^T or P^T is applied to C .

m INTEGER. The number of rows in C .

n INTEGER. The number of columns in C .

k INTEGER. One of the dimensions of A in ?gebrd:
If $vect = 'Q'$, the number of columns in A ;
If $vect = 'P'$, the number of rows in A .
Constraints: $m \geq 0$, $n \geq 0$, $k \geq 0$.

a, *c*, *work* REAL for *sormbr*
 DOUBLE PRECISION for *dormbr*.
 Arrays:
a(lda,)* is the array *a* as returned by *?gebrd*.
 Its second dimension must be at least $\max(1, \min(r,k))$
 for *vect* = 'Q', or $\max(1, r)$ for *vect* = 'P'.
c(ldc,)* holds the matrix *C*.
 Its second dimension must be at least $\max(1, n)$.
work(lwork) is a workspace array.

lda INTEGER. The first dimension of *a*. Constraints:
lda $\geq \max(1, r)$ if *vect* = 'Q';
lda $\geq \max(1, \min(r,k))$ if *vect* = 'P'.

ldc INTEGER. The first dimension of *c*; *ldc* $\geq \max(1, m)$.

tau REAL for *sormbr*
 DOUBLE PRECISION for *dormbr*.
 Array, DIMENSION at least $\max(1, \min(r, k))$.
 For *vect* = 'Q', the array *taug* as returned by *?gebrd*.
 For *vect* = 'P', the array *taup* as returned by *?gebrd*.

lwork INTEGER. The size of the *work* array. Constraints:
lwork $\geq \max(1, n)$ if *side* = 'L';
lwork $\geq \max(1, m)$ if *side* = 'R'.
 See *Application notes* for the suggested value of *lwork*.

Output Parameters

c Overwritten by the product *QC*, *Q^TC*, *CQ*, *CQ^T*, *PC*,
P^TC, *CP*, or *CP^T*, as specified by *vect*, *side*, and
trans.

work(1) If *info* = 0, on exit *work(1)* contains the minimum
 value of *lwork* required for optimum performance. Use
 this *lwork* for subsequent runs.

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using

$lwork = n * blocksize$ for $side = 'L'$, or

$lwork = m * blocksize$ for $side = 'R'$,

where $blocksize$ is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*.

If you are in doubt how much workspace to supply, use a generous value of $lwork$ for the first run. On exit, examine $work(1)$ and use this value for subsequent runs.

The computed product differs from the exact product by a matrix E such that $\|E\|_2 = O(\epsilon) \|C\|_2$.

The total number of floating-point operations is approximately

$2 * n * k * (2 * m - k)$ if $side = 'L'$ and $m \geq k$;

$2 * m * k * (2 * n - k)$ if $side = 'R'$ and $n \geq k$;

$2 * m^2 * n$ if $side = 'L'$ and $m < k$;

$2 * n^2 * m$ if $side = 'R'$ and $n < k$.

The complex counterpart of this routine is [?unmbr](#).

?ungbr

Generates the complex unitary matrix Q
or P^H determined by ?gebrd.

```
call cungbr ( vect, m, n, k, a, lda, tau, work, lwork, info )
call zungbr ( vect, m, n, k, a, lda, tau, work, lwork, info )
```

Discussion

The routine generates the whole or part of the unitary matrices Q and P^H formed by the routines `cgebrd/zgebrd` (see [page 5-76](#)). Use this routine after a call to `cgebrd/zgebrd`. All valid combinations of arguments are described in *Input Parameters*; in most cases you'll need the following:

To compute the whole m by m matrix Q :

```
call ?ungbr ( 'Q', m, m, n, a ... )
```

(note that the array a must have at least m columns).

To form the n leading columns of Q if $m > n$:

```
call ?ungbr ( 'Q', m, n, n, a ... )
```

To compute the whole n by n matrix P^H :

```
call ?ungbr ( 'P', n, n, m, a ... )
```

(note that the array a must have at least n rows).

To form the m leading rows of P^H if $m < n$:

```
call ?ungbr ( 'P', m, n, m, a ... )
```

Input Parameters

<i>vect</i>	CHARACTER*1. Must be 'Q' or 'P'. If <i>vect</i> = 'Q', the routine generates the matrix Q . If <i>vect</i> = 'P', the routine generates the matrix P^H .
<i>m</i>	INTEGER. The number of required rows of Q or P^H .
<i>n</i>	INTEGER. The number of required columns of Q or P^H .
<i>k</i>	INTEGER. One of the dimensions of A in ?gebrd: If <i>vect</i> = 'Q', the number of columns in A ; If <i>vect</i> = 'P', the number of rows in A .

Constraints: $m \geq 0$, $n \geq 0$, $k \geq 0$.
 For $vect = 'Q'$: $k \leq n \leq m$ if $m > k$, or $m = n$ if $m \leq k$.
 For $vect = 'P'$: $k \leq m \leq n$ if $n > k$, or $m = n$ if $n \leq k$.

$a, work$ COMPLEX for `cungbr`
 DOUBLE COMPLEX for `zungbr`.
 Arrays:
 $a(lda, *)$ is the array a as returned by `?gebrd`.
 The second dimension of a must be at least $\max(1, n)$.
 $work(lwork)$ is a workspace array.

lda INTEGER. The first dimension of a ; at least $\max(1, m)$.

tau COMPLEX for `cungbr`
 DOUBLE COMPLEX for `zungbr`.
 For $vect = 'Q'$, the array $tauq$ as returned by `?gebrd`.
 For $vect = 'P'$, the array $taup$ as returned by `?gebrd`.
 The dimension of tau must be at least $\max(1, \min(m, k))$
 for $vect = 'Q'$, or $\max(1, \min(m, k))$ for $vect = 'P'$.

$lwork$ INTEGER. The size of the $work$ array.
 Constraint: $lwork \geq \max(1, \min(m, n))$.
 See *Application notes* for the suggested value of $lwork$.

Output Parameters

a Overwritten by the orthogonal matrix Q or P^T (or the leading rows or columns thereof) as specified by $vect$, m , and n .

$work(1)$ If $info = 0$, on exit $work(1)$ contains the minimum value of $lwork$ required for optimum performance. Use this $lwork$ for subsequent runs.

$info$ INTEGER.
 If $info = 0$, the execution is successful.
 If $info = -i$, the i th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = \min(m, n) * blocksize$, where $blocksize$ is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*.

If you are in doubt how much workspace to supply, use a generous value of $lwork$ for the first run. On exit, examine $work(1)$ and use this value for subsequent runs.

The computed matrix Q differs from an exactly orthogonal matrix by a matrix E such that $\|E\|_2 = O(\epsilon)$.

The approximate numbers of floating-point operations for the cases listed in *Discussion* are as follows:

To form the whole of Q :

$$\begin{aligned} (16/3)n(3m^2 - 3m*n + n^2) & \quad \text{if } m > n; \\ (16/3)m^3 & \quad \text{if } m \leq n. \end{aligned}$$

To form the n leading columns of Q when $m > n$:

$$(8/3)n^2(3m - n^2) \quad \text{if } m > n.$$

To form the whole of P^T :

$$\begin{aligned} (16/3)n^3 & \quad \text{if } m \geq n; \\ (16/3)m(3n^2 - 3m*n + m^2) & \quad \text{if } m < n. \end{aligned}$$

To form the m leading columns of P^T when $m < n$:

$$(8/3)n^2(3m - n^2) \quad \text{if } m > n.$$

The real counterpart of this routine is [?orgbr](#).

?unmbr

*Multiplies an arbitrary complex matrix
by the unitary matrix Q or P determined
by ?gebrd.*

```
call cunmbr (vect,side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info)
call zunmbr (vect,side,trans,m,n,k,a,lda,tau,c,ldc,work,lwork,info)
```

Discussion

Given an arbitrary complex matrix C , this routine forms one of the matrix products QC , $Q^H C$, CQ , CQ^H , PC , $P^H C$, CP , or CP^H , where Q and P are orthogonal matrices computed by a call to [cgebrd/zgebrd](#) (see [page 5-76](#)). The routine overwrites the product on C .

Input Parameters

In the descriptions below, r denotes the order of Q or P^H :

If $side = 'L'$, $r = m$; if $side = 'R'$, $r = n$.

vect CHARACTER*1. Must be 'Q' or 'P'.
If $vect = 'Q'$, then Q or Q^H is applied to C .
If $vect = 'P'$, then P or P^H is applied to C .

side CHARACTER*1. Must be 'L' or 'R'.
If $side = 'L'$, multipliers are applied to C from the left.
If $side = 'R'$, they are applied to C from the right.

trans CHARACTER*1. Must be 'N' or 'C'.
If $trans = 'N'$, then Q or P is applied to C .
If $trans = 'C'$, then Q^H or P^H is applied to C .

m INTEGER. The number of rows in C .

n INTEGER. The number of columns in C .

k INTEGER. One of the dimensions of A in [?gebrd](#):
If $vect = 'Q'$, the number of columns in A ;
If $vect = 'P'$, the number of rows in A .
Constraints: $m \geq 0$, $n \geq 0$, $k \geq 0$.

<i>a, c, work</i>	<p>COMPLEX for <code>cunmbr</code> DOUBLE COMPLEX for <code>zunmbr</code>. Arrays: <i>a(l da, *)</i> is the array <i>a</i> as returned by <code>?gebrd</code>. Its second dimension must be at least $\max(1, \min(r, k))$ for <i>vect</i> = 'Q', or $\max(1, r)$ for <i>vect</i> = 'P'. <i>c(l dc, *)</i> holds the matrix <i>C</i>. Its second dimension must be at least $\max(1, n)$. <i>work(l work)</i> is a workspace array.</p>
<i>lda</i>	<p>INTEGER. The first dimension of <i>a</i>. Constraints: <i>lda</i> $\geq \max(1, r)$ if <i>vect</i> = 'Q'; <i>lda</i> $\geq \max(1, \min(r, k))$ if <i>vect</i> = 'P'.</p>
<i>ldc</i>	<p>INTEGER. The first dimension of <i>c</i>; <i>ldc</i> $\geq \max(1, m)$.</p>
<i>tau</i>	<p>COMPLEX for <code>cunmbr</code> DOUBLE COMPLEX for <code>zunmbr</code>. Array, DIMENSION at least $\max(1, \min(r, k))$. For <i>vect</i> = 'Q', the array <i>taug</i> as returned by <code>?gebrd</code>. For <i>vect</i> = 'P', the array <i>taup</i> as returned by <code>?gebrd</code>.</p>
<i>lwork</i>	<p>INTEGER. The size of the <i>work</i> array. Constraints: <i>lwork</i> $\geq \max(1, n)$ if <i>side</i> = 'L'; <i>lwork</i> $\geq \max(1, m)$ if <i>side</i> = 'R'. See <i>Application notes</i> for the suggested value of <i>lwork</i>.</p>

Output Parameters

<i>c</i>	<p>Overwritten by the product $QC, Q^H C, CQ, CQ^H, PC, P^H C, CP,$ or CP^H, as specified by <i>vect</i>, <i>side</i>, and <i>trans</i>.</p>
<i>work(1)</i>	<p>If <i>info</i> = 0, on exit <i>work(1)</i> contains the minimum value of <i>lwork</i> required for optimum performance. Use this <i>lwork</i> for subsequent runs.</p>
<i>info</i>	<p>INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = -<i>i</i>, the <i>i</i>th parameter had an illegal value.</p>

Application Notes

For better performance, try using

$lwork = n * blocksize$ for $side = 'L'$, or

$lwork = m * blocksize$ for $side = 'R'$,

where $blocksize$ is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*.

If you are in doubt how much workspace to supply, use a generous value of $lwork$ for the first run. On exit, examine $work(1)$ and use this value for subsequent runs.

The computed product differs from the exact product by a matrix E such that $\|E\|_2 = O(\epsilon) \|C\|_2$.

The total number of floating-point operations is approximately

$8 * n * k * (2 * m - k)$ if $side = 'L'$ and $m \geq k$;

$8 * m * k * (2 * n - k)$ if $side = 'R'$ and $n \geq k$;

$8 * m^2 * n$ if $side = 'L'$ and $m < k$;

$8 * n^2 * m$ if $side = 'R'$ and $n < k$.

The real counterpart of this routine is [?ormbr](#).

?bdsqr

Computes the singular value decomposition of a general matrix that has been reduced to bidiagonal form.

```
call sbdsqr ( uplo, n, ncv, nru, ncc, d, e, vt, ldvt, u, ldu,
             c, ldc, work, info )
call dbdsqr ( uplo, n, ncv, nru, ncc, d, e, vt, ldvt, u, ldu,
             c, ldc, work, info )
call cbdsqr ( uplo, n, ncv, nru, ncc, d, e, vt, ldvt, u, ldu,
             c, ldc, work, info )
call zbdsqr ( uplo, n, ncv, nru, ncc, d, e, vt, ldvt, u, ldu,
             c, ldc, work, info )
```

Discussion

This routine computes the singular values and, optionally, the right and/or left singular vectors from the [Singular Value Decomposition \(SVD\)](#) of a real n -by- n (upper or lower) bidiagonal matrix B using the implicit zero-shift QR algorithm. The SVD of B has the form $B = Q * S * P^H$ where S is the diagonal matrix of singular values, Q is an orthogonal matrix of left singular vectors, and P is an orthogonal matrix of right singular vectors. If left singular vectors are requested, this subroutine actually returns $U * Q$ instead of Q , and, if right singular vectors are requested, this subroutine returns $P^H * VT$ instead of P^H , for given real/complex input matrices U and VT .

When U and VT are the orthogonal/unitary matrices that reduce a general matrix A to bidiagonal form: $A = U * B * VT$, as computed by [?gebrd](#), then

$$A = (U * Q) * S * (P^H * VT)$$

is the SVD of A . Optionally, the subroutine may also compute $Q^H * C$ for a given real/complex input matrix C .

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', B is an upper bidiagonal matrix.
 If *uplo* = 'L', B is a lower bidiagonal matrix.

<i>n</i>	INTEGER . The order of the matrix <i>B</i> ($n \geq 0$).
<i>ncvt</i>	INTEGER . The number of columns of the matrix <i>VT</i> , that is, the number of right singular vectors ($ncvt \geq 0$). Set $ncvt = 0$ if no right singular vectors are required.
<i>nru</i>	INTEGER . The number of rows in <i>U</i> , that is, the number of left singular vectors ($nru \geq 0$). Set $nru = 0$ if no left singular vectors are required.
<i>ncc</i>	INTEGER . The number of columns in the matrix <i>C</i> used for computing the product $Q^H C$ ($ncc \geq 0$). Set $ncc = 0$ if no matrix <i>C</i> is supplied.
<i>d, e, work</i>	REAL for single-precision flavors DOUBLE PRECISION for double-precision flavors. Arrays: <i>d</i> (*) contains the diagonal elements of <i>B</i> . The dimension of <i>d</i> must be at least $\max(1, n)$. <i>e</i> (*) contains the $(n-1)$ off-diagonal elements of <i>B</i> . The dimension of <i>e</i> must be at least $\max(1, n)$. <i>e</i> (<i>n</i>) is used for workspace. <i>work</i> (*) is a workspace array. The dimension of <i>work</i> must be at least $\max(1, 2*n)$ if $ncvt = nru = ncc = 0$; $\max(1, 4*(n-1))$ otherwise.
<i>vt, u, c</i>	REAL for sbdsqr DOUBLE PRECISION for dbdsqr COMPLEX for cbdsqr DOUBLE COMPLEX for zbdqsqr . Arrays: <i>vt</i> (<i>ldvt</i> ,*) contains an <i>n</i> by <i>ncvt</i> matrix <i>VT</i> . The second dimension of <i>vt</i> must be at least $\max(1, ncvt)$. <i>vt</i> is not referenced if $ncvt = 0$. <i>u</i> (<i>ldu</i> ,*) contains an <i>nru</i> by <i>n</i> unit matrix <i>U</i> . The second dimension of <i>u</i> must be at least $\max(1, n)$. <i>u</i> is not referenced if $nru = 0$.

	$c(ldc, *)$ contains the matrix C for computing the product $Q^H * C$. The second dimension of c must be at least $\max(1, ncc)$. The array is not referenced if $ncc = 0$.
$ldvt$	INTEGER. The first dimension of vt . Constraints: $ldvt \geq \max(1, n)$ if $ncvt > 0$; $ldvt \geq 1$ if $ncvt = 0$.
ldu	INTEGER. The first dimension of u . Constraint: $ldu \geq \max(1, nru)$.
ldc	INTEGER. The first dimension of c . Constraints: $ldc \geq \max(1, n)$ if $ncc > 0$; $ldc \geq 1$ otherwise.

Output Parameters

d	On exit, if $info = 0$, overwritten by the singular values in decreasing order (see $info$).
e	On exit, if $info = 0$, e is destroyed. See also $info$ below.
c	Overwritten by the product $Q^H * C$.
vt	On exit, this array is overwritten by $P^H * VT$.
u	On exit, this array is overwritten by $U * Q$.
$info$	INTEGER. If $info = 0$, the execution is successful. If $info = -i$, the i th parameter had an illegal value. If $info = i$, the algorithm failed to converge; i specifies how many off-diagonals did not converge. In this case, d and e contain on exit the diagonal and off-diagonal elements, respectively, of a bidiagonal matrix orthogonally equivalent to B .

Application Notes

Each singular value and singular vector is computed to high relative accuracy. However, the reduction to bidiagonal form (prior to calling the routine) may decrease the relative accuracy in the small singular values of the original matrix if its singular values vary widely in magnitude.

If σ_i is an exact singular value of B , and s_i is the corresponding computed value, then

$$|s_i - \sigma_i| \leq p(m, n)\epsilon\sigma_i$$

where $p(m, n)$ is a modestly increasing function of m and n , and ϵ is the machine precision. If only singular values are computed, they are computed more accurately than when some singular vectors are also computed (that is, the function $p(m, n)$ is smaller).

If u_i is the corresponding exact left singular vector of B , and w_i is the corresponding computed left singular vector, then the angle $\theta(u_i, w_i)$ between them is bounded as follows:

$$\theta(u_i, w_i) \leq p(m, n)\epsilon / \min_{i \neq j} (|\sigma_i - \sigma_j| / |\sigma_i + \sigma_j|).$$

Here $\min_{i \neq j} (|\sigma_i - \sigma_j| / |\sigma_i + \sigma_j|)$ is the *relative gap* between σ_i and the other singular values. A similar error bound holds for the right singular vectors.

The total number of real floating-point operations is roughly proportional to n^2 if only the singular values are computed. About $6n^2 * nru$ additional operations ($12n^2 * nru$ for complex flavors) are required to compute the left singular vectors and about $6n^2 * ncvt$ operations ($12n^2 * ncvt$ for complex flavors) to compute the right singular vectors.

?bdsdc

Computes the singular value decomposition of a real bidiagonal matrix using a divide and conquer method.

```
call sbdsdc ( uplo, compq, n, d, e, u, ldu, vt, ldvt, q, iq, work,
             iwork, info )
call dbdsdc ( uplo, compq, n, d, e, u, ldu, vt, ldvt, q, iq, work,
             iwork, info )
```

Discussion

This routine computes the [Singular Value Decomposition](#) (SVD) of a real n -by- n (upper or lower) bidiagonal matrix B : $B = U \Sigma V^T$, using a divide and conquer method, where Σ is a diagonal matrix with non-negative diagonal elements (the singular values of B), and U and V are orthogonal matrices of left and right singular vectors, respectively. `?bdsdc` can be used to compute all singular values, and optionally, singular vectors or singular vectors in compact form.

Input Parameters

`uplo` CHARACTER*1. Must be 'U' or 'L'.
 If `uplo` = 'U', B is an upper bidiagonal matrix.
 If `uplo` = 'L', B is a lower bidiagonal matrix.

`compq` CHARACTER*1. Must be 'N', 'P', or 'I'.
 If `compq` = 'N', compute singular values only.
 If `compq` = 'P', compute singular values and compute singular vectors in compact form.
 If `compq` = 'I', compute singular values and singular vectors.

`n` INTEGER. The order of the matrix B ($n \geq 0$).

`d, e, work` REAL for `sbdsdc`
 DOUBLE PRECISION for `dbdsdc`.
 Arrays:

$d(*)$ contains the n diagonal elements of the bidiagonal matrix B . The dimension of d must be at least $\max(1, n)$.

$e(*)$ contains the off-diagonal elements of the bidiagonal matrix B . The dimension of e must be at least $\max(1, n)$.

$work(*)$ is a workspace array.

The dimension of $work$ must be at least:

$\max(1, 4*n)$, if $compq = 'N'$;

$\max(1, 6*n)$, if $compq = 'P'$;

$\max(1, 3*n^2+4*n)$, if $compq = 'I'$.

ldu **INTEGER**. The first dimension of the output array u ;
 $ldu \geq 1$. If singular vectors are desired, then
 $ldu \geq \max(1, n)$.

$ldvt$ **INTEGER**. The first dimension of the output array vt ;
 $ldvt \geq 1$. If singular vectors are desired, then
 $ldvt \geq \max(1, n)$.

$iwork$ **INTEGER**.
 Workspace array, dimension at least $\max(1, 8*n)$.

Output Parameters

d If $info = 0$, overwritten by the singular values of B .

e On exit, e is overwritten.

u, vt, q **REAL** for `sbdsdc`
DOUBLE PRECISION for `sbdsdc`.
 Arrays: $u(ldu, *)$, $vt(ldvt, *)$, $q(*)$.
 If $compq = 'I'$, then on exit u contains the left singular vectors of the bidiagonal matrix B , unless $info \neq 0$ (see *info*). For other values of $compq$, u is not referenced. The second dimension of u must be at least $\max(1, n)$.
 If $compq = 'I'$, then on exit vt contains the right singular vectors of the bidiagonal matrix B , unless $info \neq 0$ (see *info*). For other values of $compq$, vt is not referenced. The second dimension of vt must be at least $\max(1, n)$.

If `compq = 'P'`, then on exit, if `info = 0`, `q` and `iq` contain the left and right singular vectors in a compact form. Specifically, `q` contains all the `REAL` (for `sbdsc`) or `DOUBLE PRECISION` (for `dbdsc`) data for singular vectors. For other values of `compq`, `q` is not referenced. See *Application notes* for details.

`iq`

`INTEGER`.

Array: `iq(*)`.

If `compq = 'P'`, then on exit, if `info = 0`, `q` and `iq` contain the left and right singular vectors in a compact form. Specifically, `iq` contains all the `INTEGER` data for singular vectors. For other values of `compq`, `iq` is not referenced. See *Application notes* for details.

`info`

`INTEGER`.

If `info = 0`, the execution is successful.

If `info = -i`, the `i`th parameter had an illegal value.

If `info = i`, the algorithm failed to compute a singular value. The update process of divide and conquer failed.

Symmetric Eigenvalue Problems

Symmetric eigenvalue problems are posed as follows: given an n by n real symmetric or complex Hermitian matrix A , find the *eigenvalues* λ and the corresponding *eigenvectors* z that satisfy the equation

$$Az = \lambda z. \text{ (or, equivalently, } z^H A = \lambda z^H \text{).}$$

In such eigenvalue problems, all n eigenvalues are real not only for real symmetric but also for complex Hermitian matrices A , and there exists an orthonormal system of n eigenvectors. If A is a symmetric or Hermitian positive-definite matrix, all eigenvalues are positive.

To solve a symmetric eigenvalue problem with LAPACK, you usually need to reduce the matrix to tridiagonal form and then solve the eigenvalue problem with the tridiagonal matrix obtained. LAPACK includes routines for reducing the matrix to a tridiagonal form by an orthogonal (or unitary) similarity transformation $A = QTQ^H$ as well as for solving tridiagonal symmetric eigenvalue problems. These routines are listed in [Table 5-3](#).

There are different routines for symmetric eigenvalue problems, depending on whether you need all eigenvectors or only some of them or eigenvalues only, whether the matrix A is positive-definite or not, and so on.

These routines are based on three primary algorithms for computing eigenvalues and eigenvectors of symmetric problems: the divide and conquer algorithm, the QR algorithm, and bisection followed by inverse iteration. The divide and conquer algorithm is generally more efficient and is recommended for computing all eigenvalues and eigenvectors.

Furthermore, to solve an eigenvalue problem using the divide and conquer algorithm, you need to call only one routine. In general, more than one routine has to be called if the QR algorithm or bisection followed by inverse iteration is used.

Decision tree in [Figure 5-2](#) will help you choose the right routine or sequence of routines for eigenvalue problems with real symmetric matrices. A similar decision tree for complex Hermitian matrices is presented in [Figure 5-3](#).

Figure 5-2 Decision Tree: Real Symmetric Eigenvalue Problems

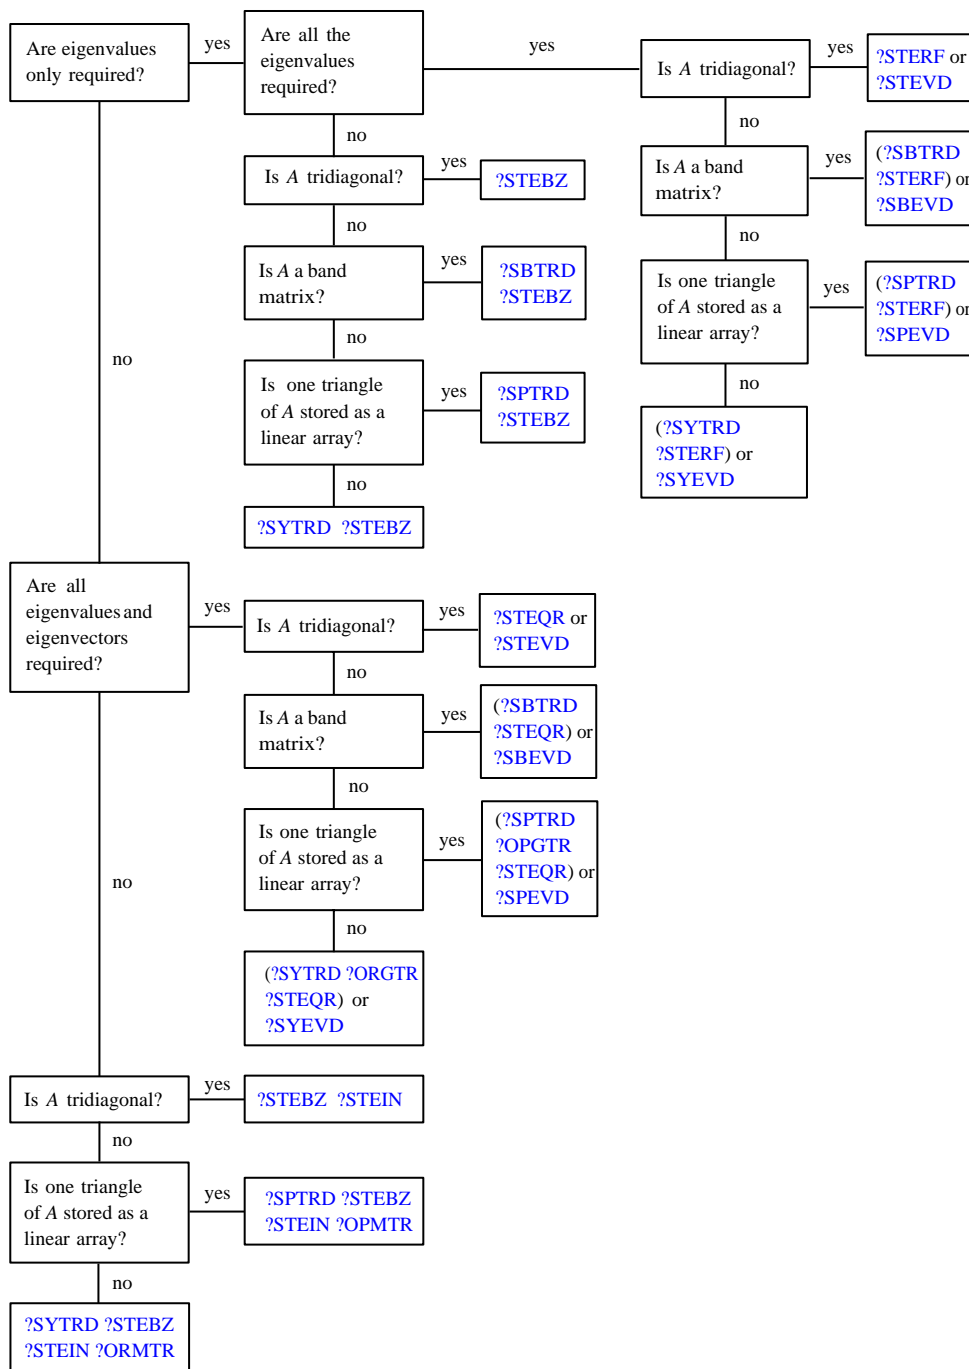


Figure 5-3 Decision Tree: Complex Hermitian Eigenvalue Problems

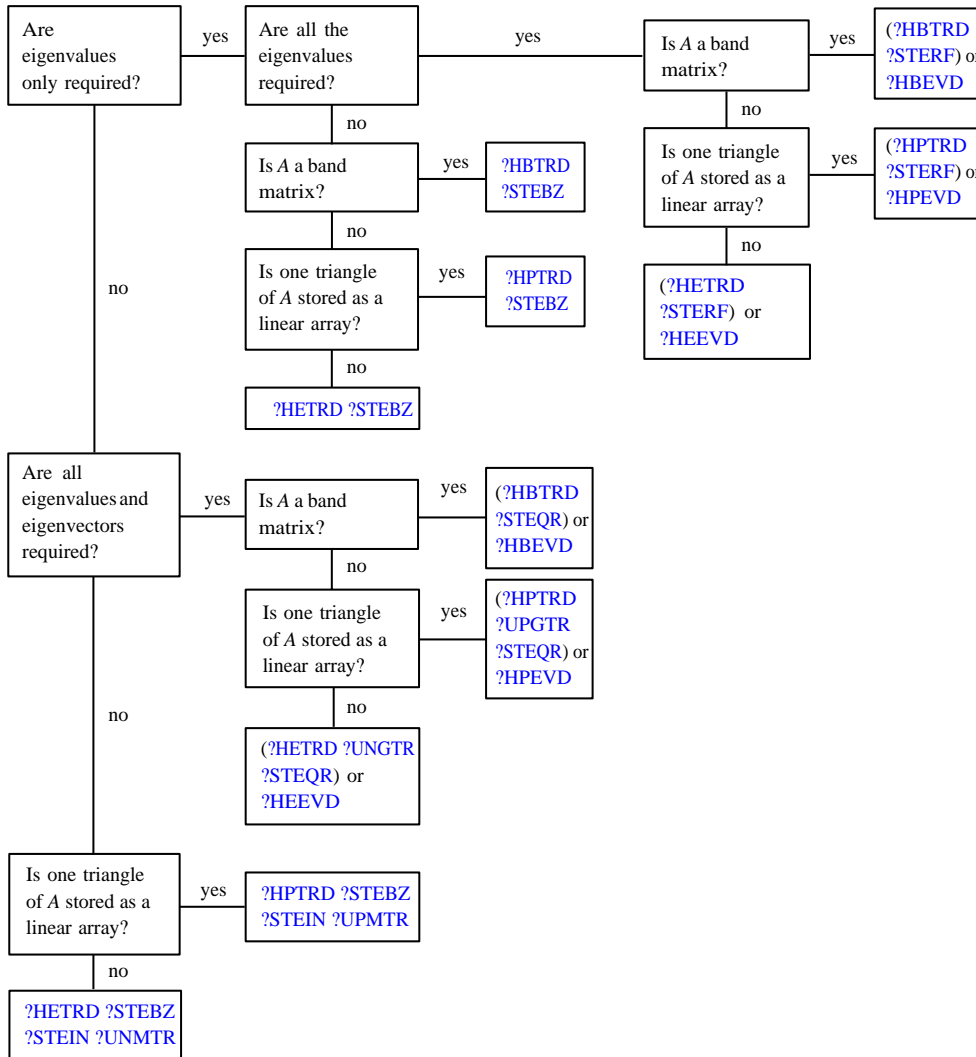


Table 5-3 Computational Routines for Solving Symmetric Eigenvalue Problems

Operation	Real symmetric matrices	Complex Hermitian matrices		
Reduce to tridiagonal form $A = QTQ^H$ (full storage)	<u>?sytrd</u>	<u>?hetrd</u>		
Reduce to tridiagonal form $A = QTQ^H$ (packed storage)	<u>?sptrd</u>	<u>?hptrd</u>		
Reduce to tridiagonal form $A = QTQ^H$ (band storage).	<u>?sbtrd</u>	<u>?hbtrd</u>		
Generate matrix Q (full storage)	<u>?orgtr</u>	<u>?ungtr</u>		
Generate matrix Q (packed storage)	<u>?opgtr</u>	<u>?upgtr</u>		
Apply matrix Q (full storage)	<u>?ormtr</u>	<u>?unmtr</u>		
Apply matrix Q (packed storage)	<u>?opmtr</u>	<u>?upmtr</u>		
Find all eigenvalues of a tridiagonal matrix T	<u>?sterf</u>			
Find all eigenvalues and eigenvectors of a tridiagonal matrix T	<u>?stegr</u>	<u>?stedc</u>	<u>?stegr</u>	<u>?stedc</u>
Find all eigenvalues and eigenvectors of a tridiagonal positive-definite matrix T .	<u>?ptegr</u>	<u>?ptegr</u>		
Find selected eigenvalues of a tridiagonal matrix T	<u>?stebz</u> <u>?stegr</u>	<u>?stegr</u>		
Find selected eigenvectors of a tridiagonal matrix T	<u>?stein</u> <u>?stegr</u>	<u>?stein</u> <u>?stegr</u>		
Compute the reciprocal condition numbers for the eigenvectors	<u>?disna</u>	<u>?disna</u>		

?sytrd

Reduces a real symmetric matrix to tridiagonal form.

```
call ssytrd ( uplo,n,a,lda,d,e,tau,work,lwork,info )
call dsytrd ( uplo,n,a,lda,d,e,tau,work,lwork,info )
```

Discussion

This routine reduces a real symmetric matrix A to symmetric tridiagonal form T by an orthogonal similarity transformation: $A = QTQ^T$. The orthogonal matrix Q is not formed explicitly but is represented as a product of $n-1$ elementary reflectors. Routines are provided for working with Q in this representation. (They are described later in this section.)

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', *a* stores the upper triangular part of A .
 If *uplo* = 'L', *a* stores the lower triangular part of A .

n INTEGER. The order of the matrix A ($n \geq 0$).

a, *work* REAL for *ssytrd*
 DOUBLE PRECISION for *dsytrd*.
a(*lda*,*) is an array containing either upper or lower triangular part of the matrix A , as specified by *uplo*.
 The second dimension of *a* must be at least $\max(1, n)$.
work(*lwork*) is a workspace array.

lda INTEGER. The first dimension of *a*; at least $\max(1, n)$.

lwork INTEGER. The size of the *work* array ($lwork \geq n$)
 See *Application notes* for the suggested value of *lwork*.

Output Parameters

a Overwritten by the tridiagonal matrix T and details of the orthogonal matrix Q , as specified by *uplo*.

<i>d</i> , <i>e</i> , <i>tau</i>	REAL for <i>ssytrd</i> DOUBLE PRECISION for <i>dsytrd</i> . Arrays: <i>d</i> (*) contains the diagonal elements of the matrix <i>T</i> . The dimension of <i>d</i> must be at least $\max(1, n)$. <i>e</i> (*) contains the off-diagonal elements of <i>T</i> . The dimension of <i>e</i> must be at least $\max(1, n-1)$. <i>tau</i> (*) stores further details of the orthogonal matrix <i>Q</i> . The dimension of <i>tau</i> must be at least $\max(1, n-1)$.
<i>work</i> (1)	If <i>info</i> =0, on exit <i>work</i> (1) contains the minimum value of <i>lwork</i> required for optimum performance. Use this <i>lwork</i> for subsequent runs.
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = n * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work*(1) and use this value for subsequent runs.

The computed matrix *T* is exactly similar to a matrix $A + E$, where $\|E\|_2 = c(n)\epsilon \|A\|_2$, $c(n)$ is a modestly increasing function of *n*, and ϵ is the machine precision.

The approximate number of floating-point operations is $(4/3)n^3$.

After calling this routine, you can call the following:

[?orgtr](#) to form the computed matrix *Q* explicitly;

[?ormtr](#) to multiply a real matrix by *Q*.

The complex counterpart of this routine is [?hetrd](#).

?orgtr

Generates the real orthogonal matrix Q determined by ?sytrd.

```
call sorgtr ( uplo, n, a, lda, tau, work, lwork, info )
call dorgtr ( uplo, n, a, lda, tau, work, lwork, info )
```

Discussion

The routine explicitly generates the n by n orthogonal matrix Q formed by ?sytrd (see [page 5-105](#)) when reducing a real symmetric matrix A to tridiagonal form: $A = QTQ^T$. Use this routine after a call to ?sytrd.

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Use the same **uplo** as supplied to ?sytrd.

n INTEGER. The order of the matrix Q ($n \geq 0$).

a, tau, work REAL for **sorgtr**
DOUBLE PRECISION for **dorgtr**.
Arrays:
a(lda,*) is the array **a** as returned by ?sytrd.
The second dimension of **a** must be at least $\max(1, n)$.
tau(*) is the array **tau** as returned by ?sytrd.
The dimension of **tau** must be at least $\max(1, n-1)$.
work(lwork) is a workspace array.

lda INTEGER. The first dimension of **a**; at least $\max(1, n)$.

lwork INTEGER. The size of the **work** array ($lwork \geq n$)
See *Application notes* for the suggested value of **lwork**.

Output Parameters

a Overwritten by the orthogonal matrix Q .

work(1) If *info* = 0, on exit *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = (n-1) * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*.

If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

The computed matrix *Q* differs from an exactly orthogonal matrix by a matrix *E* such that $\|E\|_2 = O(\epsilon)$, where ϵ is the machine precision.

The approximate number of floating-point operations is $(4/3)n^3$.

The complex counterpart of this routine is [?ungtr](#).

?ormtr

Multiplies a real matrix by the real orthogonal matrix Q determined by ?sytrd.

```
call sormtr ( side,uplo,trans,m,n,a,lda,tau,c,ldc,work,lwork,info )
call dormtr ( side,uplo,trans,m,n,a,lda,tau,c,ldc,work,lwork,info )
```

Discussion

The routine multiplies a real matrix C by Q or Q^T , where Q is the orthogonal matrix Q formed by ?sytrd (see [page 5-105](#)) when reducing a real symmetric matrix A to tridiagonal form: $A = QTQ^T$. Use this routine after a call to ?sytrd.

Depending on the parameters *side* and *trans*, the routine can form one of the matrix products QC , Q^TC , CQ , or CQ^T (overwriting the result on C).

Input Parameters

In the descriptions below, r denotes the order of Q :

If *side* = 'L', $r = m$; if *side* = 'R', $r = n$.

<i>side</i>	CHARACTER*1. Must be either 'L' or 'R'. If <i>side</i> = 'L', Q or Q^T is applied to C from the left. If <i>side</i> = 'R', Q or Q^T is applied to C from the right.
<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. Use the same <i>uplo</i> as supplied to ?sytrd.
<i>trans</i>	CHARACTER*1. Must be either 'N' or 'T'. If <i>trans</i> = 'N', the routine multiplies C by Q . If <i>trans</i> = 'T', the routine multiplies C by Q^T .
<i>m</i>	INTEGER. The number of rows in the matrix C ($m \geq 0$).
<i>n</i>	INTEGER. The number of columns in C ($n \geq 0$).
<i>a,work,tau,c</i>	REAL for sormtr DOUBLE PRECISION for dormtr. <i>a(lda,*)</i> and <i>tau</i> are the arrays returned by ?sytrd.

The second dimension of *a* must be at least $\max(1, r)$.
 The dimension of *tau* must be at least $\max(1, r-1)$.

c(ldc,)* contains the matrix *C*.

The second dimension of *c* must be at least $\max(1, n)$

work(lwork) is a workspace array.

lda **INTEGER.** The first dimension of *a*; $lda \geq \max(1, r)$.

ldc **INTEGER.** The first dimension of *c*; $ldc \geq \max(1, n)$.

lwork **INTEGER.** The size of the *work* array. Constraints:

lwork $\geq \max(1, n)$ if *side* = 'L';

lwork $\geq \max(1, m)$ if *side* = 'R'.

See *Application notes* for the suggested value of *lwork*.

Output Parameters

c Overwritten by the product *QC*, *Q^TC*, *CQ*, or *CQ^T*
 (as specified by *side* and *trans*).

work(1) If *info* = 0, on exit *work(1)* contains the minimum
 value of *lwork* required for optimum performance. Use
 this *lwork* for subsequent runs.

info **INTEGER.**
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = n * blocksize$ for *side* = 'L', or
 $lwork = m * blocksize$ for *side* = 'R', where *blocksize* is a
 machine-dependent value (typically, 16 to 64) required for optimum
 performance of the *blocked algorithm*. If you are in doubt how much
 workspace to supply, use a generous value of *lwork* for the first run. On
 exit, examine *work(1)* and use this value for subsequent runs.

The computed product differs from the exact product by a matrix *E* such
 that $\|E\|_2 = O(\epsilon) \|C\|_2$.

The total number of floating-point operations is approximately $2 * m^2 * n$ if
side = 'L' or $2 * n^2 * m$ if *side* = 'R'.

The complex counterpart of this routine is [?unmtr](#).

?hetrd

Reduces a complex Hermitian matrix to tridiagonal form.

```
call chetrd ( uplo,n,a,lda,d,e,tau,work,lwork,info )
call zhetrd ( uplo,n,a,lda,d,e,tau,work,lwork,info )
```

Discussion

This routine reduces a complex Hermitian matrix A to symmetric tridiagonal form T by a unitary similarity transformation: $A = QTQ^H$. The unitary matrix Q is not formed explicitly but is represented as a product of $n-1$ elementary reflectors. Routines are provided to work with Q in this representation. (They are described later in this section.)

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
If *uplo* = 'U', *a* stores the upper triangular part of A .
If *uplo* = 'L', *a* stores the lower triangular part of A .

n INTEGER. The order of the matrix A ($n \geq 0$).

a, *work* COMPLEX for *chetrd*
DOUBLE COMPLEX for *zhetrd*.
a(*lda*,*) is an array containing either upper or lower triangular part of the matrix A , as specified by *uplo*.
The second dimension of *a* must be at least $\max(1, n)$.
work(*lwork*) is a workspace array.

lda INTEGER. The first dimension of *a*; at least $\max(1, n)$.

lwork INTEGER. The size of the *work* array ($lwork \geq n$)
See *Application notes* for the suggested value of *lwork*.

Output Parameters

a Overwritten by the tridiagonal matrix T and details of the unitary matrix Q , as specified by *uplo*.

<i>d</i> , <i>e</i>	<p>REAL for <code>chetrd</code> DOUBLE PRECISION for <code>zhetr</code>d.</p> <p>Arrays: <i>d</i>(*) contains the diagonal elements of the matrix <i>T</i>. The dimension of <i>d</i> must be at least $\max(1, n)$. <i>e</i>(*) contains the off-diagonal elements of <i>T</i>. The dimension of <i>e</i> must be at least $\max(1, n-1)$.</p>
<i>tau</i>	<p>COMPLEX for <code>chetrd</code> DOUBLE COMPLEX for <code>zhetr</code>d.</p> <p>Array, DIMENSION at least $\max(1, n-1)$. Stores further details of the unitary matrix <i>Q</i>.</p>
<i>work</i> (1)	<p>If <i>info</i> = 0, on exit <i>work</i>(1) contains the minimum value of <i>lwork</i> required for optimum performance. Use this <i>lwork</i> for subsequent runs.</p>
<i>info</i>	<p>INTEGER.</p> <p>If <i>info</i> = 0, the execution is successful. If <i>info</i> = -<i>i</i>, the <i>i</i>th parameter had an illegal value.</p>

Application Notes

For better performance, try using $lwork = n * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work*(1) and use this value for subsequent runs.

The computed matrix *T* is exactly similar to a matrix $A + E$, where $\|E\|_2 = c(n)\epsilon \|A\|_2$, $c(n)$ is a modestly increasing function of *n*, and ϵ is the machine precision.

The approximate number of floating-point operations is $(16/3)n^3$.

After calling this routine, you can call the following:

[?ungtr](#) to form the computed matrix *Q* explicitly;

[?unmtr](#) to multiply a complex matrix by *Q*.

The real counterpart of this routine is [?sytrd](#).

?ungtr

Generates the complex unitary matrix Q determined by ?hetrd.

```
call cungr ( uplo, n, a, lda, tau, work, lwork, info )
call zungr ( uplo, n, a, lda, tau, work, lwork, info )
```

Discussion

The routine explicitly generates the n by n unitary matrix Q formed by ?hetrd (see [page 5-111](#)) when reducing a complex Hermitian matrix A to tridiagonal form: $A = QTQ^H$. Use this routine after a call to ?hetrd.

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Use the same **uplo** as supplied to ?hetrd.

n INTEGER. The order of the matrix Q ($n \geq 0$).

a, tau, work COMPLEX for cungr
DOUBLE COMPLEX for zungr.
Arrays:
a(lda,*) is the array **a** as returned by ?hetrd.
The second dimension of **a** must be at least $\max(1, n)$.
tau(*) is the array **tau** as returned by ?hetrd.
The dimension of **tau** must be at least $\max(1, n-1)$.
work(lwork) is a workspace array.

lda INTEGER. The first dimension of **a**; at least $\max(1, n)$.

lwork INTEGER. The size of the **work** array ($lwork \geq n$)
See *Application notes* for the suggested value of **lwork**.

Output Parameters

a Overwritten by the unitary matrix Q .

work(1) If *info* = 0, on exit *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = (n-1) * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*.

If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

The computed matrix *Q* differs from an exactly unitary matrix by a matrix *E* such that $\|E\|_2 = O(\epsilon)$, where ϵ is the machine precision.

The approximate number of floating-point operations is $(16/3)n^3$.

The real counterpart of this routine is [?orgtr](#).

?unmtr

Multiplies a complex matrix by the complex unitary matrix Q determined by ?hetrd.

```
call cummtr ( side,uplo,trans,m,n,a,lda,tau,c,ldc,work,lwork,info )
call zummtr ( side,uplo,trans,m,n,a,lda,tau,c,ldc,work,lwork,info )
```

Discussion

The routine multiplies a complex matrix C by Q or Q^H , where Q is the unitary matrix Q formed by ?hetrd (see [page 5-111](#)) when reducing a complex Hermitian matrix A to tridiagonal form: $A = QTQ^H$. Use this routine after a call to ?hetrd.

Depending on the parameters *side* and *trans*, the routine can form one of the matrix products QC , Q^HC , CQ , or CQ^H (overwriting the result on C).

Input Parameters

In the descriptions below, r denotes the order of Q :

If *side* = 'L', $r = m$; if *side* = 'R', $r = n$.

<i>side</i>	CHARACTER*1. Must be either 'L' or 'R'. If <i>side</i> = 'L', Q or Q^H is applied to C from the left. If <i>side</i> = 'R', Q or Q^H is applied to C from the right.
<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. Use the same <i>uplo</i> as supplied to ?hetrd.
<i>trans</i>	CHARACTER*1. Must be either 'N' or 'T'. If <i>trans</i> = 'N', the routine multiplies C by Q . If <i>trans</i> = 'T', the routine multiplies C by Q^H .
<i>m</i>	INTEGER. The number of rows in the matrix C ($m \geq 0$).
<i>n</i>	INTEGER. The number of columns in C ($n \geq 0$).
<i>a,work,tau,c</i>	COMPLEX for cummtr DOUBLE COMPLEX for zummtr. <i>a(lda,*)</i> and <i>tau</i> are the arrays returned by ?hetrd.

The second dimension of a must be at least $\max(1, r)$.
The dimension of τ must be at least $\max(1, r-1)$.

$c(ldc, *)$ contains the matrix C .

The second dimension of c must be at least $\max(1, n)$

$work(lwork)$ is a workspace array.

lda INTEGER. The first dimension of a ; $lda \geq \max(1, r)$.

ldc INTEGER. The first dimension of c ; $ldc \geq \max(1, n)$.

$lwork$ INTEGER. The size of the $work$ array. Constraints:

$lwork \geq \max(1, n)$ if $side = 'L'$;

$lwork \geq \max(1, m)$ if $side = 'R'$.

See *Application notes* for the suggested value of $lwork$.

Output Parameters

c Overwritten by the product QC , $Q^H C$, CQ , or CQ^H (as specified by $side$ and $trans$).

$work(1)$ If $info = 0$, on exit $work(1)$ contains the minimum value of $lwork$ required for optimum performance. Use this $lwork$ for subsequent runs.

$info$ INTEGER.

If $info = 0$, the execution is successful.

If $info = -i$, the i th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = n * blocksize$ (for $side = 'L'$) or $lwork = m * blocksize$ (for $side = 'R'$) where $blocksize$ is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of $lwork$ for the first run. On exit, examine $work(1)$ and use this value for subsequent runs.

The computed product differs from the exact product by a matrix E such that $\|E\|_2 = O(\epsilon) \|C\|_2$, where ϵ is the machine precision.

The total number of floating-point operations is approximately $8 * m^2 * n$ if $side = 'L'$ or $8 * n^2 * m$ if $side = 'R'$.

The real counterpart of this routine is [?ormtr](#).

?sptrd

Reduces a real symmetric matrix to tridiagonal form using packed storage.

```
call ssptdr ( uplo,n,ap,d,e,tau,info )
call dsptdr ( uplo,n,ap,d,e,tau,info )
```

Discussion

This routine reduces a packed real symmetric matrix A to symmetric tridiagonal form T by an orthogonal similarity transformation: $A = QTQ^T$. The orthogonal matrix Q is not formed explicitly but is represented as a product of $n-1$ elementary reflectors. Routines are provided for working with Q in this representation. (They are described later in this section.)

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', *ap* stores the packed upper triangle of A .
 If *uplo* = 'L', *ap* stores the packed lower triangle of A .

n INTEGER. The order of the matrix A ($n \geq 0$).

ap REAL for *ssptdr*
 DOUBLE PRECISION for *dsptdr*.
 Array, DIMENSION at least $\max(1, n(n+1)/2)$.
 Contains either upper or lower triangle of A (as specified by *uplo*) in packed form.

Output Parameters

ap Overwritten by the tridiagonal matrix T and details of the orthogonal matrix Q , as specified by *uplo*.

d, *e*, *tau* REAL for *ssptdr*
 DOUBLE PRECISION for *dsptdr*.
 Arrays:
d(*) contains the diagonal elements of the matrix T .
 The dimension of *d* must be at least $\max(1, n)$.

$e(*)$ contains the off-diagonal elements of T .
The dimension of e must be at least $\max(1, n-1)$.

$tau(*)$ stores further details of the matrix Q .
The dimension of tau must be at least $\max(1, n-1)$.

$info$

INTEGER.

If $info = 0$, the execution is successful.

If $info = -i$, the i th parameter had an illegal value.

Application Notes

The computed matrix T is exactly similar to a matrix $A + E$, where $\|E\|_2 = c(n)\epsilon \|A\|_2$, $c(n)$ is a modestly increasing function of n , and ϵ is the machine precision.

The approximate number of floating-point operations is $(4/3)n^3$.

After calling this routine, you can call the following:

[?opgtr](#) to form the computed matrix Q explicitly;

[?opmtr](#) to multiply a real matrix by Q .

The complex counterpart of this routine is [?hptrd](#).

?opgtr

Generates the real orthogonal matrix Q determined by ?sptrd.

```
call sopgtr ( uplo, n, ap, tau, q, ldq, work, info )
call dogpgr ( uplo, n, ap, tau, q, ldq, work, info )
```

Discussion

The routine explicitly generates the n by n orthogonal matrix Q formed by ?sptrd (see [page 5-117](#)) when reducing a packed real symmetric matrix A to tridiagonal form: $A = QTQ^T$. Use this routine after a call to ?sptrd.

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
Use the same **uplo** as supplied to ?sptrd.

n INTEGER. The order of the matrix Q ($n \geq 0$).

ap, tau REAL for **sopgtr**
DOUBLE PRECISION for **dogpgr**.
Arrays **ap** and **tau**, as returned by ?sptrd.
The dimension of **ap** must be at least $\max(1, n(n+1)/2)$.
The dimension of **tau** must be at least $\max(1, n-1)$.

ldq INTEGER. The first dimension of the output array **q**;
at least $\max(1, n)$.

work REAL for **sopgtr**
DOUBLE PRECISION for **dogpgr**.
Workspace array, DIMENSION at least $\max(1, n-1)$.

Output Parameters

q REAL for **sopgtr**
DOUBLE PRECISION for **dogpgr**.
Array, DIMENSION (**ldq**, *).
Contains the computed matrix Q .
The second dimension of **q** must be at least $\max(1, n)$.

info

INTEGER.

If *info* = 0, the execution is successful.If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The computed matrix Q differs from an exactly orthogonal matrix by a matrix E such that $\|E\|_2 = O(\epsilon)$, where ϵ is the machine precision.

The approximate number of floating-point operations is $(4/3)n^3$.

The complex counterpart of this routine is [?upgtr](#).

?opmtr

Multiplies a real matrix by the real orthogonal matrix Q determined by [?sptrd](#).

```
call sopmtr (side,uplo,trans,m,n,ap,tau,c,ldc,work,info)
call dopmtr (side,uplo,trans,m,n,ap,tau,c,ldc,work,info)
```

Discussion

The routine multiplies a real matrix C by Q or Q^T , where Q is the orthogonal matrix Q formed by [?sptrd](#) (see [page 5-117](#)) when reducing a packed real symmetric matrix A to tridiagonal form: $A = QTQ^T$. Use this routine after a call to [?sptrd](#).

Depending on the parameters *side* and *trans*, the routine can form one of the matrix products QC , Q^TC , CQ , or CQ^T (overwriting the result on C).

Input Parameters

In the descriptions below, r denotes the order of Q :

If *side* = 'L', $r = m$; if *side* = 'R', $r = n$.

side

CHARACTER*1. Must be either 'L' or 'R'.

If *side* = 'L', Q or Q^T is applied to C from the left.If *side* = 'R', Q or Q^T is applied to C from the right.

uplo CHARACTER*1. Must be 'U' or 'L'.
Use the same *uplo* as supplied to `?sptrd`.

trans CHARACTER*1. Must be either 'N' or 'T'.
If *trans* = 'N', the routine multiplies C by Q .
If *trans* = 'T', the routine multiplies C by Q^T .

m INTEGER. The number of rows in the matrix C ($m \geq 0$).

n INTEGER. The number of columns in C ($n \geq 0$).

ap,work,tau,c REAL for `sopmtr`
DOUBLE PRECISION for `dopmtr`.
ap and *tau* are the arrays returned by `?sptrd`.
The dimension of *ap* must be at least $\max(1, r(r+1)/2)$.
The dimension of *tau* must be at least $\max(1, r-1)$.
c(ldc,)* contains the matrix C .
The second dimension of *c* must be at least $\max(1, n)$.
work()* is a workspace array.
The dimension of *work* must be at least
 $\max(1, n)$ if *side* = 'L';
 $\max(1, m)$ if *side* = 'R'.

ldc INTEGER. The first dimension of *c*; $ldc \geq \max(1, n)$.

Output Parameters

c Overwritten by the product QC , Q^TC , CQ , or CQ^T
(as specified by *side* and *trans*).

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = $-i$, the *i*th parameter had an illegal value.

Application Notes

The computed product differs from the exact product by a matrix E such that $\|E\|_2 = O(\epsilon) \|C\|_2$, where ϵ is the machine precision.

The total number of floating-point operations is approximately $2 * m^2 * n$ if *side* = 'L' or $2 * n^2 * m$ if *side* = 'R'.

The complex counterpart of this routine is `?upmtr`.

?hptrd

Reduces a complex Hermitian matrix to tridiagonal form using packed storage.

```
call chptrd ( uplo,n,ap,d,e,tau,info )
call zhptrd ( uplo,n,ap,d,e,tau,info )
```

Discussion

This routine reduces a packed complex Hermitian matrix A to symmetric tridiagonal form T by a unitary similarity transformation: $A = QTQ^H$. The unitary matrix Q is not formed explicitly but is represented as a product of $n-1$ elementary reflectors. Routines are provided for working with Q in this representation. (They are described later in this section.)

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', *ap* stores the packed upper triangle of A .
 If *uplo* = 'L', *ap* stores the packed lower triangle of A .

n INTEGER. The order of the matrix A ($n \geq 0$).

ap COMPLEX for `chptrd`
 DOUBLE COMPLEX for `zhptrd`.
 Array, DIMENSION at least $\max(1, n(n+1)/2)$.
 Contains either upper or lower triangle of A (as specified by *uplo*) in packed form.

Output Parameters

ap Overwritten by the tridiagonal matrix T and details of the orthogonal matrix Q , as specified by *uplo*.

d, e REAL for `chptrd`
 DOUBLE PRECISION for `zhptrd`.
 Arrays:
d(*) contains the diagonal elements of the matrix T .
 The dimension of *d* must be at least $\max(1, n)$.

e(*) contains the off-diagonal elements of *T*.
The dimension of *e* must be at least $\max(1, n-1)$.

tau **COMPLEX** for **chptrd**
 DOUBLE COMPLEX for **zhptrd**.
Arrays, DIMENSION at least $\max(1, n-1)$.
Contains further details of the orthogonal matrix *Q*.

info **INTEGER**.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The computed matrix *T* is exactly similar to a matrix $A + E$, where $\|E\|_2 = c(n)\epsilon \|A\|_2$, $c(n)$ is a modestly increasing function of n , and ϵ is the machine precision.

The approximate number of floating-point operations is $(16/3)n^3$.

After calling this routine, you can call the following:

[?upgtr](#) to form the computed matrix *Q* explicitly;

[?upmtr](#) to multiply a complex matrix by *Q*.

The real counterpart of this routine is [?sptrd](#).

?upgtr

Generates the complex unitary matrix Q determined by ?hptrd.

```
call cupgtr ( uplo, n, ap, tau, q, ldq, work, info )
call zupgtr ( uplo, n, ap, tau, q, ldq, work, info )
```

Discussion

The routine explicitly generates the n by n unitary matrix Q formed by ?hptrd (see [page 5-122](#)) when reducing a packed complex Hermitian matrix A to tridiagonal form: $A = QTQ^H$. Use this routine after a call to ?hptrd.

Input Parameters

<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. Use the same <i>uplo</i> as supplied to ?sptrd.
<i>n</i>	INTEGER. The order of the matrix Q ($n \geq 0$).
<i>ap, tau</i>	COMPLEX for cupgtr DOUBLE COMPLEX for zupgtr. Arrays <i>ap</i> and <i>tau</i> , as returned by ?hptrd. The dimension of <i>ap</i> must be at least $\max(1, n(n+1)/2)$. The dimension of <i>tau</i> must be at least $\max(1, n-1)$.
<i>ldq</i>	INTEGER. The first dimension of the output array <i>q</i> ; at least $\max(1, n)$.
<i>work</i>	COMPLEX for cupgtr DOUBLE COMPLEX for zupgtr. Workspace array, DIMENSION at least $\max(1, n-1)$.

Output Parameters

- q* **COMPLEX** for `cupgtr`
 DOUBLE COMPLEX for `zupgtr`.
 Array, **DIMENSION** (`ldq, *`).
 Contains the computed matrix Q .
 The second dimension of *q* must be at least $\max(1, n)$.
- info* **INTEGER**.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The computed matrix Q differs from an exactly orthogonal matrix by a matrix E such that $\|E\|_2 = O(\epsilon)$, where ϵ is the machine precision.

The approximate number of floating-point operations is $(16/3)n^3$.

The real counterpart of this routine is [?opgtr](#).

?upmtr

Multiplies a complex matrix by the unitary matrix Q determined by ?hptrd.

```
call cupmtr (side,uplo,trans,m,n,ap,tau,c,ldc,work,info)
call zupmtr (side,uplo,trans,m,n,ap,tau,c,ldc,work,info)
```

Discussion

The routine multiplies a complex matrix C by Q or Q^H , where Q is the unitary matrix Q formed by [?hptrd](#) (see [page 5-122](#)) when reducing a packed complex Hermitian matrix A to tridiagonal form: $A = QTQ^H$. Use this routine after a call to [?hptrd](#).

Depending on the parameters *side* and *trans*, the routine can form one of the matrix products QC , Q^HC , CQ , or CQ^H (overwriting the result on C).

Input Parameters

In the descriptions below, r denotes the order of Q :

If $side = 'L'$, $r = m$; if $side = 'R'$, $r = n$.

$side$	CHARACTER*1. Must be either 'L' or 'R'. If $side = 'L'$, Q or Q^H is applied to C from the left. If $side = 'R'$, Q or Q^H is applied to C from the right.
$uplo$	CHARACTER*1. Must be 'U' or 'L'. Use the same $uplo$ as supplied to <code>?hptrd</code> .
$trans$	CHARACTER*1. Must be either 'N' or 'T'. If $trans = 'N'$, the routine multiplies C by Q . If $trans = 'T'$, the routine multiplies C by Q^H .
m	INTEGER. The number of rows in the matrix C ($m \geq 0$).
n	INTEGER. The number of columns in C ($n \geq 0$).
$ap, tau, c, work$	COMPLEX for <code>cupmtr</code> DOUBLE COMPLEX for <code>zupmtr</code> . ap and tau are the arrays returned by <code>?hptrd</code> . The dimension of ap must be at least $\max(1, r(r+1)/2)$. The dimension of tau must be at least $\max(1, r-1)$. $c(ldc, *)$ contains the matrix C . The second dimension of c must be at least $\max(1, n)$. $work(*)$ is a workspace array. The dimension of $work$ must be at least $\max(1, n)$ if $side = 'L'$; $\max(1, m)$ if $side = 'R'$.
ldc	INTEGER. The first dimension of c ; $ldc \geq \max(1, n)$.

Output Parameters

c	Overwritten by the product QC , $Q^H C$, CQ , or CQ^H (as specified by $side$ and $trans$).
$info$	INTEGER. If $info = 0$, the execution is successful. If $info = -i$, the i th parameter had an illegal value.

Application Notes

The computed product differs from the exact product by a matrix E such that $\|E\|_2 = O(\epsilon) \|C\|_2$, where ϵ is the machine precision.

The total number of floating-point operations is approximately $8 * m^2 * n$ if $side = 'L'$ or $8 * n^2 * m$ if $side = 'R'$.

The real counterpart of this routine is [?opmtr](#).

?sbtrd

Reduces a real symmetric band matrix to tridiagonal form.

```
call ssbtrd (vect,uplo,n,kd,ab,ldab,d,e,q,ldq,work,info)
call dsbtrd (vect,uplo,n,kd,ab,ldab,d,e,q,ldq,work,info)
```

Discussion

This routine reduces a real symmetric band matrix A to symmetric tridiagonal form T by an orthogonal similarity transformation: $A = QTQ^T$. The orthogonal matrix Q is determined as a product of Givens rotations. If required, the routine can also form the matrix Q explicitly.

Input Parameters

vect CHARACTER*1. Must be 'V' or 'N'.
If *vect* = 'V', the routine returns the explicit matrix Q .
If *vect* = 'N', the routine does not return Q .

uplo CHARACTER*1. Must be 'U' or 'L'.
If *uplo* = 'U', *ab* stores the upper triangular part of A .
If *uplo* = 'L', *ab* stores the lower triangular part of A .

n INTEGER. The order of the matrix A ($n \geq 0$).

kd INTEGER. The number of super- or sub-diagonals in A ($kd \geq 0$).

ab, work REAL for *ssbtrd*
DOUBLE PRECISION for *dsbtrd*.
ab (*ldab*,*) is an array containing either upper or lower triangular part of the matrix A (as specified by *uplo*) in band storage format.
The second dimension of *ab* must be at least $\max(1, n)$.
work (*) is a workspace array.
The dimension of *work* must be at least $\max(1, n)$.

ldab INTEGER. The first dimension of *ab*; at least $kd+1$.

ldq **INTEGER**. The first dimension of *q*. Constraints:
 ldq ≥ max(1, *n*) if *vect* = 'V';
 ldq ≥ 1 if *vect* = 'N'.

Output Parameters

ab On exit, the array *ab* is overwritten.

d, *e*, *q* **REAL** for *ssbtrd*
 DOUBLE PRECISION for *dsbtrd*.
 Arrays:
d(*) contains the diagonal elements of the matrix *T*.
 The dimension of *d* must be at least max(1, *n*).
e(*) contains the off-diagonal elements of *T*.
 The dimension of *e* must be at least max(1, *n*-1).
q(*ldq*,*) is not referenced if *vect* = 'N'.
 If *vect* = 'V', *q* contains the *n* by *n* matrix *Q*.
 The second dimension of *q* must be:
 at least max(1, *n*) if *vect* = 'V';
 at least 1 if *vect* = 'N'.

info **INTEGER**.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The computed matrix *T* is exactly similar to a matrix *A* + *E*, where $\|E\|_2 = c(n)\epsilon \|A\|_2$, *c*(*n*) is a modestly increasing function of *n*, and ϵ is the machine precision. The computed matrix *Q* differs from an exactly orthogonal matrix by a matrix *E* such that $\|E\|_2 = O(\epsilon)$.

The total number of floating-point operations is approximately $6n^2 * kd$ if *vect* = 'N', with $3n^3 * (kd-1) / kd$ additional operations if *vect* = 'V'.

The complex counterpart of this routine is [?hbtrd](#).

?hbtrd

Reduces a complex Hermitian band matrix to tridiagonal form.

```
call chbtrd (vect,uplo,n,kd,ab,ldab,d,e,q,ldq,work,info)
call zhbtrd (vect,uplo,n,kd,ab,ldab,d,e,q,ldq,work,info)
```

Discussion

This routine reduces a complex Hermitian band matrix A to symmetric tridiagonal form T by a unitary similarity transformation: $A = QTQ^H$. The unitary matrix Q is determined as a product of Givens rotations. If required, the routine can also form the matrix Q explicitly.

Input Parameters

<i>vect</i>	CHARACTER*1. Must be 'V' or 'N'. If <i>vect</i> = 'V', the routine returns the explicit matrix Q . If <i>vect</i> = 'N', the routine does not return Q .
<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. If <i>uplo</i> = 'U', <i>ab</i> stores the upper triangular part of A . If <i>uplo</i> = 'L', <i>ab</i> stores the lower triangular part of A .
<i>n</i>	INTEGER. The order of the matrix A ($n \geq 0$).
<i>kd</i>	INTEGER. The number of super- or sub-diagonals in A ($kd \geq 0$).
<i>ab, work</i>	COMPLEX for chbtrd DOUBLE COMPLEX for zhbtrd. <i>ab</i> (<i>ldab</i> ,*) is an array containing either upper or lower triangular part of the matrix A (as specified by <i>uplo</i>) in band storage format. The second dimension of <i>ab</i> must be at least $\max(1, n)$. <i>work</i> (*) is a workspace array. The dimension of <i>work</i> must be at least $\max(1, n)$.
<i>ldab</i>	INTEGER. The first dimension of <i>ab</i> ; at least $kd+1$.

ldq **INTEGER**. The first dimension of *q*. Constraints:
ldq ≥ max(1, *n*) if *vect* = 'V';
ldq ≥ 1 if *vect* = 'N'.

Output Parameters

ab On exit, the array *ab* is overwritten.

d, *e* **REAL** for **chbtrd**
DOUBLE PRECISION for **zhbtrd**.
 Arrays:
d(*) contains the diagonal elements of the matrix *T*.
 The dimension of *d* must be at least max(1, *n*).
e(*) contains the off-diagonal elements of *T*.
 The dimension of *e* must be at least max(1, *n*-1).

q **COMPLEX** for **chbtrd**
DOUBLE COMPLEX for **zhbtrd**.
 Array, **DIMENSION** (*ldq*,*).
 If *vect* = 'N', *q* is not referenced.
 If *vect* = 'V', *q* contains the *n* by *n* matrix *Q*.
 The second dimension of *q* must be:
 at least max(1, *n*) if *vect* = 'V';
 at least 1 if *vect* = 'N'.

info **INTEGER**.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The computed matrix *T* is exactly similar to a matrix *A* + *E*, where $\| | E | \|_2 = c(n)\epsilon$ $\| | A | \|_2$, *c*(*n*) is a modestly increasing function of *n*, and ϵ is the machine precision. The computed matrix *Q* differs from an exactly unitary matrix by a matrix *E* such that $\| | E | \|_2 = O(\epsilon)$.

The total number of floating-point operations is approximately $20n^2 * kd$ if *vect* = 'N', with $10n^3 * (kd-1)/kd$ additional operations if *vect* = 'V'.

The real counterpart of this routine is [?sbtrd](#).

?sterf

Computes all eigenvalues of a real symmetric tridiagonal matrix using *QR* algorithm.

```
call ssterf ( n, d, e, info )
call dsterf ( n, d, e, info )
```

Discussion

This routine computes all the eigenvalues of a real symmetric tridiagonal matrix T (which can be obtained by reducing a symmetric or Hermitian matrix to tridiagonal form). The routine uses a square-root-free variant of the *QR* algorithm.

If you need not only the eigenvalues but also the eigenvectors, call [?stegr](#) ([page 5-134](#)).

Input Parameters

n **INTEGER**. The order of the matrix T ($n \geq 0$).

d, e **REAL** for **ssterf**
DOUBLE PRECISION for **dsterf**.

Arrays:

$d(*)$ contains the diagonal elements of T .
The dimension of d must be at least $\max(1, n)$.

$e(*)$ contains the off-diagonal elements of T .
The dimension of e must be at least $\max(1, n-1)$.

Output Parameters

d The n eigenvalues in ascending order, unless $info > 0$.
See also *info*.

e On exit, the array is overwritten; see *info*.

info

INTEGER.

If *info* = 0, the execution is successful.If *info* = *i*, the algorithm failed to find all the eigenvalues after 30*n* iterations: *i* off-diagonal elements have not converged to zero. On exit, *d* and *e* contain, respectively, the diagonal and off-diagonal elements of a tridiagonal matrix orthogonally similar to *T*.If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The computed eigenvalues and eigenvectors are exact for a matrix $T + E$ such that $\|E\|_2 = O(\epsilon) \|T\|_2$, where ϵ is the machine precision.

If λ_i is an exact eigenvalue, and μ_i is the corresponding computed value, then

$$|\mu_i - \lambda_i| \leq c(n)\epsilon \|T\|_2$$

where $c(n)$ is a modestly increasing function of n .

The total number of floating-point operations depends on how rapidly the algorithm converges. Typically, it is about $14n^2$.

?steqr

Computes all eigenvalues and eigenvectors of a symmetric or Hermitian matrix reduced to tridiagonal form (QR algorithm).

```
call ssteqr ( compz, n, d, e, z, ldz, work, info )
call dsteqr ( compz, n, d, e, z, ldz, work, info )
call csteqr ( compz, n, d, e, z, ldz, work, info )
call zsteqr ( compz, n, d, e, z, ldz, work, info )
```

Discussion

This routine computes all the eigenvalues and (optionally) all the eigenvectors of a real symmetric tridiagonal matrix T . In other words, the routine can compute the spectral factorization: $T = Z\Lambda Z^T$.

Here Λ is a diagonal matrix whose diagonal elements are the eigenvalues λ_i ; Z is an orthogonal matrix whose columns are eigenvectors. Thus,

$$Tz_i = \lambda_i z_i \text{ for } i = 1, 2, \dots, n.$$

(The routine normalizes the eigenvectors so that $\|z_i\|_2 = 1$.)

You can also use the routine for computing the eigenvalues and eigenvectors of an arbitrary real symmetric (or complex Hermitian) matrix A reduced to tridiagonal form $T: A = QTQ^H$. In this case, the spectral factorization is as follows: $A = QTQ^H = (QZ)\Lambda(QZ)^H$. Before calling `?steqr`, you must reduce A to tridiagonal form and generate the explicit matrix Q by calling the following routines:

	for real matrices:	for complex matrices:
full storage	<code>?sytrd, ?orgtr</code>	<code>?hetrd, ?ungtr</code>
packed storage	<code>?sptrd, ?opgtr</code>	<code>?hptrd, ?upgtr</code>
band storage	<code>?sbtrd (vect='V')</code>	<code>?hbtrd (vect='V')</code>

If you need eigenvalues only, it's more efficient to call `?sterf` ([page 5-132](#)). If T is positive-definite, `?pteqr` ([page 5-146](#)) can compute small eigenvalues more accurately than `?steqr`.

To solve the problem by a single call, use one of the divide and conquer routines `?stevd`, `?syevd`, `?spevd`, or `?sbevd` for real symmetric matrices or `?heevd`, `?hpevd`, or `?hbevd` for complex Hermitian matrices.

Input Parameters

<i>compz</i>	<p>CHARACTER*1. Must be 'N' or 'I' or 'V'.</p> <p>If <i>compz</i> = 'N', the routine computes eigenvalues only.</p> <p>If <i>compz</i> = 'I', the routine computes the eigenvalues and eigenvectors of the tridiagonal matrix <i>T</i>.</p> <p>If <i>compz</i> = 'V', the routine computes the eigenvalues and eigenvectors of <i>A</i> (and the array <i>z</i> must contain the matrix <i>Q</i> on entry).</p>
<i>n</i>	<p>INTEGER. The order of the matrix <i>T</i> ($n \geq 0$).</p>
<i>d,e,work</i>	<p>REAL for single-precision flavors</p> <p>DOUBLE PRECISION for double-precision flavors.</p> <p>Arrays:</p> <p><i>d</i>(*) contains the diagonal elements of <i>T</i>. The dimension of <i>d</i> must be at least $\max(1, n)$.</p> <p><i>e</i>(*) contains the off-diagonal elements of <i>T</i>. The dimension of <i>e</i> must be at least $\max(1, n-1)$.</p> <p><i>work</i>(*) is a workspace array. The dimension of <i>work</i> must be: at least 1 if <i>compz</i> = 'N'; at least $\max(1, 2*n-2)$ if <i>compz</i> = 'V' or 'I'.</p>
<i>z</i>	<p>REAL for <i>ssteqr</i></p> <p>DOUBLE PRECISION for <i>dsteqr</i></p> <p>COMPLEX for <i>csteqr</i></p> <p>DOUBLE COMPLEX for <i>zsteqr</i>.</p> <p>Array, DIMENSION (<i>ldz</i>, *)</p> <p>If <i>compz</i> = 'N' or 'I', <i>z</i> need not be set.</p> <p>If <i>vect</i> = 'V', <i>z</i> must contain the <i>n</i> by <i>n</i> matrix <i>Q</i>. The second dimension of <i>z</i> must be: at least 1 if <i>compz</i> = 'N'; at least $\max(1, n)$ if <i>compz</i> = 'V' or 'I'.</p> <p><i>work</i>(<i>lwork</i>) is a workspace array.</p>
<i>ldz</i>	<p>INTEGER. The first dimension of <i>z</i>. Constraints:</p> <p><i>ldz</i> \geq 1 if <i>compz</i> = 'N';</p> <p><i>ldz</i> \geq $\max(1, n)$ if <i>compz</i> = 'V' or 'I'.</p>

Output Parameters

<i>d</i>	The <i>n</i> eigenvalues in ascending order, unless <i>info</i> > 0. See also <i>info</i> .
<i>e</i>	On exit, the array is overwritten; see <i>info</i> .
<i>z</i>	If <i>info</i> = 0, contains the <i>n</i> orthonormal eigenvectors, stored by columns. (The <i>i</i> th column corresponds to the <i>i</i> th eigenvalue.)
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = <i>i</i> , the algorithm failed to find all the eigenvalues after 30 <i>n</i> iterations: <i>i</i> off-diagonal elements have not converged to zero. On exit, <i>d</i> and <i>e</i> contain, respectively, the diagonal and off-diagonal elements of a tridiagonal matrix orthogonally similar to <i>T</i> . If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value.

Application Notes

The computed eigenvalues and eigenvectors are exact for a matrix $T + E$ such that $\|E\|_2 = O(\epsilon) \|T\|_2$, where ϵ is the machine precision.

If λ_i is an exact eigenvalue, and μ_i is the corresponding computed value, then

$$|\mu_i - \lambda_i| \leq c(n)\epsilon \|T\|_2$$

where $c(n)$ is a modestly increasing function of n .

If z_i is the corresponding exact eigenvector, and w_i is the corresponding computed vector, then the angle $\theta(z_i, w_i)$ between them is bounded as follows:

$$\theta(z_i, w_i) \leq c(n)\epsilon \|T\|_2 / \min_{i \neq j} |\lambda_i - \lambda_j|.$$

The total number of floating-point operations depends on how rapidly the algorithm converges. Typically, it is about

$$\begin{aligned} &24n^2 \text{ if } \textit{compz} = \text{'N'}; \\ &7n^3 \text{ (for complex flavors, } 14n^3) \text{ if } \textit{compz} = \text{'V'} \text{ or } \text{'I'}. \end{aligned}$$

?stedc

Computes all eigenvalues and eigenvectors of a symmetric tridiagonal matrix using the divide and conquer method.

```
call sstedc(compz, n, d, e, z, ldz, work, lwork, iwork, liwork, info)
call dstedc(compz, n, d, e, z, ldz, work, lwork, iwork, liwork, info)
call cstedc(compz, n, d, e, z, ldz, work, lwork, rwork, lrwork,
            iwork, liwork, info)
call zstedc(compz, n, d, e, z, ldz, work, lwork, rwork, lrwork,
            iwork, liwork, info)
```

Discussion

This routine computes all the eigenvalues and (optionally) all the eigenvectors of a symmetric tridiagonal matrix using the divide and conquer method.

The eigenvectors of a full or band real symmetric or complex Hermitian matrix can also be found if `?sytrd/?hetrd` or `?sptrd/?hptrd` or `?sbtrd/?hbtrd` has been used to reduce this matrix to tridiagonal form.

Input Parameters

`compz` CHARACTER*1. Must be 'N' or 'I' or 'V'.
If `compz = 'N'`, the routine computes eigenvalues only.
If `compz = 'I'`, the routine computes the eigenvalues and eigenvectors of the tridiagonal matrix.
If `compz = 'V'`, the routine computes the eigenvalues and eigenvectors of original symmetric/Hermitian matrix. On entry, the array `z` must contain the orthogonal/unitary matrix used to reduce the original matrix to tridiagonal form.

`n` INTEGER. The order of the symmetric tridiagonal matrix ($n \geq 0$).

d, *e*, *rwork* REAL for single-precision flavors
DOUBLE PRECISION for double-precision flavors.
Arrays:
d(*) contains the diagonal elements of the tridiagonal matrix. The dimension of *d* must be at least $\max(1, n)$.
e(*) contains the subdiagonal elements of the tridiagonal matrix. The dimension of *e* must be at least $\max(1, n-1)$.
rwork(*lrwork*) is a workspace array used in complex flavors only.

z, *work* REAL for *sstedc*
DOUBLE PRECISION for *dstedc*
COMPLEX for *cstedc*
DOUBLE COMPLEX for *zstedc*.
Arrays: *z*(*ldz*, *), *work*(*).
If *compz* = 'V', then, on entry, *z* must contain the orthogonal/unitary matrix used to reduce the original matrix to tridiagonal form.
The second dimension of *z* must be at least $\max(1, n)$.
work(*lwork*) is a workspace array.

ldz INTEGER. The first dimension of *z*. Constraints:
 $ldz \geq 1$ if *compz* = 'N';
 $ldz \geq \max(1, n)$ if *compz* = 'V' or 'I'.

lwork INTEGER. The dimension of the array *work*.
See *Application Notes* for the required value of *lwork*.

lrwork INTEGER. The dimension of the array *rwork* (used for complex flavors only).
See *Application Notes* for the required value of *lrwork*.

iwork INTEGER. Workspace array, DIMENSION (*liwork*).

liwork INTEGER. The dimension of the array *iwork*.
See *Application Notes* for the required value of *liwork*.

Output Parameters

<i>d</i>	The <i>n</i> eigenvalues in ascending order, unless <i>info</i> ≠ 0. See also <i>info</i> .
<i>e</i>	On exit, the array is overwritten; see <i>info</i> .
<i>z</i>	If <i>info</i> = 0, then if <i>compz</i> = 'V', <i>z</i> contains the orthonormal eigenvectors of the original symmetric/Hermitian matrix, and if <i>compz</i> = 'I', <i>z</i> contains the orthonormal eigenvectors of the symmetric tridiagonal matrix. If <i>compz</i> = 'N', <i>z</i> is not referenced.
<i>work(1)</i>	On exit, if <i>info</i> = 0, then <i>work(1)</i> returns the optimal <i>lwork</i> .
<i>rwork(1)</i>	On exit, if <i>info</i> = 0, then <i>rwork(1)</i> returns the optimal <i>lrwork</i> (for complex flavors only).
<i>iwork(1)</i>	On exit, if <i>info</i> = 0, then <i>iwork(1)</i> returns the optimal <i>liwork</i> .
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = <i>i</i> , the algorithm failed to compute an eigenvalue while working on the submatrix lying in rows and columns <i>i</i> /(<i>n</i> +1) through mod(<i>i</i> , <i>n</i> +1).

Application Notes

The required size of workspace arrays must be as follows.

For *sstedc/dstedc*:

If *compz* = 'N' or *n* ≤ 1 then *lwork* must be at least 1.

If *compz* = 'V' and *n* > 1 then *lwork* must be at least $(1 + 3n + 2n \cdot \lg n + 3n^2)$, where $\lg(n)$ = smallest integer *k* such that $2^k \geq n$.

If *compz* = 'I' and *n* > 1 then *lwork* must be at least $(1 + 4n + n^2)$.

If *compz* = 'N' or *n* ≤ 1 then *liwork* must be at least 1.

If *compz* = 'V' and *n* > 1 then *liwork* must be at least $(6 + 6n + 5n \cdot \lg n)$.

If *compz* = 'I' and *n* > 1 then *liwork* must be at least $(3 + 5n)$.

For *cstedc/zstedc*:

If *compz* = 'N' or 'I', or $n \leq 1$, *lwork* must be at least 1.

If *compz* = 'V' and $n > 1$, *lwork* must be at least n^2 .

If *compz* = 'N' or $n \leq 1$, *lrwork* must be at least 1.

If *compz* = 'V' and $n > 1$, *lrwork* must be at least

$(1 + 3n + 2n \cdot \lg n + 3n^2)$, where $\lg(n)$ = smallest integer k such that $2^k \geq n$.

If *compz* = 'I' and $n > 1$, *lrwork* must be at least $(1 + 4n + 2n^2)$.

The required value of *liwork* for complex flavors is the same as for real flavors.

?stegr

Computes selected eigenvalues and eigenvectors of a real symmetric tridiagonal matrix.

```
call sstegr (jobz, range, n, d, e, vl, vu, il, iu, abstol, m, w, z,
            ldz, isuppz, work, lwork, iwork, liwork, info)
call dstegr (jobz, range, n, d, e, vl, vu, il, iu, abstol, m, w, z,
            ldz, isuppz, work, lwork, iwork, liwork, info)
call cstegr (jobz, range, n, d, e, vl, vu, il, iu, abstol, m, w, z,
            ldz, isuppz, work, lwork, iwork, liwork, info)
call zstegr (jobz, range, n, d, e, vl, vu, il, iu, abstol, m, w, z,
            ldz, isuppz, work, lwork, iwork, liwork, info)
```

Discussion

This routine computes selected eigenvalues and, optionally, eigenvectors of a real symmetric tridiagonal matrix T . Eigenvalues and eigenvectors can be selected by specifying either a range of values or a range of indices for the desired eigenvalues. The eigenvalues are computed by the *dqds* algorithm, while orthogonal eigenvectors are computed from various “good” LDL^T representations (also known as Relatively Robust Representations). Gram-Schmidt orthogonalization is avoided as far as possible. More specifically, the various steps of the algorithm are as follows. For the i -th unreduced block of T ,

- (a) Compute $T - \mathfrak{Q}_i = L_i D_i L_i^T$, such that $L_i D_i L_i^T$ is a relatively robust representation;
- (b) Compute the eigenvalues, λ_j , of $L_i D_i L_i^T$ to high relative accuracy by the *dqds* algorithm;
- (c) If there is a cluster of close eigenvalues, “choose” \mathfrak{Q}_i close to the cluster, and go to step (a);
- (d) Given the approximate eigenvalue λ_j of $L_i D_i L_i^T$, compute the corresponding eigenvector by forming a rank-revealing twisted factorization.

The desired accuracy of the output can be specified by the input parameter *abstol*.

Input Parameters

<i>jobz</i>	<p>CHARACTER*1. Must be 'N' or 'V'.</p> <p>If <i>job</i> = 'N', then only eigenvalues are computed.</p> <p>If <i>job</i> = 'V', then eigenvalues and eigenvectors are computed.</p>
<i>range</i>	<p>CHARACTER*1. Must be 'A' or 'V' or 'I'.</p> <p>If <i>range</i> = 'A', the routine computes all eigenvalues.</p> <p>If <i>range</i> = 'V', the routine computes eigenvalues λ_i in the half-open interval: $vl < \lambda_i \leq vu$.</p> <p>If <i>range</i> = 'I', the routine computes eigenvalues with indices <i>il</i> to <i>iu</i>.</p>
<i>n</i>	<p>INTEGER. The order of the matrix <i>T</i> ($n \geq 0$).</p>
<i>d, e, work</i>	<p>REAL for single precision flavors</p> <p>DOUBLE PRECISION for double precision flavors.</p> <p>Arrays:</p> <p><i>d</i>(*) contains the diagonal elements of <i>T</i>.</p> <p>The dimension of <i>d</i> must be at least $\max(1, n)$.</p> <p><i>e</i>(*) contains the subdiagonal elements of <i>T</i> in elements 1 to <i>n</i>-1; <i>e</i>(<i>n</i>) need not be set.</p> <p>The dimension of <i>e</i> must be at least $\max(1, n)$.</p> <p><i>work</i>(<i>lwork</i>) is a workspace array.</p>
<i>vl, vu</i>	<p>REAL for single precision flavors</p> <p>DOUBLE PRECISION for double precision flavors.</p> <p>If <i>range</i> = 'V', the lower and upper bounds of the interval to be searched for eigenvalues.</p> <p>Constraint: $vl < vu$.</p> <p>If <i>range</i> = 'A' or 'I', <i>vl</i> and <i>vu</i> are not referenced.</p>
<i>il, iu</i>	<p>INTEGER.</p> <p>If <i>range</i> = 'I', the indices in ascending order of the smallest and largest eigenvalues to be returned.</p> <p>Constraint: $1 \leq il \leq iu \leq n$, if $n > 0$; <i>il</i>=1 and <i>iu</i>=0 if $n = 0$.</p>

If *range* = 'A' or 'V', *il* and *iu* are not referenced.

abstol REAL for single precision flavors
DOUBLE PRECISION for double precision flavors.
The absolute tolerance to which each eigenvalue/eigenvector is required.
If *jobz* = 'V', the eigenvalues and eigenvectors output have residual norms bounded by *abstol*, and the dot products between different eigenvectors are bounded by *abstol*. If $abstol < n\epsilon ||T||_1$, then $n\epsilon ||T||_1$ will be used in its place, where ϵ is the machine precision. The eigenvalues are computed to an accuracy of $\epsilon ||T||_1$ irrespective of *abstol*. If high relative accuracy is important, set *abstol* to *?lamch* ('Safe minimum').

ldz INTEGER. The leading dimension of the output array *z*.
Constraints:
 $ldz \geq 1$ if *jobz* = 'N';
 $ldz \geq \max(1, n)$ if *jobz* = 'V'.

lwork INTEGER. The dimension of the array *work*,
 $lwork \geq \max(1, 18n)$.

iwork INTEGER.
Workspace array, DIMENSION (*liwork*).

liwork INTEGER. The dimension of the array *iwork*,
 $lwork \geq \max(1, 10n)$.

Output Parameters

d, *e* On exit, *d* and *e* are overwritten.

m INTEGER. The total number of eigenvalues found,
 $0 \leq m \leq n$. If *range* = 'A', $m = n$, and if *range* = 'I',
 $m = iu - il + 1$.

w REAL for single precision flavors
DOUBLE PRECISION for double precision flavors.
Array, DIMENSION at least $\max(1, n)$.
The selected eigenvalues in ascending order, stored in *w*(1) to *w*(*m*).

<i>z</i>	<p>REAL for <code>sstegr</code> DOUBLE PRECISION for <code>dstegr</code> COMPLEX for <code>cstegr</code> DOUBLE COMPLEX for <code>zstegr</code>.</p> <p>Array <code>z(ldz, *)</code>, the second dimension of <code>z</code> must be at least $\max(1, m)$.</p> <p>If <code>jobz = 'V'</code>, then if <code>info = 0</code>, the first m columns of <code>z</code> contain the orthonormal eigenvectors of the matrix T corresponding to the selected eigenvalues, with the i-th column of <code>z</code> holding the eigenvector associated with $w(i)$. If <code>jobz = 'N'</code>, then <code>z</code> is not referenced.</p> <p>Note: you must ensure that at least $\max(1, m)$ columns are supplied in the array <code>z</code>; if <code>range = 'V'</code>, the exact value of m is not known in advance and an upper bound must be used.</p>
<i>isuppz</i>	<p>INTEGER.</p> <p>Array, DIMENSION at least $2 * \max(1, m)$.</p> <p>The support of the eigenvectors in <code>z</code>, i.e., the indices indicating the nonzero elements in <code>z</code>. The i-th eigenvector is nonzero only in elements <code>isuppz(2i-1)</code> through <code>isuppz(2i)</code>.</p>
<i>work(1)</i>	<p>On exit, if <code>info = 0</code>, then <code>work(1)</code> returns the required minimal size of <code>lwork</code>.</p>
<i>iwork(1)</i>	<p>On exit, if <code>info = 0</code>, then <code>iwork(1)</code> returns the required minimal size of <code>liwork</code>.</p>
<i>info</i>	<p>INTEGER.</p> <p>If <code>info = 0</code>, the execution is successful.</p> <p>If <code>info = -i</code>, the ith parameter had an illegal value.</p> <p>If <code>info = 1</code>, internal error in <code>slarre</code> occurred,</p> <p>If <code>info = 2</code>, internal error in <code>?larrv</code> occurred.</p>

Application Notes

Currently `?stegr` is only set up to find *all* the n eigenvalues and eigenvectors of T in $O(n^2)$ time, that is, only `range = 'A'` is supported.

Currently the routine `?stein` is called when an appropriate \mathfrak{G} cannot be chosen in step (c) above. `?stein` invokes modified Gram-Schmidt when eigenvalues are close.

`?stegr` works only on machines which follow IEEE-754 floating-point standard in their handling of infinities and NaNs. Normal execution of `?stegr` may create NaNs and infinities and hence may abort due to a floating point exception in environments which do not conform to the IEEE-754 standard.

?pteqr

Computes all eigenvalues and (optionally) all eigenvectors of a real symmetric positive-definite tridiagonal matrix.

```
call spteqr ( compz, n, d, e, z, ldz, work, info )
call dpteqr ( compz, n, d, e, z, ldz, work, info )
call cpteqr ( compz, n, d, e, z, ldz, work, info )
call zpteqr ( compz, n, d, e, z, ldz, work, info )
```

Discussion

This routine computes all the eigenvalues and (optionally) all the eigenvectors of a real symmetric positive-definite tridiagonal matrix T . In other words, the routine can compute the spectral factorization: $T = Z\Lambda Z^T$. Here Λ is a diagonal matrix whose diagonal elements are the eigenvalues λ_i ; Z is an orthogonal matrix whose columns are eigenvectors. Thus,

$$Tz_i = \lambda_i z_i \text{ for } i = 1, 2, \dots, n.$$

(The routine normalizes the eigenvectors so that $\|z_i\|_2 = 1$.)

You can also use the routine for computing the eigenvalues and eigenvectors of real symmetric (or complex Hermitian) positive-definite matrices A reduced to tridiagonal form T : $A = QTQ^H$. In this case, the spectral factorization is as follows: $A = QTQ^H = (QZ)\Lambda(QZ)^H$. Before calling `?pteqr`, you must reduce A to tridiagonal form and generate the explicit matrix Q by calling the following routines:

	for real matrices:	for complex matrices:
full storage	<code>?sytrd, ?orgtr</code>	<code>?hetrd, ?ungtr</code>
packed storage	<code>?sptrd, ?opgtr</code>	<code>?hptrd, ?upgtr</code>
band storage	<code>?sbtrd (vect='V')</code>	<code>?hbtrd (vect='V')</code>

The routine first factorizes T as LDL^H where L is a unit lower bidiagonal matrix, and D is a diagonal matrix. Then it forms the bidiagonal matrix $B = LD^{1/2}$ and calls `?bdsqr` to compute the singular values of B , which are the same as the eigenvalues of T .

Input Parameters

<i>compz</i>	<p>CHARACTER*1. Must be 'N' or 'I' or 'V'.</p> <p>If <i>compz</i> = 'N', the routine computes eigenvalues only.</p> <p>If <i>compz</i> = 'I', the routine computes the eigenvalues and eigenvectors of the tridiagonal matrix <i>T</i>.</p> <p>If <i>compz</i> = 'V', the routine computes the eigenvalues and eigenvectors of <i>A</i> (and the array <i>z</i> must contain the matrix <i>Q</i> on entry).</p>
<i>n</i>	<p>INTEGER. The order of the matrix <i>T</i> ($n \geq 0$).</p>
<i>d,e,work</i>	<p>REAL for single-precision flavors</p> <p>DOUBLE PRECISION for double-precision flavors.</p> <p>Arrays:</p> <p><i>d</i>(*) contains the diagonal elements of <i>T</i>. The dimension of <i>d</i> must be at least $\max(1, n)$.</p> <p><i>e</i>(*) contains the off-diagonal elements of <i>T</i>. The dimension of <i>e</i> must be at least $\max(1, n-1)$.</p> <p><i>work</i>(*) is a workspace array. The dimension of <i>work</i> must be: at least 1 if <i>compz</i> = 'N'; at least $\max(1, 4*n-4)$ if <i>compz</i> = 'V' or 'I'.</p>
<i>z</i>	<p>REAL for <i>spteqr</i></p> <p>DOUBLE PRECISION for <i>dpteqr</i></p> <p>COMPLEX for <i>cpteqr</i></p> <p>DOUBLE COMPLEX for <i>zpteqr</i>.</p> <p>Array, DIMENSION (<i>ldz</i>, *)</p> <p>If <i>compz</i> = 'N' or 'I', <i>z</i> need not be set.</p> <p>If <i>vect</i> = 'V', <i>z</i> must contain the <i>n</i> by <i>n</i> matrix <i>Q</i>. The second dimension of <i>z</i> must be: at least 1 if <i>compz</i> = 'N'; at least $\max(1, n)$ if <i>compz</i> = 'V' or 'I'.</p>
<i>ldz</i>	<p>INTEGER. The first dimension of <i>z</i>. Constraints:</p> <p><i>ldz</i> \geq 1 if <i>compz</i> = 'N';</p> <p><i>ldz</i> \geq $\max(1, n)$ if <i>compz</i> = 'V' or 'I'.</p>

Output Parameters

<i>d</i>	The <i>n</i> eigenvalues in descending order, unless <i>info</i> > 0. See also <i>info</i> .
<i>e</i>	On exit, the array is overwritten.
<i>z</i>	If <i>info</i> = 0, contains the <i>n</i> orthonormal eigenvectors, stored by columns. (The <i>i</i> th column corresponds to the <i>i</i> th eigenvalue.)
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = <i>i</i> , the leading minor of order <i>i</i> (and hence <i>T</i> itself) is not positive-definite. If <i>info</i> = <i>n</i> + <i>i</i> , the algorithm for computing singular values failed to converge; <i>i</i> off-diagonal elements have not converged to zero. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value.

Application Notes

If λ_i is an exact eigenvalue, and μ_i is the corresponding computed value, then

$$|\mu_i - \lambda_i| \leq c(n)\varepsilon K \lambda_i$$

where $c(n)$ is a modestly increasing function of n , ε is the machine precision, and $K = \| |DTD| \|_2 \| (DTD)^{-1} \|_2$, D is diagonal with $d_{ii} = t_{ii}^{-1/2}$.

If z_i is the corresponding exact eigenvector, and w_i is the corresponding computed vector, then the angle $\theta(z_i, w_i)$ between them is bounded as follows:

$$\theta(z_i, w_i) \leq c(n)\varepsilon K / \min_{i \neq j} (|\lambda_i - \lambda_j| / |\lambda_i + \lambda_j|).$$

Here $\min_{i \neq j} (|\lambda_i - \lambda_j| / |\lambda_i + \lambda_j|)$ is the *relative gap* between λ_i and the other eigenvalues.

The total number of floating-point operations depends on how rapidly the algorithm converges. Typically, it is about

$$\begin{aligned} & 30n^2 \text{ if } \textit{compz} = \text{'N'}; \\ & 6n^3 \text{ (for complex flavors, } 12n^3 \text{) if } \textit{compz} = \text{'V'} \text{ or } \text{'I'}. \end{aligned}$$

?stebz

Computes selected eigenvalues of a real symmetric tridiagonal matrix by bisection.

```
call sstebz (range, order, n, vl, vu, il, iu, abstol,
            d, e, m, nsplit, w, iblock, isplit, work, iwork, info)
call dstebz (range, order, n, vl, vu, il, iu, abstol,
            d, e, m, nsplit, w, iblock, isplit, work, iwork, info)
```

Discussion

This routine computes some (or all) of the eigenvalues of a real symmetric tridiagonal matrix T by bisection. The routine searches for zero or negligible off-diagonal elements to see if T splits into block-diagonal form $T = \text{diag}(T_1, T_2, \dots)$. Then it performs bisection on each of the blocks T_i and returns the block index of each computed eigenvalue, so that a subsequent call to `?stein` can also take advantage of the block structure.

Input Parameters

range CHARACTER*1. Must be 'A' or 'V' or 'I'.
 If *range* = 'A', the routine computes all eigenvalues.
 If *range* = 'V', the routine computes eigenvalues λ_i in the half-open interval: $vl < \lambda_i \leq vu$.
 If *range* = 'I', the routine computes eigenvalues with indices *il* to *iu*.

order CHARACTER*1. Must be 'B' or 'E'.
 If *order* = 'B', the eigenvalues are to be ordered from smallest to largest within each split-off block.
 If *order* = 'E', the eigenvalues for the entire matrix are to be ordered from smallest to largest.

n INTEGER. The order of the matrix T ($n \geq 0$).

<i>vl, vu</i>	<p>REAL for <i>sstebz</i> DOUBLE PRECISION for <i>dstebz</i>.</p> <p>If <i>range</i> = 'V', the routine computes eigenvalues λ_i in the half-open interval: $vl < \lambda_i \leq vu$.</p> <p>If <i>range</i> = 'A' or 'I', <i>vl</i> and <i>vu</i> are not referenced.</p>
<i>il, iu</i>	<p>INTEGER. Constraint: $1 \leq il \leq iu \leq n$.</p> <p>If <i>range</i> = 'I', the routine computes eigenvalues λ_i such that $il \leq i \leq iu$ (assuming that the eigenvalues λ_i are in ascending order).</p> <p>If <i>range</i> = 'A' or 'V', <i>il</i> and <i>iu</i> are not referenced.</p>
<i>abstol</i>	<p>REAL for <i>sstebz</i> DOUBLE PRECISION for <i>dstebz</i>.</p> <p>The absolute tolerance to which each eigenvalue is required. An eigenvalue (or cluster) is considered to have converged if it lies in an interval of width <i>abstol</i>. If <i>abstol</i> ≤ 0.0, then the tolerance is taken as $\epsilon T _1$, where ϵ is the machine precision.</p>
<i>d, e</i>	<p>REAL for <i>sstebz</i> DOUBLE PRECISION for <i>dstebz</i>.</p> <p>Arrays: <i>d</i>(*) contains the diagonal elements of <i>T</i>. The dimension of <i>d</i> must be at least $\max(1, n)$. <i>e</i>(*) contains the off-diagonal elements of <i>T</i>. The dimension of <i>e</i> must be at least $\max(1, n-1)$.</p>
<i>iwork</i>	<p>INTEGER. Workspace. Array, DIMENSION at least $\max(1, 3n)$.</p>

Output Parameters

<i>m</i>	INTEGER. The actual number of eigenvalues found.
<i>nsplit</i>	INTEGER. The number of diagonal blocks detected in <i>T</i> .
<i>w</i>	<p>REAL for <i>sstebz</i> DOUBLE PRECISION for <i>dstebz</i>.</p> <p>Array, DIMENSION at least $\max(1, n)$. The computed eigenvalues, stored in <i>w</i>(1) to <i>w</i>(<i>m</i>).</p>

iblock, isplit INTEGER.

Arrays, DIMENSION at least $\max(1, n)$.

A positive value *iblock*(*i*) is the block number of the eigenvalue stored in *w*(*i*) (see also *info*).

The leading *nsplit* elements of *isplit* contain points at which *T* splits into blocks T_i as follows: the block T_1 contains rows/columns 1 to *isplit*(1); the block T_2 contains rows/columns *isplit*(1)+1 to *isplit*(2), and so on.

info INTEGER.

If *info* = 0, the execution is successful.

If *info* = 1, for *range* = 'A' or 'V', the algorithm failed to compute some of the required eigenvalues to the desired accuracy; *iblock*(*i*) < 0 indicates that the eigenvalue stored in *w*(*i*) failed to converge.

If *info* = 2, for *range* = 'I', the algorithm failed to compute some of the required eigenvalues. Try calling the routine again with *range* = 'A'.

If *info* = 3:

for *range* = 'A' or 'V', same as *info* = 1;

for *range* = 'I', same as *info* = 2.

If *info* = 4, no eigenvalues have been computed. The floating-point arithmetic on the computer is not behaving as expected.

If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The eigenvalues of *T* are computed to high relative accuracy which means that if they vary widely in magnitude, then any small eigenvalues will be computed more accurately than, for example, with the standard *QR* method. However, the reduction to tridiagonal form (prior to calling the routine) may exclude the possibility of obtaining high relative accuracy in the small eigenvalues of the original matrix if its eigenvalues vary widely in magnitude.

?stein

Computes the eigenvectors corresponding to specified eigenvalues of a real symmetric tridiagonal matrix.

```
call sstein ( n, d, e, m, w, iblock, isplit, z, ldz,  
            work, iwork, ifailv, info )  
call dstein ( n, d, e, m, w, iblock, isplit, z, ldz,  
            work, iwork, ifailv, info )  
call cstein ( n, d, e, m, w, iblock, isplit, z, ldz,  
            work, iwork, ifailv, info )  
call zstein ( n, d, e, m, w, iblock, isplit, z, ldz,  
            work, iwork, ifailv, info )
```

Discussion

This routine computes the eigenvectors of a real symmetric tridiagonal matrix T corresponding to specified eigenvalues, by inverse iteration. It is designed to be used in particular after the specified eigenvalues have been computed by `?stebz` with `order='B'`, but may also be used when the eigenvalues have been computed by other routines. If you use this routine after `?stebz`, it can take advantage of the block structure by performing inverse iteration on each block T_i separately, which is more efficient than using the whole matrix T .

If T has been formed by reduction of a full symmetric or Hermitian matrix A to tridiagonal form, you can transform eigenvectors of T to eigenvectors of A by calling `?ormtr` or `?opmtr` (for real flavors) or by calling `?unmtr` or `?upmtr` (for complex flavors).

Input Parameters

`n` **INTEGER.** The order of the matrix T ($n \geq 0$).

`m` **INTEGER.** The number of eigenvectors to be returned.

d, *e*, *w* **REAL** for single-precision flavors
DOUBLE PRECISION for double-precision flavors.
 Arrays:
d(*) contains the diagonal elements of *T*.
 The dimension of *d* must be at least $\max(1, n)$.
e(*) contains the off-diagonal elements of *T*.
 The dimension of *e* must be at least $\max(1, n-1)$.
w(*) contains the eigenvalues of *T*, stored in *w*(1)
 to *w*(*m*) (as returned by `?stebz`, see [page 5-149](#)).
 Eigenvalues of T_1 must be supplied first, in
 non-decreasing order; then those of T_2 , again in
 non-decreasing order, and so on. Constraint:
 if $iblock(i) = iblock(i+1)$, $w(i) \leq w(i+1)$.
 The dimension of *w* must be at least $\max(1, n)$.

iblock, isplit **INTEGER**.
 Arrays, **DIMENSION** at least $\max(1, n)$.
 The arrays *iblock* and *isplit*, as returned by `?stebz`
 with *order* = 'B'.
 If you did not call `?stebz` with *order* = 'B', set all
 elements of *iblock* to 1, and *isplit*(1) to *n*.

ldz **INTEGER**. The first dimension of the output array *z*;
 $ldz \geq \max(1, n)$.

work **REAL** for single-precision flavors
DOUBLE PRECISION for double-precision flavors.
 Workspace array, **DIMENSION** at least $\max(1, 5n)$.

iwork **INTEGER**.
 Workspace array, **DIMENSION** at least $\max(1, n)$.

Output Parameters

z **REAL** for `sstein`
DOUBLE PRECISION for `dstein`
COMPLEX for `cstein`
DOUBLE COMPLEX for `zstein`.
 Array, **DIMENSION** (*ldz*, *).

If *info* = 0, *z* contains the *m* orthonormal eigenvectors, stored by columns. (The *i*th column corresponds to the *i*th specified eigenvalue.)

ifailv INTEGER. Array, DIMENSION at least max(1, *m*).
If *info* = *i* > 0, the first *i* elements of *ifailv* contain the indices of any eigenvectors that failed to converge.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = *i*, then *i* eigenvectors (as indicated by the parameter *ifailv*) each failed to converge in 5 iterations. The current iterates are stored in the corresponding columns of the array *z*.
If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

Each computed eigenvector z_i is an exact eigenvector of a matrix $T + E_i$, where $\|E_i\|_2 = O(\epsilon) \|T\|_2$. However, a set of eigenvectors computed by this routine may not be orthogonal to so high a degree of accuracy as those computed by `?steqr`.

?disna

Computes the reciprocal condition numbers for the eigenvectors of a symmetric/ Hermitian matrix or for the left or right singular vectors of a general matrix.

```
call sdisna (job, m, n, d, sep, info)
call ddisna (job, m, n, d, sep, info)
```

Discussion

This routine computes the reciprocal condition numbers for the eigenvectors of a real symmetric or complex Hermitian matrix or for the left or right singular vectors of a general *m*-by-*n* matrix.

The reciprocal condition number is the 'gap' between the corresponding eigenvalue or singular value and the nearest other one.

The bound on the error, measured by angle in radians, in the i -th computed vector is given by

$$\text{slamch}('E') * (\text{anorm} / \text{sep}(i))$$

where $\text{anorm} = \|A\|_2 = \max(|d(j)|)$. $\text{sep}(i)$ is not allowed to be smaller than $\text{slamch}('E') * \text{anorm}$ in order to limit the size of the error bound.

`?disna` may also be used to compute error bounds for eigenvectors of the generalized symmetric definite eigenproblem.

Input Parameters

- job** CHARACTER*1. Must be 'E', 'L', or 'R'.
Specifies for which problem the reciprocal condition numbers should be computed:
job = 'E': for the eigenvectors of a symmetric/Hermitian matrix ;
job = 'L': for the left singular vectors of a general matrix;
job = 'R': for the right singular vectors of a general matrix .
- m** INTEGER. The number of rows of the matrix ($m \geq 0$).
- n** INTEGER. If **job = 'L'**, or 'R', the number of columns of the matrix ($n \geq 0$). Ignored if **job = 'E'**.
- d** REAL for `sdisna`
DOUBLE PRECISION for `ddisna`.
Array, dimension at least $\max(1, m)$ if **job = 'E'**, and at least $\max(1, \min(m, n))$ if **job = 'L'** or 'R'.
This array must contain the eigenvalues (if **job = 'E'**) or singular values (if **job = 'L'** or 'R') of the matrix, in either increasing or decreasing order. If singular values, they must be non-negative.

Output Parameters

<i>sep</i>	REAL for <i>sdisna</i> DOUBLE PRECISION for <i>ddisna</i> . Array, dimension at least $\max(1,m)$ if <i>job</i> = 'E', and at least $\max(1, \min(m,n))$ if <i>job</i> = 'L' or 'R'. The reciprocal condition numbers of the vectors.
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value.

Generalized Symmetric-Definite Eigenvalue Problems

Generalized symmetric-definite eigenvalue problems are as follows: find the eigenvalues λ and the corresponding eigenvectors z that satisfy one of these equations:

$$Az = \lambda Bz, \quad ABz = \lambda z, \quad \text{or} \quad BAz = \lambda z$$

where A is an n by n symmetric or Hermitian matrix, and B is an n by n symmetric positive-definite or Hermitian positive-definite matrix.

In these problems, there exist n real eigenvectors corresponding to real eigenvalues (even for complex Hermitian matrices A and B).

Routines described in this section allow you to reduce the above generalized problems to standard symmetric eigenvalue problem $Cy = \lambda y$, which you can solve by calling LAPACK routines described earlier in this chapter (see [page 5-101](#)).

Different routines allow the matrices to be stored either conventionally or in packed storage. Prior to reduction, the positive-definite matrix B must first be factorized using either [?potrf](#) or [?pptrf](#).

The reduction routine for the banded matrices A and B uses a split Cholesky factorization for which a specific routine [?pbstf](#) is provided. This refinement halves the amount of work required to form matrix C .

Table 5-4 Computational Routines for Reducing Generalized Eigenproblems to Standard Problems

Matrix type	Reduce to standard problems (full storage)	Reduce to standard problems (packed storage)	Reduce to standard problems (band matrices)	Factorize band matrix
real symmetric matrices	?sygst	?spgst	?sbgst	?pbstf
complex Hermitian matrices	?hegst /	?hpgst	?hbgst	?pbstf

?sygst

Reduces a real symmetric-definite generalized eigenvalue problem to the standard form.

```
call ssgst ( itype, uplo, n, a, lda, b, ldb, info )
call dsgst ( itype, uplo, n, a, lda, b, ldb, info )
```

Discussion

This routine reduces real symmetric-definite generalized eigenproblems

$$Az = \lambda Bz, \quad ABz = \lambda z, \quad \text{or} \quad BAz = \lambda z$$

to the standard form $Cy = \lambda y$. Here A is a real symmetric matrix, and B is a real symmetric positive-definite matrix. Before calling this routine, call [?potrf](#) to compute the Cholesky factorization: $B = U^T U$ or $B = LL^T$ (see [page 4-14](#)).

Input Parameters

itype **INTEGER**. Must be 1 or 2 or 3.
 If *itype* = 1, the generalized eigenproblem is $Az = \lambda Bz$;
 for *uplo* = 'U': $C = U^{-T}AU^{-1}$, $z = U^{-1}y$;
 for *uplo* = 'L': $C = L^{-1}AL^{-T}$, $z = L^{-T}y$.
 If *itype* = 2, the generalized eigenproblem is $ABz = \lambda z$;
 for *uplo* = 'U': $C = UAU^T$, $z = U^{-1}y$;
 for *uplo* = 'L': $C = L^TAL$, $z = L^{-T}y$.
 If *itype* = 3, the generalized eigenproblem is $BAz = \lambda z$;
 for *uplo* = 'U': $C = UAU^T$, $z = U^T y$;
 for *uplo* = 'L': $C = L^TAL$, $z = Ly$.

uplo **CHARACTER*1**. Must be 'U' or 'L'.
 If *uplo* = 'U', the array *a* stores the upper triangle of A ;
 you must supply B in the factored form $B = U^T U$.
 If *uplo* = 'L', the array *a* stores the lower triangle of A ;
 you must supply B in the factored form $B = LL^T$.

n **INTEGER**. The order of the matrices A and B ($n \geq 0$).

a, *b* REAL for *ssygst*
DOUBLE PRECISION for *dsygst*.
Arrays:
a(*lda*,*) contains the upper or lower triangle of *A*.
The second dimension of *a* must be at least $\max(1, n)$.
b(*ldb*,*) contains the Cholesky-factored matrix *B*:
 $B = U^T U$ or $B = LL^T$ (as returned by *?potrf*).
The second dimension of *b* must be at least $\max(1, n)$.
lda INTEGER. The first dimension of *a*; at least $\max(1, n)$.
ldb INTEGER. The first dimension of *b*; at least $\max(1, n)$.

Output Parameters

a The upper or lower triangle of *A* is overwritten by the upper or lower triangle of *C*, as specified by the arguments *itype* and *uplo*.
info INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

Forming the reduced matrix *C* is a stable procedure. However, it involves implicit multiplication by B^{-1} (if *itype* = 1) or *B* (if *itype* = 2 or 3). When the routine is used as a step in the computation of eigenvalues and eigenvectors of the original problem, there may be a significant loss of accuracy if *B* is ill-conditioned with respect to inversion.

The approximate number of floating-point operations is n^3 .

?hegst

Reduces a complex Hermitian-definite generalized eigenvalue problem to the standard form.

```
call chegst ( itype, uplo, n, a, lda, b, ldb, info )
call zhegst ( itype, uplo, n, a, lda, b, ldb, info )
```

Discussion

This routine reduces complex Hermitian-definite generalized eigenvalue problems

$$Az = \lambda Bz, \quad ABz = \lambda z, \quad \text{or} \quad BAz = \lambda z$$

to the standard form $Cy = \lambda y$. Here the matrix A is complex Hermitian, and B is complex Hermitian positive-definite. Before calling this routine, you must call `?potrf` to compute the Cholesky factorization: $B = U^H U$ or $B = LL^H$ (see [page 4-14](#)).

Input Parameters

itype **INTEGER**. Must be 1 or 2 or 3.
 If **itype** = 1, the generalized eigenproblem is $Az = \lambda Bz$;
 for **uplo** = 'U': $C = U^{-H}AU^{-1}$, $z = U^{-1}y$;
 for **uplo** = 'L': $C = L^{-1}AL^{-H}$, $z = L^{-H}y$.
 If **itype** = 2, the generalized eigenproblem is $ABz = \lambda z$;
 for **uplo** = 'U': $C = UAU^H$, $z = U^{-1}y$;
 for **uplo** = 'L': $C = L^H AL$, $z = L^{-H}y$.
 If **itype** = 3, the generalized eigenproblem is $BAz = \lambda z$;
 for **uplo** = 'U': $C = UAU^H$, $z = U^H y$;
 for **uplo** = 'L': $C = L^H AL$, $z = Ly$.

uplo **CHARACTER*1**. Must be 'U' or 'L'.
 If **uplo** = 'U', the array **a** stores the upper triangle of A ;
 you must supply B in the factored form $B = U^H U$.
 If **uplo** = 'L', the array **a** stores the lower triangle of A ;
 you must supply B in the factored form $B = LL^H$.

n **INTEGER**. The order of the matrices *A* and *B* ($n \geq 0$).

a, *b* **COMPLEX** for **chegst**
DOUBLE COMPLEX for **zhgst**.
Arrays:
a(*lda*,*) contains the upper or lower triangle of *A*.
The second dimension of *a* must be at least $\max(1, n)$.
b(*ldb*,*) contains the Cholesky-factored matrix *B*:
 $B = U^H U$ or $B = LL^H$ (as returned by **?potrf**).
The second dimension of *b* must be at least $\max(1, n)$.

lda **INTEGER**. The first dimension of *a*; at least $\max(1, n)$.

ldb **INTEGER**. The first dimension of *b*; at least $\max(1, n)$.

Output Parameters

a The upper or lower triangle of *A* is overwritten by the upper or lower triangle of *C*, as specified by the arguments *itype* and *uplo*.

info **INTEGER**.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

Forming the reduced matrix *C* is a stable procedure. However, it involves implicit multiplication by B^{-1} (if *itype* = 1) or *B* (if *itype* = 2 or 3). When the routine is used as a step in the computation of eigenvalues and eigenvectors of the original problem, there may be a significant loss of accuracy if *B* is ill-conditioned with respect to inversion.

The approximate number of floating-point operations is n^3 .

?spgst

Reduces a real symmetric-definite generalized eigenvalue problem to the standard form using packed storage.

```
call sspgst ( itype, uplo, n, ap, bp, info )
call dspgst ( itype, uplo, n, ap, bp, info )
```

Discussion

This routine reduces real symmetric-definite generalized eigenproblems

$$Az = \lambda Bz, \quad ABz = \lambda z, \quad \text{or} \quad BAz = \lambda z$$

to the standard form $Cy = \lambda y$, using packed matrix storage. Here A is a real symmetric matrix, and B is a real symmetric positive-definite matrix.

Before calling this routine, call [?pptrf](#) to compute the Cholesky factorization: $B = U^T U$ or $B = LL^T$ (see [page 4-16](#)).

Input Parameters

itype **INTEGER**. Must be 1 or 2 or 3.
 If *itype* = 1, the generalized eigenproblem is $Az = \lambda Bz$;
 for *uplo* = 'U': $C = U^{-T}AU^{-1}$, $z = U^{-1}y$;
 for *uplo* = 'L': $C = L^{-1}AL^{-T}$, $z = L^{-T}y$.
 If *itype* = 2, the generalized eigenproblem is $ABz = \lambda z$;
 for *uplo* = 'U': $C = UAU^T$, $z = U^{-1}y$;
 for *uplo* = 'L': $C = L^TAL$, $z = L^{-T}y$.
 If *itype* = 3, the generalized eigenproblem is $BAz = \lambda z$;
 for *uplo* = 'U': $C = UAU^T$, $z = U^T y$;
 for *uplo* = 'L': $C = L^TAL$, $z = Ly$.

uplo **CHARACTER*1**. Must be 'U' or 'L'.
 If *uplo* = 'U', *ap* stores the packed upper triangle of A ;
 you must supply B in the factored form $B = U^T U$.
 If *uplo* = 'L', *ap* stores the packed lower triangle of A ;
 you must supply B in the factored form $B = LL^T$.

n **INTEGER**. The order of the matrices A and B ($n \geq 0$).

ap, *bp* REAL for *sspgst*
 DOUBLE PRECISION for *dspgst*.
Arrays:
ap(*) contains the packed upper or lower triangle of *A*.
The dimension of *ap* must be at least $\max(1, n*(n+1)/2)$.
bp(*) contains the packed Cholesky factor of *B*
(as returned by *?pptrf* with the same *uplo* value).
The dimension of *bp* must be at least $\max(1, n*(n+1)/2)$.

Output Parameters

ap The upper or lower triangle of *A* is overwritten by the
 upper or lower triangle of *C*, as specified by the
 arguments *itype* and *uplo*.
info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

Forming the reduced matrix *C* is a stable procedure. However, it involves implicit multiplication by B^{-1} (if *itype* = 1) or *B* (if *itype* = 2 or 3). When the routine is used as a step in the computation of eigenvalues and eigenvectors of the original problem, there may be a significant loss of accuracy if *B* is ill-conditioned with respect to inversion.
The approximate number of floating-point operations is n^3 .

?hpgst

Reduces a complex Hermitian-definite generalized eigenvalue problem to the standard form using packed storage.

```
call chpgst ( itype, uplo, n, ap, bp, info )
call zhpgst ( itype, uplo, n, ap, bp, info )
```

Discussion

This routine reduces real symmetric-definite generalized eigenproblems

$$Az = \lambda Bz, \quad ABz = \lambda z, \quad \text{or} \quad BAz = \lambda z$$

to the standard form $Cy = \lambda y$, using packed matrix storage. Here A is a real symmetric matrix, and B is a real symmetric positive-definite matrix.

Before calling this routine, you must call [?pptrf](#) to compute the Cholesky factorization: $B = U^H U$ or $B = LL^H$ (see [page 4-16](#)).

Input Parameters

itype **INTEGER**. Must be 1 or 2 or 3.
 If **itype** = 1, the generalized eigenproblem is $Az = \lambda Bz$;
 for **uplo** = 'U': $C = U^{-H}AU^{-1}$, $z = U^{-1}y$;
 for **uplo** = 'L': $C = L^{-1}AL^{-H}$, $z = L^{-H}y$.
 If **itype** = 2, the generalized eigenproblem is $ABz = \lambda z$;
 for **uplo** = 'U': $C = UAU^H$, $z = U^{-1}y$;
 for **uplo** = 'L': $C = L^H AL$, $z = L^{-H}y$.
 If **itype** = 3, the generalized eigenproblem is $BAz = \lambda z$;
 for **uplo** = 'U': $C = UAU^H$, $z = U^H y$;
 for **uplo** = 'L': $C = L^H AL$, $z = Ly$.

uplo **CHARACTER*1**. Must be 'U' or 'L'.
 If **uplo** = 'U', **ap** stores the packed upper triangle of A ;
 you must supply B in the factored form $B = U^H U$.
 If **uplo** = 'L', **ap** stores the packed lower triangle of A ;
 you must supply B in the factored form $B = LL^H$.

n **INTEGER**. The order of the matrices A and B ($n \geq 0$).

ap, *bp* COMPLEX for *chpgst*
 DOUBLE COMPLEX for *zhpgst*.
Arrays:
ap(*) contains the packed upper or lower triangle of *A*.
The dimension of *a* must be at least $\max(1, n*(n+1)/2)$.
bp(*) contains the packed Cholesky factor of *B*
(as returned by *?pptrf* with the same *uplo* value).
The dimension of *b* must be at least $\max(1, n*(n+1)/2)$.

Output Parameters

ap The upper or lower triangle of *A* is overwritten by the
 upper or lower triangle of *C*, as specified by the
 arguments *itype* and *uplo*.
info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

Forming the reduced matrix *C* is a stable procedure. However, it involves implicit multiplication by B^{-1} (if *itype* = 1) or *B* (if *itype* = 2 or 3). When the routine is used as a step in the computation of eigenvalues and eigenvectors of the original problem, there may be a significant loss of accuracy if *B* is ill-conditioned with respect to inversion.

The approximate number of floating-point operations is n^3 .

?sbgst

Reduces a real symmetric-definite generalized eigenproblem for banded matrices to the standard form using the factorization performed by ?pbstf.

```
call ssgst ( vect, uplo, n, ka, kb, ab, ldab, bb, ldbb, x, ldx,
            work, info )
call dsbgst ( vect, uplo, n, ka, kb, ab, ldab, bb, ldbb, x, ldx,
            work, info )
```

Discussion

To reduce the real symmetric-definite generalized eigenproblem $Az = \lambda Bz$ to the standard form $Cy = \lambda y$, where A , B and C are banded, this routine must be preceded by a call to [spbstf/dpbstf](#), which computes the split Cholesky factorization of the positive-definite matrix B : $B = S^T S$. The split Cholesky factorization, compared with the ordinary Cholesky factorization, allows the work to be approximately halved.

This routine overwrites A with $C = X^T A X$, where $X = S^{-1} Q$ and Q is an orthogonal matrix chosen (implicitly) to preserve the bandwidth of A . The routine also has an option to allow the accumulation of X , and then, if z is an eigenvector of C , Xz is an eigenvector of the original system.

Input Parameters

vect CHARACTER*1. Must be 'N' or 'V'.
If **vect** = 'N', then matrix X is not returned;
If **vect** = 'V', then matrix X is returned.

uplo CHARACTER*1. Must be 'U' or 'L'.
If **uplo** = 'U', **ab** stores the upper triangular part of A .
If **uplo** = 'L', **ab** stores the lower triangular part of A .

n INTEGER. The order of the matrices A and B ($n \geq 0$).

ka INTEGER. The number of super- or sub-diagonals in A ($ka \geq 0$).

<i>kb</i>	INTEGER. The number of super- or sub-diagonals in <i>B</i> ($ka \geq kb \geq 0$).
<i>ab,bb,work</i>	REAL for <i>ssbgst</i> DOUBLE PRECISION for <i>dsbgst</i> <i>ab</i> (<i>ldab</i> ,*) is an array containing either upper or lower triangular part of the symmetric matrix <i>A</i> (as specified by <i>uplo</i>) in band storage format. The second dimension of the array <i>ab</i> must be at least $\max(1, n)$. <i>bb</i> (<i>lbbb</i> ,*) is an array containing the banded split Cholesky factor of <i>B</i> as specified by <i>uplo</i> , <i>n</i> and <i>kb</i> and returned by <i>spbstf</i> / <i>dpbstf</i> . The second dimension of the array <i>bb</i> must be at least $\max(1, n)$. <i>work</i> (*) is a workspace array, DIMENSION at least $\max(1, 2*n)$
<i>ldab</i>	INTEGER. The first dimension of the array <i>ab</i> ; must be at least $ka+1$.
<i>lbbb</i>	INTEGER. The first dimension of the array <i>bb</i> ; must be at least $kb+1$.
<i>ldx</i>	The first dimension of the output array <i>x</i> . Constraints: if <i>vect</i> = 'N' , then $ldx \geq 1$; if <i>vect</i> = 'V' , then $ldx \geq \max(1, n)$.

Output Parameters

<i>ab</i>	On exit, this array is overwritten by the upper or lower triangle of <i>C</i> as specified by <i>uplo</i> .
<i>x</i>	REAL for <i>ssbgst</i> DOUBLE PRECISION for <i>dsbgst</i> Array. If <i>vect</i> = 'V' , then <i>x</i> (<i>ldx</i> ,*) contains the <i>n</i> by <i>n</i> matrix $X = S^{-1}Q$. If <i>vect</i> = 'N' , then <i>x</i> is not referenced. The second dimension of <i>x</i> must be: at least $\max(1, n)$, if <i>vect</i> = 'V' ; at least 1, if <i>vect</i> = 'N' .

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

Forming the reduced matrix *C* involves implicit multiplication by B^{-1} . When the routine is used as a step in the computation of eigenvalues and eigenvectors of the original problem, there may be a significant loss of accuracy if *B* is ill-conditioned with respect to inversion.

The total number of floating-point operations is approximately $6n^2 * kb$, when *vect* = 'N'. Additional $(3/2)n^3 * (kb/ka)$ operations are required when *vect* = 'V'. All these estimates assume that both *ka* and *kb* are much less than *n*.

?hbgst

Reduces a complex Hermitian-definite generalized eigenproblem for banded matrices to the standard form using the factorization performed by ?pbstf.

```
call chbgst ( vect, uplo, n, ka, kb, ab, ldab, bb, ldbb, x, ldx,
             work, rwork, info )
call zhbgst ( vect, uplo, n, ka, kb, ab, ldab, bb, ldbb, x, ldx,
             work, rwork, info )
```

Discussion

To reduce the complex Hermitian-definite generalized eigenproblem $Az = \lambda Bz$ to the standard form $Cy = \lambda y$, where A , B and C are banded, this routine must be preceded by a call to [cpbstf/zpbstf](#), which computes the split Cholesky factorization of the positive-definite matrix B : $B = S^H S$. The split Cholesky factorization, compared with the ordinary Cholesky factorization, allows the work to be approximately halved.

This routine overwrites A with $C = X^H A X$, where $X = S^{-1} Q$ and Q is a unitary matrix chosen (implicitly) to preserve the bandwidth of A . The routine also has an option to allow the accumulation of X , and then, if z is an eigenvector of C , Xz is an eigenvector of the original system.

Input Parameters

vect CHARACTER*1. Must be 'N' or 'V'.
 If **vect** = 'N', then matrix X is not returned;
 If **vect** = 'V', then matrix X is returned.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If **uplo** = 'U', **ab** stores the upper triangular part of A .
 If **uplo** = 'L', **ab** stores the lower triangular part of A .

n INTEGER. The order of the matrices A and B ($n \geq 0$).

ka INTEGER. The number of super- or sub-diagonals in A ($ka \geq 0$).

<i>kb</i>	INTEGER. The number of super- or sub-diagonals in <i>B</i> ($ka \geq kb \geq 0$).
<i>ab,bb,work</i>	COMPLEX for <i>chbgst</i> DOUBLE COMPLEX for <i>zhbgst</i> <i>ab (ldab,*)</i> is an array containing either upper or lower triangular part of the Hermitian matrix <i>A</i> (as specified by <i>uplo</i>) in band storage format. The second dimension of the array <i>ab</i> must be at least $\max(1, n)$. <i>bb (ldb,*)</i> is an array containing the banded split Cholesky factor of <i>B</i> as specified by <i>uplo, n</i> and <i>kb</i> and returned by <i>cpbstf/zpbstf</i> . The second dimension of the array <i>bb</i> must be at least $\max(1, n)$. <i>work(*)</i> is a workspace array, DIMENSION at least $\max(1, n)$
<i>ldab</i>	INTEGER. The first dimension of the array <i>ab</i> ; must be at least $ka+1$.
<i>ldb</i>	INTEGER. The first dimension of the array <i>bb</i> ; must be at least $kb+1$.
<i>ldx</i>	The first dimension of the output array <i>x</i> . Constraints: if <i>vect</i> = 'N' , then $ldx \geq 1$; if <i>vect</i> = 'V' , then $ldx \geq \max(1, n)$.
<i>rwork</i>	REAL for <i>chbgst</i> DOUBLE PRECISION for <i>zhbgst</i> Workspace array, DIMENSION at least $\max(1, n)$

Output Parameters

<i>ab</i>	On exit, this array is overwritten by the upper or lower triangle of <i>C</i> as specified by <i>uplo</i> .
<i>x</i>	COMPLEX for <i>chbgst</i> DOUBLE COMPLEX for <i>zhbgst</i> Array. If <i>vect</i> = 'V', then <i>x (ldx,*)</i> contains the <i>n</i> by <i>n</i> matrix $X = S^{-1}Q$. If <i>vect</i> = 'N', then <i>x</i> is not referenced.

The second dimension of x must be:
at least $\max(1, n)$, if $vect = 'V'$;
at least 1, if $vect = 'N'$.

info

INTEGER.

If $info = 0$, the execution is successful.

If $info = -i$, the i th parameter had an illegal value.

Application Notes

Forming the reduced matrix C involves implicit multiplication by B^{-1} . When the routine is used as a step in the computation of eigenvalues and eigenvectors of the original problem, there may be a significant loss of accuracy if B is ill-conditioned with respect to inversion.

The total number of floating-point operations is approximately $20n^2 * kb$, when $vect = 'N'$. Additional $5n^3 * (kb/ka)$ operations are required when $vect = 'V'$. All these estimates assume that both ka and kb are much less than n .

?pbstf

Computes a split Cholesky factorization
of a real symmetric or complex
Hermitian positive-definite banded
matrix used in ?sbgst/?hbgst .

```
call spbstf ( uplo, n, kb, bb, ldbb, info )
call dpbstf ( uplo, n, kb, bb, ldbb, info )
call cpbstf ( uplo, n, kb, bb, ldbb, info )
call zpbstf ( uplo, n, kb, bb, ldbb, info )
```

Discussion

This routine computes a split Cholesky factorization of a real symmetric or complex Hermitian positive-definite band matrix B . It is to be used in conjunction with ?sbgst/?hbgst.

The factorization has the form $B = S^T S$ (or $B = S^H S$ for complex flavors), where S is a band matrix of the same bandwidth as B and the following structure: S is upper triangular in the first $(n+kb)/2$ rows and lower triangular in the remaining rows.

Input Parameters

uplo CHARACTER*1. Must be 'U' or 'L'.
If **uplo** = 'U', **bb** stores the upper triangular part of B .
If **uplo** = 'L', **bb** stores the lower triangular part of B .

n INTEGER. The order of the matrix B ($n \geq 0$).

kb INTEGER. The number of super- or sub-diagonals in B ($kb \geq 0$).

bb REAL for spbstf
DOUBLE PRECISION for dpbstf
COMPLEX for cpbstf
DOUBLE COMPLEX for zpbstf.
bb (**ldbb**,*) is an array containing either upper or lower triangular part of the matrix B (as specified by

uplo) in band storage format.

The second dimension of the array *bb* must be at least $\max(1, n)$.

ldbb **INTEGER.** The first dimension of *bb*; must be at least $kb+1$.

Output Parameters

bb On exit, this array is overwritten by the elements of the split Cholesky factor *S*.

info **INTEGER.**
 If *info* = 0, the execution is successful.
 If *info* = *i*, then the factorization could not be completed, because the updated element b_{ii} would be the square root of a negative number; hence the matrix *B* is not positive-definite.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The computed factor *S* is the exact factor of a perturbed matrix $B + E$, where

$$|E| \leq c(kb + 1)\epsilon |S^H| |S|, \quad |e_{ij}| \leq c(kb + 1)\epsilon \sqrt{b_{ii}b_{jj}}$$

$c(n)$ is a modest linear function of *n*, and ϵ is the machine precision.

The total number of floating-point operations for real flavors is approximately $n(kb+1)^2$. The number of operations for complex flavors is 4 times greater. All these estimates assume that *kb* is much less than *n*.

After calling this routine, you can call [?sbgst/?hbgst](#) to solve the generalized eigenproblem $Az = \lambda Bz$, where *A* and *B* are banded and *B* is positive-definite.

Nonsymmetric Eigenvalue Problems

This section describes LAPACK routines for solving nonsymmetric eigenvalue problems, computing the Schur factorization of general matrices, as well as performing a number of related computational tasks.

A *nonsymmetric eigenvalue problem* is as follows: given a nonsymmetric (or non-Hermitian) matrix A , find the *eigenvalues* λ and the corresponding *eigenvectors* z that satisfy the equation

$$Az = \lambda z \quad (\text{right eigenvectors } z)$$

or the equation

$$z^H A = \lambda z^H \quad (\text{left eigenvectors } z).$$

Nonsymmetric eigenvalue problems have the following properties:

- The number of eigenvectors may be less than the matrix order (but is not less than the number of *distinct eigenvalues* of A).
- Eigenvalues may be complex even for a real matrix A .
- If a real nonsymmetric matrix has a complex eigenvalue $a+bi$ corresponding to an eigenvector z , then $a-bi$ is also an eigenvalue. The eigenvalue $a-bi$ corresponds to the eigenvector whose elements are complex conjugate to the elements of z .

To solve a nonsymmetric eigenvalue problem with LAPACK, you usually need to reduce the matrix to the upper Hessenberg form and then solve the eigenvalue problem with the Hessenberg matrix obtained. [Table 5-5](#) lists LAPACK routines for reducing the matrix to the upper Hessenberg form by an orthogonal (or unitary) similarity transformation $A = QHQ^H$ as well as routines for solving eigenvalue problems with Hessenberg matrices, forming the Schur factorization of such matrices and computing the corresponding condition numbers.

Decision tree in [Figure 5-4](#) helps you choose the right routine or sequence of routines for an eigenvalue problem with a real nonsymmetric matrix. If you need to solve an eigenvalue problem with a complex non-Hermitian matrix, use the decision tree shown in [Figure 5-5](#).

Table 5-5 Computational Routines for Solving Nonsymmetric Eigenvalue Problems

Operation performed	Routines for real matrices	Routines for complex matrices
Reduce to Hessenberg form $A = QHQ^H$?gehrd ,	?gehrd
Generate the matrix Q	?orghr	?unghr
Apply the matrix Q	?ormhr	?unmhr
Balance matrix	?gebal	?gebal
Transform eigenvectors of balanced matrix to those of the original matrix	?gebak	?gebak
Find eigenvalues and Schur factorization (QR algorithm)	?hseqr	?hseqr
Find eigenvectors from Hessenberg form (inverse iteration)	?hsein	?hsein
Find eigenvectors from Schur factorization	?trevc	?trevc
Estimate sensitivities of eigenvalues and eigenvectors	?trsna	?trsna
Reorder Schur factorization	?trexc	?trexc
Reorder Schur factorization, find the invariant subspace and estimate sensitivities	?trsen	?trsen
Solves Sylvester's equation.	?trsyl	?trsyl

Figure 5-4 Decision Tree: Real Nonsymmetric Eigenvalue Problems

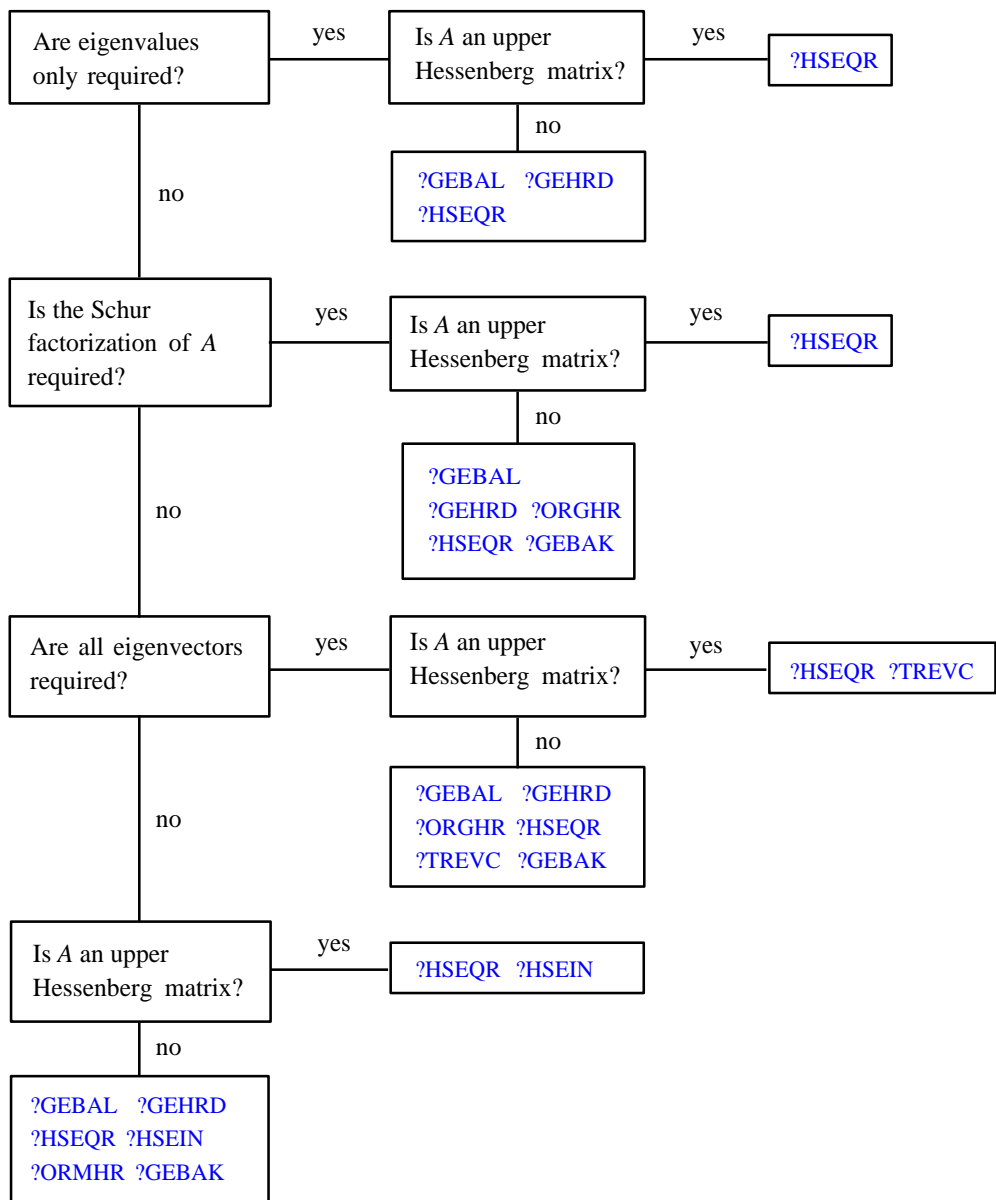
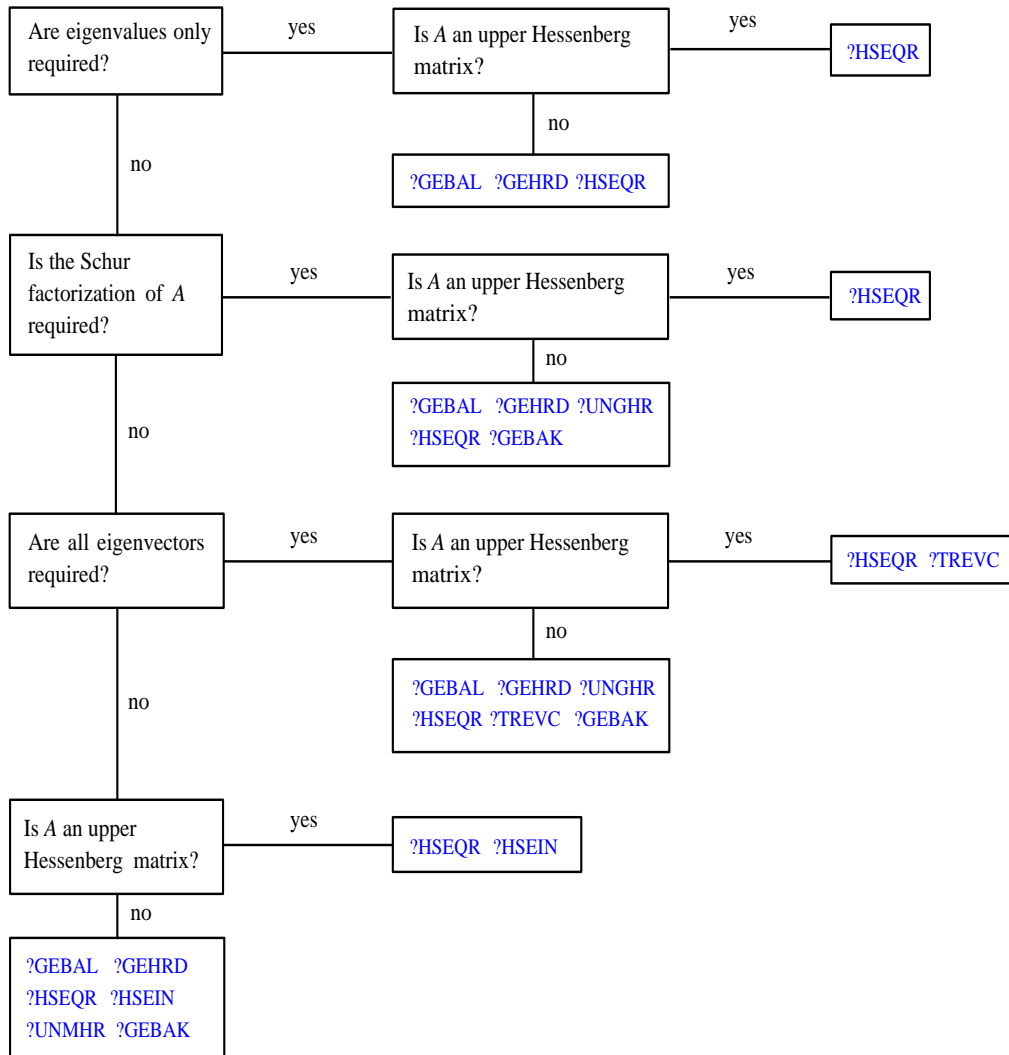


Figure 5-5 Decision Tree: Complex Non-Hermitian Eigenvalue Problems



?gehrd

Reduces a general matrix to upper Hessenberg form.

```
call sgehrd ( n, ilo, ihi, a, lda, tau, work, lwork, info )
call dgehrd ( n, ilo, ihi, a, lda, tau, work, lwork, info )
call cgehrd ( n, ilo, ihi, a, lda, tau, work, lwork, info )
call zgehrd ( n, ilo, ihi, a, lda, tau, work, lwork, info )
```

Discussion

The routine reduces a general matrix A to upper Hessenberg form H by an orthogonal or unitary similarity transformation $A = QHQ^H$. Here H has real subdiagonal elements.

The routine does not form the matrix Q explicitly. Instead, Q is represented as a product of *elementary reflectors*. Routines are provided to work with Q in this representation.

Input Parameters

n **INTEGER.** The order of the matrix A ($n \geq 0$).

ilo, ihi **INTEGER.** If A has been output by ?gebal, then *ilo* and *ihi* must contain the values returned by that routine. Otherwise *ilo* = 1 and *ihi* = n . (If $n > 0$, then $1 \leq ilo \leq ihi \leq n$; if $n = 0$, *ilo* = 1 and *ihi* = 0.)

a, work **REAL** for sgehrd
DOUBLE PRECISION for dgehrd
COMPLEX for cgehrd
DOUBLE COMPLEX for zgehrd.
Arrays:
a (*lda*,*) contains the matrix A .
The second dimension of *a* must be at least $\max(1, n)$.
work (*lwork*) is a workspace array.

lda **INTEGER.** The first dimension of *a*; at least $\max(1, n)$.

lwork INTEGER. The size of the *work* array; at least $\max(1, n)$. See *Application notes* for the suggested value of *lwork*.

Output Parameters

a Overwritten by the upper Hessenberg matrix H and details of the matrix Q . The subdiagonal elements of H are real.

tau REAL for *sgehrd*
DOUBLE PRECISION for *dgehrd*
COMPLEX for *cgehrd*
DOUBLE COMPLEX for *zgehrd*.
Array, DIMENSION at least $\max(1, n-1)$.
Contains additional information on the matrix Q .

work(1) If *info* = 0, on exit *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = $-i$, the i th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = n * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

The computed Hessenberg matrix H is exactly similar to a nearby matrix $A + E$, where $\|E\|_2 < c(n)\epsilon \|A\|_2$, $c(n)$ is a modestly increasing function of n , and ϵ is the machine precision.

The approximate number of floating-point operations for real flavors is $(2/3)(ihi - ilo)^2(2ihi + 2ilo + 3n)$; for complex flavors it is 4 times greater.

?orghr

Generates the real orthogonal matrix Q determined by ?gehrd.

```
call sorghr ( n, ilo, ihi, a, lda, tau, work, lwork, info )
call dorghr ( n, ilo, ihi, a, lda, tau, work, lwork, info )
```

Discussion

This routine explicitly generates the orthogonal matrix Q that has been determined by a preceding call to `sgehrd/dgehrd`. (The routine `?gehrd` reduces a real general matrix A to upper Hessenberg form H by an orthogonal similarity transformation, $A = QHQ^T$, and represents the matrix Q as a product of $ihi - ilo$ elementary reflectors. Here ilo and ihi are values determined by `sgebal/dgebal` when balancing the matrix; if the matrix has not been balanced, $ilo = 1$ and $ihi = n$.)

The matrix Q generated by `?orghr` has the structure:

$$Q = \begin{bmatrix} I & 0 & 0 \\ 0 & Q_{22} & 0 \\ 0 & 0 & I \end{bmatrix}$$

where Q_{22} occupies rows and columns ilo to ihi .

Input Parameters

n **INTEGER.** The order of the matrix Q ($n \geq 0$).

ilo, ihi **INTEGER.** These must be the same parameters ilo and ihi , respectively, as supplied to `?gehrd`. (If $n > 0$, then $1 \leq ilo \leq ihi \leq n$; if $n = 0$, $ilo = 1$ and $ihi = 0$.)

$a, tau, work$ **REAL** for `sorghr`
DOUBLE PRECISION for `dorghr`
Arrays:

$a(lda, *)$ contains details of the vectors which define the elementary reflectors, as returned by `?gehrd`.

The second dimension of a must be at least $\max(1, n)$.

$tau(*)$ contains further details of the elementary reflectors, as returned by `?gehrd`.

The dimension of tau must be at least $\max(1, n-1)$.

$work(lwork)$ is a workspace array.

lda INTEGER. The first dimension of a ; at least $\max(1, n)$.

$lwork$ INTEGER. The size of the $work$ array;
 $lwork \geq \max(1, ihi-ilo)$.
 See *Application notes* for the suggested value of $lwork$.

Output Parameters

a Overwritten by the n by n orthogonal matrix Q .

$work(1)$ If $info = 0$, on exit $work(1)$ contains the minimum value of $lwork$ required for optimum performance. Use this $lwork$ for subsequent runs.

$info$ INTEGER.
 If $info = 0$, the execution is successful.
 If $info = -i$, the i th parameter had an illegal value.

Application Notes

For better performance, try using $lwork = (ihi-ilo) * blocksize$, where $blocksize$ is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of $lwork$ for the first run. On exit, examine $work(1)$ and use this value for subsequent runs.

The computed matrix Q differs from the exact result by a matrix E such that $\|E\|_2 = O(\epsilon)$, where ϵ is the machine precision.

The approximate number of floating-point operations is $(4/3)(ihi-ilo)^3$.

The complex counterpart of this routine is [?unghr](#).

?ormhr

Multiplies an arbitrary real matrix C by the real orthogonal matrix Q determined by ?gehrd.

```
call sormhr ( side, trans, m, n, ilo, ihi, a, lda, tau, c, ldc,
             work, lwork, info )
call dormhr ( side, trans, m, n, ilo, ihi, a, lda, tau, c, ldc,
             work, lwork, info )
```

Discussion

This routine multiplies a matrix C by the orthogonal matrix Q that has been determined by a preceding call to sgehrd/dgehrd. (The routine ?gehrd reduces a real general matrix A to upper Hessenberg form H by an orthogonal similarity transformation, $A = QHQ^T$, and represents the matrix Q as a product of $ihi-iilo$ elementary reflectors. Here ilo and ihi are values determined by sgebal/dgebal when balancing the matrix; if the matrix has not been balanced, $ilo = 1$ and $ihi = n$.)

With ?ormhr, you can form one of the matrix products QC , Q^TC , CQ , or CQ^T , overwriting the result on C (which may be any real rectangular matrix).

A common application of ?ormhr is to transform a matrix V of eigenvectors of H to the matrix QV of eigenvectors of A .

Input Parameters

<i>side</i>	CHARACTER*1. Must be 'L' or 'R'. If <i>side</i> = 'L', then the routine forms QC or Q^TC . If <i>side</i> = 'R', then the routine forms CQ or CQ^T .
<i>trans</i>	CHARACTER*1. Must be 'N' or 'T'. If <i>trans</i> = 'N', then Q is applied to C . If <i>trans</i> = 'T', then Q^T is applied to C .
<i>m</i>	INTEGER. The number of rows in C ($m \geq 0$).
<i>n</i>	INTEGER. The number of columns in C ($n \geq 0$).

ilo, ihi **INTEGER.** These must be the same parameters *ilo* and *ihi*, respectively, as supplied to `?gehrd`.
 If $m > 0$ and *side* = 'L', then $1 \leq ilo \leq ihi \leq m$.
 If $m = 0$ and *side* = 'L', then *ilo* = 1 and *ihi* = 0.
 If $n > 0$ and *side* = 'R', then $1 \leq ilo \leq ihi \leq n$.
 If $n = 0$ and *side* = 'R', then *ilo* = 1 and *ihi* = 0.

a, tau, c, work **REAL** for `sormhr`
 DOUBLE PRECISION for `dormhr`
 Arrays:
a(lda,)* contains details of the vectors which define the *elementary reflectors*, as returned by `?gehrd`.
 The second dimension of *a* must be at least $\max(1, m)$ if *side* = 'L' and at least $\max(1, n)$ if *side* = 'R'.
tau()* contains further details of the *elementary reflectors*, as returned by `?gehrd`.
 The dimension of *tau* must be at least $\max(1, m-1)$ if *side* = 'L' and at least $\max(1, n-1)$ if *side* = 'R'.
c(ldc,)* contains the m by n matrix *C*.
 The second dimension of *c* must be at least $\max(1, n)$.
work(lwork) is a workspace array.

lda **INTEGER.** The first dimension of *a*; at least $\max(1, m)$ if *side* = 'L' and at least $\max(1, n)$ if *side* = 'R'.

ldc **INTEGER.** The first dimension of *c*; at least $\max(1, m)$.

lwork **INTEGER.** The size of the *work* array.
 If *side* = 'L', *lwork* $\geq \max(1, n)$.
 If *side* = 'R', *lwork* $\geq \max(1, m)$.
 See *Application notes* for the suggested value of *lwork*.

Output Parameters

c *C* is overwritten by QC or $Q^T C$ or CQ^T or CQ as specified by *side* and *trans*.

work(1) If *info* = 0, on exit *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, *lwork* should be at least $n \cdot \text{blocksize}$ if *side* = 'L' and at least $m \cdot \text{blocksize}$ if *side* = 'R', where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

The computed matrix *Q* differs from the exact result by a matrix *E* such that $\|E\|_2 = O(\epsilon)\|C\|_2$, where ϵ is the machine precision.

The approximate number of floating-point operations is

$2n(ihi-ilo)^2$ if *side* = 'L';

$2m(ihi-ilo)^2$ if *side* = 'R'.

The complex counterpart of this routine is [?unmhr](#).

?unghr

Generates the complex unitary matrix Q determined by ?gehrd.

```
call cunghr ( n, ilo, ihi, a, lda, tau, work, lwork, info )
call zunghr ( n, ilo, ihi, a, lda, tau, work, lwork, info )
```

Discussion

This routine is intended to be used following a call to cgehrd/zgehrd, which reduces a complex matrix A to upper Hessenberg form H by a unitary similarity transformation: $A = QHQ^H$. ?gehrd represents the matrix Q as a product of $ihi-ilo$ elementary reflectors. Here ilo and ihi are values determined by cgebal/zgebal when balancing the matrix; if the matrix has not been balanced, $ilo = 1$ and $ihi = n$.

Use the routine ?unghr to generate Q explicitly as a square matrix. The matrix Q has the structure:

$$Q = \begin{bmatrix} I & 0 & 0 \\ 0 & Q_{22} & 0 \\ 0 & 0 & I \end{bmatrix}$$

where Q_{22} occupies rows and columns ilo to ihi .

Input Parameters

n **INTEGER**. The order of the matrix Q ($n \geq 0$).

ilo, ihi **INTEGER**. These must be the same parameters ilo and ihi , respectively, as supplied to ?gehrd. (If $n > 0$, then $1 \leq ilo \leq ihi \leq n$. If $n = 0$, then $ilo = 1$ and $ihi = 0$.)

$a, tau, work$ **COMPLEX** for cunghr
 DOUBLE COMPLEX for zunghr.

Arrays:

`a(lda,*)` contains details of the vectors which define the *elementary reflectors*, as returned by `?gehrd`.

The second dimension of `a` must be at least $\max(1, n)$.

`tau(*)` contains further details of the *elementary reflectors*, as returned by `?gehrd`.

The dimension of `tau` must be at least $\max(1, n-1)$.

`work(lwork)` is a workspace array.

`lda` INTEGER. The first dimension of `a`; at least $\max(1, n)$.

`lwork` INTEGER. The size of the `work` array;
`lwork` $\geq \max(1, ihi-ilo)$.

See *Application notes* for the suggested value of `lwork`.

Output Parameters

`a` Overwritten by the n by n unitary matrix Q .

`work(1)` If `info` = 0, on exit `work(1)` contains the minimum value of `lwork` required for optimum performance. Use this `lwork` for subsequent runs.

`info` INTEGER.
 If `info` = 0, the execution is successful.
 If `info` = $-i$, the i th parameter had an illegal value.

Application Notes

For better performance, try using `lwork = (ihi-ilo)*blocksize`, where `blocksize` is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of `lwork` for the first run. On exit, examine `work(1)` and use this value for subsequent runs.

The computed matrix Q differs from the exact result by a matrix E such that $\|E\|_2 = O(\epsilon)$, where ϵ is the machine precision.

The approximate number of real floating-point operations is $(16/3)(ihi-ilo)^3$.

The real counterpart of this routine is [?orghr](#).

?unmhr

Multiplies an arbitrary complex matrix C by the complex unitary matrix Q determined by ?gehrd.

```
call cummhr ( side, trans, m, n, ilo, ihi, a, lda, tau, c, ldc,
              work, lwork, info )
call zummhr ( side, trans, m, n, ilo, ihi, a, lda, tau, c, ldc,
              work, lwork, info )
```

Discussion

This routine multiplies a matrix C by the unitary matrix Q that has been determined by a preceding call to cgehrd/zgehrd. (The routine ?gehrd reduces a real general matrix A to upper Hessenberg form H by an orthogonal similarity transformation, $A = QHQ^H$, and represents the matrix Q as a product of $ihi-iilo$ elementary reflectors. Here ilo and ihi are values determined by cgebal/zgebal when balancing the matrix; if the matrix has not been balanced, $ilo = 1$ and $ihi = n$.)

With ?unmhr, you can form one of the matrix products QC , Q^HC , CQ , or CQ^H , overwriting the result on C (which may be any complex rectangular matrix). A common application of this routine is to transform a matrix V of eigenvectors of H to the matrix QV of eigenvectors of A .

Input Parameters

side CHARACTER*1. Must be 'L' or 'R'.
 If **side** = 'L', then the routine forms QC or Q^HC .
 If **side** = 'R', then the routine forms CQ or CQ^H .

trans CHARACTER*1. Must be 'N' or 'C'.
 If **trans** = 'N', then Q is applied to C .
 If **trans** = 'T', then Q^H is applied to C .

m INTEGER. The number of rows in C ($m \geq 0$).

n INTEGER. The number of columns in C ($n \geq 0$).

ilo, ihi **INTEGER**. These must be the same parameters *ilo* and *ihi*, respectively, as supplied to `?gehrd`.
 If $m > 0$ and *side* = 'L', then $1 \leq ilo \leq ihi \leq m$.
 If $m = 0$ and *side* = 'L', then *ilo* = 1 and *ihi* = 0.
 If $n > 0$ and *side* = 'R', then $1 \leq ilo \leq ihi \leq n$.
 If $n = 0$ and *side* = 'R', then *ilo* = 1 and *ihi* = 0.

a, tau, c, work **COMPLEX** for `cunmhr`
DOUBLE COMPLEX for `zunmhr`.
 Arrays:
a (*lda*, *) contains details of the vectors which define the elementary reflectors, as returned by `?gehrd`.
 The second dimension of *a* must be at least $\max(1, m)$ if *side* = 'L' and at least $\max(1, n)$ if *side* = 'R'.
tau(*) contains further details of the elementary reflectors, as returned by `?gehrd`.
 The dimension of *tau* must be at least $\max(1, m-1)$ if *side* = 'L' and at least $\max(1, n-1)$ if *side* = 'R'.
c (*ldc*, *) contains the m by n matrix *C*.
 The second dimension of *c* must be at least $\max(1, n)$.
work (*lwork*) is a workspace array.

lda **INTEGER**. The first dimension of *a*; at least $\max(1, m)$ if *side* = 'L' and at least $\max(1, n)$ if *side* = 'R'.

ldc **INTEGER**. The first dimension of *c*; at least $\max(1, m)$.

lwork **INTEGER**. The size of the *work* array.
 If *side* = 'L', $lwork \geq \max(1, n)$.
 If *side* = 'R', $lwork \geq \max(1, m)$.
 See *Application notes* for the suggested value of *lwork*.

Output Parameters

c *C* is overwritten by QC or $Q^H C$ or CQ^H or CQ as specified by *side* and *trans*.

work(1) If *info* = 0, on exit *work*(1) contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, *lwork* should be at least $n \cdot \text{blocksize}$ if *side* = 'L' and at least $m \cdot \text{blocksize}$ if *side* = 'R', where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

The computed matrix *Q* differs from the exact result by a matrix *E* such that $\|E\|_2 = O(\epsilon) \|C\|_2$, where ϵ is the machine precision.

The approximate number of floating-point operations is

$8n(ihi-ilo)^2$ if *side* = 'L';

$8m(ihi-ilo)^2$ if *side* = 'R'.

The real counterpart of this routine is [?ormhr](#).

?gebal

Balances a general matrix to improve the accuracy of computed eigenvalues and eigenvectors.

```
call sgebal ( job, n, a, lda, ilo, ihi, scale, info )
call dgebal ( job, n, a, lda, ilo, ihi, scale, info )
call cgebal ( job, n, a, lda, ilo, ihi, scale, info )
call zgebal ( job, n, a, lda, ilo, ihi, scale, info )
```

Discussion

This routine *balances* a matrix A by performing either or both of the following two similarity transformations:

- (1) The routine first attempts to permute A to block upper triangular form:

$$PAP^T = A' = \begin{bmatrix} A'_{11} & A'_{12} & A'_{13} \\ 0 & A'_{22} & A'_{23} \\ 0 & 0 & A'_{33} \end{bmatrix}$$

where P is a permutation matrix, and A'_{11} and A'_{33} are upper triangular. The diagonal elements of A'_{11} and A'_{33} are eigenvalues of A . The rest of the eigenvalues of A are the eigenvalues of the central diagonal block A'_{22} , in rows and columns *ilo* to *ihi*. Subsequent operations to compute the eigenvalues of A (or its Schur factorization) need only be applied to these rows and columns; this can save a significant amount of work if *ilo* > 1 and *ihi* < n . If no suitable permutation exists (as is often the case), the routine sets *ilo* = 1 and *ihi* = n , and A'_{22} is the whole of A .

- (2) The routine applies a diagonal similarity transformation to A' , to make the rows and columns of A'_{22} as close in norm as possible:

$$A'' = DA'D^{-1} = \begin{bmatrix} I & 0 & 0 \\ 0 & D_{22} & 0 \\ 0 & 0 & I \end{bmatrix} \times \begin{bmatrix} A'_{11} & A'_{12} & A'_{13} \\ 0 & A'_{22} & A'_{23} \\ 0 & 0 & A'_{33} \end{bmatrix} \times \begin{bmatrix} I & 0 & 0 \\ 0 & D_{22}^{-1} & 0 \\ 0 & 0 & I \end{bmatrix}$$

This scaling can reduce the norm of the matrix (that is, $\|A'_{22}\| < \|A'_{22}\|$), and hence reduce the effect of rounding errors on the accuracy of computed eigenvalues and eigenvectors.

Input Parameters

job CHARACTER*1. Must be 'N' or 'P' or 'S' or 'B'.
 If *job* = 'N', then *A* is neither permuted nor scaled (but *ilo*, *ihi*, and *scale* get their values).
 If *job* = 'P', then *A* is permuted but not scaled.
 If *job* = 'S', then *A* is scaled but not permuted.
 If *job* = 'B', then *A* is both scaled and permuted.

n INTEGER. The order of the matrix *A* ($n \geq 0$).

a REAL for *sgebal*
 DOUBLE PRECISION for *dgebal*
 COMPLEX for *cgebal*
 DOUBLE COMPLEX for *zgebal*.
 Arrays:
a (*lda*, *) contains the matrix *A*.
 The second dimension of *a* must be at least $\max(1, n)$.
a is not referenced if *job* = 'N'.

lda INTEGER. The first dimension of *a*; at least $\max(1, n)$.

Output Parameters

a Overwritten by the balanced matrix (*a* is not referenced if *job* = 'N').

ilo, *ihi* INTEGER. The values *ilo* and *ihi* such that on exit *a*(*i*, *j*) is zero if $i > j$ and $1 \leq j < ilo$ or $ihi < i \leq n$.
 If *job* = 'N' or 'S', then *ilo* = 1 and *ihi* = *n*.

scale REAL for single-precision flavors
 DOUBLE PRECISION for double-precision flavors
 Array, DIMENSION at least $\max(1, n)$.
 Contains details of the permutations and scaling factors.

More precisely, if p_j is the index of the row and column interchanged with row and column j , and d_j is the scaling factor used to balance row and column j , then $scale(j) = p_j$ for $j = 1, 2, \dots, ilo-1, ihi+1, \dots, n$; $scale(j) = d_j$ for $j = ilo, ilo + 1, \dots, ihi$. The order in which the interchanges are made is n to $ihi+1$, then 1 to $ilo-1$.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = $-i$, the i th parameter had an illegal value.

Application Notes

The errors are negligible, compared with those in subsequent computations.

If the matrix A is balanced by this routine, then any eigenvectors computed subsequently are eigenvectors of the matrix A'' and hence you must call `?gebak` (see [page 5-193](#)) to transform them back to eigenvectors of A .

If the Schur vectors of A are required, do not call this routine with *job* = 'S' or 'B', because then the balancing transformation is not orthogonal (not unitary for complex flavors). If you call this routine with *job* = 'P', then any Schur vectors computed subsequently are Schur vectors of the matrix A'' , and you'll need to call `?gebak` (with *side* = 'R') to transform them back to Schur vectors of A .

The total number of floating-point operations is proportional to n^2 .

?gebak

Transforms eigenvectors of a balanced matrix to those of the original nonsymmetric matrix.

```
call sgebak ( job,side,n,ilo,ihi,scale,m,v,ldv,info )
call dgebak ( job,side,n,ilo,ihi,scale,m,v,ldv,info )
call cgebak ( job,side,n,ilo,ihi,scale,m,v,ldv,info )
call zgebak ( job,side,n,ilo,ihi,scale,m,v,ldv,info )
```

Discussion

This routine is intended to be used after a matrix A has been balanced by a call to [?gebal](#), and eigenvectors of the balanced matrix A'_{22} have subsequently been computed.

For a description of balancing, see [?gebal](#) (page 5-190). The balanced matrix A'' is obtained as $A'' = DPAP^TD^{-1}$, where P is a permutation matrix and D is a diagonal scaling matrix. This routine transforms the eigenvectors as follows:

if x is a right eigenvector of A'' , then $P^TD^{-1}x$ is a right eigenvector of A ;
 if x is a left eigenvector of A'' , then P^TDy is a left eigenvector of A .

Input Parameters

job CHARACTER*1. Must be 'N' or 'P' or 'S' or 'B'.
 The same parameter *job* as supplied to [?gebal](#).

side CHARACTER*1. Must be 'L' or 'R'.
 If *side* = 'L', then left eigenvectors are transformed.
 If *side* = 'R', then right eigenvectors are transformed.

n INTEGER. The number of rows of the matrix of eigenvectors ($n \geq 0$).

ilo, ihi INTEGER. The values *ilo* and *ihi*, as returned by [?gebal](#). (If $n > 0$, then $1 \leq ilo \leq ihi \leq n$; if $n = 0$, then *ilo* = 1 and *ihi* = 0.)

<i>scale</i>	<p>REAL for single-precision flavors DOUBLE PRECISION for double-precision flavors Array, DIMENSION at least $\max(1, n)$.</p> <p>Contains details of the permutations and/or the scaling factors used to balance the original general matrix, as returned by <code>?gebal</code>.</p>
<i>m</i>	<p>INTEGER. The number of columns of the matrix of eigenvectors ($m \geq 0$).</p>
<i>v</i>	<p>REAL for <code>sgebak</code> DOUBLE PRECISION for <code>dgebak</code> COMPLEX for <code>cgebak</code> DOUBLE COMPLEX for <code>zgebak</code>.</p> <p>Arrays: <i>v</i> (<i>ldv</i>, *) contains the matrix of left or right eigenvectors to be transformed. The second dimension of <i>v</i> must be at least $\max(1, m)$.</p>
<i>ldv</i>	<p>INTEGER. The first dimension of <i>v</i>; at least $\max(1, n)$.</p>

Output Parameters

<i>v</i>	Overwritten by the transformed eigenvectors.
<i>info</i>	<p>INTEGER.</p> <p>If <i>info</i> = 0, the execution is successful. If <i>info</i> = -<i>i</i>, the <i>i</i>th parameter had an illegal value.</p>

Application Notes

The errors in this routine are negligible.

The approximate number of floating-point operations is approximately proportional to $m * n$.

?hseqr

Computes all eigenvalues and (optionally) the Schur factorization of a matrix reduced to Hessenberg form.

```
call shseqr (job,compz,n,ilo,ihi,h,ldh,wr,wi,z,ldz,work,lwork,info)
call dhseqr (job,compz,n,ilo,ihi,h,ldh,wr,wi,z,ldz,work,lwork,info)
call chseqr (job,compz,n,ilo,ihi,h,ldh,w,z,ldz,work,lwork,info)
call zhseqr (job,compz,n,ilo,ihi,h,ldh,w,z,ldz,work,lwork,info)
```

Discussion

This routine computes all the eigenvalues, and optionally the Schur factorization, of an upper Hessenberg matrix H : $H = ZTZ^H$, where T is an upper triangular (or, for real flavors, quasi-triangular) matrix (the Schur form of H), and Z is the unitary or orthogonal matrix whose columns are the Schur vectors z_i .

You can also use this routine to compute the Schur factorization of a general matrix A which has been reduced to upper Hessenberg form H :

$A = QHQ^H$, where Q is unitary (orthogonal for real flavors);

$A = (QZ)T(QZ)^H$.

In this case, after reducing A to Hessenberg form by [?gehrd \(page 5-178\)](#), call [?orghr](#) to form Q explicitly ([page 5-180](#)) and then pass Q to [?hseqr](#) with `compz = 'V'`.

You can also call [?gebal \(page 5-190\)](#) to balance the original matrix before reducing it to Hessenberg form by [?hseqr](#), so that the Hessenberg matrix H will have the structure:

$$\begin{bmatrix} H_{11} & H_{12} & H_{13} \\ 0 & H_{22} & H_{23} \\ 0 & 0 & H_{33} \end{bmatrix}$$

where H_{11} and H_{33} are upper triangular.

If so, only the central diagonal block H_{22} (in rows and columns ilo to ihi) needs to be further reduced to Schur form (the blocks H_{12} and H_{23} are also affected). Therefore the values of ilo and ihi can be supplied to `?hseqr` directly. Also, after calling this routine you must call `?gebak` ([page 5-193](#)) to permute the Schur vectors of the balanced matrix to those of the original matrix.

If `?gebal` has not been called, however, then ilo must be set to 1 and ihi to n . Note that if the Schur factorization of A is required, `?gebal` must not be called with $job = 'S'$ or $'B'$, because the balancing transformation is not unitary (for real flavors, it is not orthogonal).

`?hseqr` uses a multishift form of the upper Hessenberg QR algorithm. The Schur vectors are normalized so that $\|z_i\|_2 = 1$, but are determined only to within a complex factor of absolute value 1 (for the real flavors, to within a factor ± 1).

Input Parameters

job	CHARACTER*1. Must be 'E' or 'S'. If $job = 'E'$, then eigenvalues only are required. If $job = 'S'$, then the Schur form T is required.
$compz$	CHARACTER*1. Must be 'N' or 'I' or 'V'. If $compz = 'N'$, then no Schur vectors are computed (and the array z is not referenced). If $compz = 'I'$, then the Schur vectors of H are computed (and the array z is initialized by the routine). If $compz = 'V'$, then the Schur vectors of A are computed (and the array z must contain the matrix Q on entry).
n	INTEGER. The order of the matrix H ($n \geq 0$).
ilo, ihi	INTEGER. If A has been balanced by <code>?gebal</code> , then ilo and ihi must contain the values returned by <code>?gebal</code> . Otherwise, ilo must be set to 1 and ihi to n .
$h, z, work$	REAL for <code>shseqr</code> DOUBLE PRECISION for <code>dhseqr</code> COMPLEX for <code>chseqr</code> DOUBLE COMPLEX for <code>zhseqr</code> .

Arrays:

$h(ldh, *)$ The n by n upper Hessenberg matrix H .
The second dimension of h must be at least $\max(1, n)$.

$z(ldz, *)$

If $compz = 'V'$, then z must contain the matrix Q from the reduction to Hessenberg form.

If $compz = 'I'$, then z need not be set.

If $compz = 'N'$, then z is not referenced.

The second dimension of z must be
at least $\max(1, n)$ if $compz = 'V'$ or $'I'$;
at least 1 if $compz = 'N'$.

$work(lwork)$ is a workspace array.

The dimension of $work$ must be at least $\max(1, n)$.

ldh INTEGER. The first dimension of h ; at least $\max(1, n)$.

ldz INTEGER. The first dimension of z ;
If $compz = 'N'$, then $ldz \geq 1$.
If $compz = 'V'$ or $'I'$, then $ldz \geq \max(1, n)$.

$lwork$ INTEGER. The dimension of the array $work$.
 $lwork \geq \max(1, n)$. If $lwork = -1$, then a workspace query is assumed; the routine only calculates the optimal size of the $work$ array, returns this value as the first entry of the $work$ array, and no error message related to $lwork$ is issued by `xerbla`.

Output Parameters

w COMPLEX for `chseqqr`
DOUBLE COMPLEX for `zhseqqr`.
Array, DIMENSION at least $\max(1, n)$.
Contains the computed eigenvalues, unless $info > 0$. The eigenvalues are stored in the same order as on the diagonal of the Schur form T (if computed).

wr, wi REAL for `shseqqr`
DOUBLE PRECISION for `dhseqqr`
Arrays, DIMENSION at least $\max(1, n)$ each.
Contain the real and imaginary parts, respectively, of the

computed eigenvalues, unless *info* > 0. Complex conjugate pairs of eigenvalues appear consecutively with the eigenvalue having positive imaginary part first. The eigenvalues are stored in the same order as on the diagonal of the Schur form *T* (if computed).

z If *compz* = 'V' or 'I', then *z* contains the unitary (orthogonal) matrix of the required Schur vectors, unless *info* > 0.
 If *compz* = 'N', then *z* is not referenced.

work(1) On exit, if *info* = 0, then *work(1)* returns the optimal *lwork*.

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.
 If *info* > 0, the algorithm has failed to find all the eigenvalues after a total 30(*ihi*-*ilo*+1) iterations. If *info* = *i*, elements 1,2, ..., *ilo*-1 and *i*+1, *i*+2, ..., *n* of *wr* and *wi* contain the real and imaginary parts of the eigenvalues which have been found.

Application Notes

The computed Schur factorization is the exact factorization of a nearby matrix $H + E$, where $\|E\|_2 < O(\epsilon) \|H\|_2 / s_i$, and ϵ is the machine precision. If λ_i is an exact eigenvalue, and μ_i is the corresponding computed value, then $|\lambda_i - \mu_i| \leq c(n)\epsilon \|H\|_2 / s_i$ where $c(n)$ is a modestly increasing function of n , and s_i is the reciprocal condition number of λ_i . You can compute the condition numbers s_i by calling `?trsna` (see [page 5-210](#)).

The total number of floating-point operations depends on how rapidly the algorithm converges; typical numbers are as follows.

If only eigenvalues are computed:	$7n^3$ for real flavors
	$25n^3$ for complex flavors.
If the Schur form is computed:	$10n^3$ for real flavors
	$35n^3$ for complex flavors.
If the full Schur factorization is computed:	$20n^3$ for real flavors
	$70n^3$ for complex flavors.

?hsein

Computes selected eigenvectors of an upper Hessenberg matrix that correspond to specified eigenvalues.

```
call shsein ( job, eigsrc, initv, select, n, h, ldh, wr, wi, vl,
             ldvl, vr, ldvr, mm, m, work, ifaill, ifailr, info )
call dhsein ( job, eigsrc, initv, select, n, h, ldh, wr, wi, vl,
             ldvl, vr, ldvr, mm, m, work, ifaill, ifailr, info )
call chsein ( job, eigsrc, initv, select, n, h, ldh, w, vl,
             ldvl, vr, ldvr, mm, m, work, rwork, ifaill, ifailr, info )
call zhsein ( job, eigsrc, initv, select, n, h, ldh, w, vl,
             ldvl, vr, ldvr, mm, m, work, rwork, ifaill, ifailr, info )
```

Discussion

This routine computes left and/or right eigenvectors of an upper Hessenberg matrix H , corresponding to selected eigenvalues.

The right eigenvector x and the left eigenvector y , corresponding to an eigenvalue λ , are defined by: $Hx = \lambda x$ and $y^H H = \lambda y^H$ (or $H^H y = \lambda^* y$). Here λ^* denotes the conjugate of λ .

The eigenvectors are computed by inverse iteration. They are scaled so that, for a real eigenvector x , $\max |x_i| = 1$, and for a complex eigenvector, $\max(|\operatorname{Re}x_i| + |\operatorname{Im}x_i|) = 1$.

If H has been formed by reduction of a general matrix A to upper Hessenberg form, then eigenvectors of H may be transformed to eigenvectors of A by [?ormhr \(page 5-182\)](#) or [?unmhr \(page 5-187\)](#).

Input Parameters

job CHARACTER*1. Must be 'R' or 'L' or 'B'.
 If *job* = 'R', then only right eigenvectors are computed.
 If *job* = 'L', then only left eigenvectors are computed.
 If *job* = 'B', then all eigenvectors are computed.

<i>eigsrc</i>	<p>CHARACTER*1. Must be 'Q' or 'N'.</p> <p>If <i>eigsrc</i> = 'Q', then the eigenvalues of H were found using ?hseqr (see page 5-195); thus if H has any zero sub-diagonal elements (and so is block triangular), then the jth eigenvalue can be assumed to be an eigenvalue of the block containing the jth row/column. This property allows the routine to perform inverse iteration on just one diagonal block.</p> <p>If <i>eigsrc</i> = 'N', then no such assumption is made and the routine performs inverse iteration using the whole matrix.</p>
<i>initv</i>	<p>CHARACTER*1. Must be 'N' or 'U'.</p> <p>If <i>initv</i> = 'N', then no initial estimates for the selected eigenvectors are supplied.</p> <p>If <i>initv</i> = 'U', then initial estimates for the selected eigenvectors are supplied in <i>vl</i> and/or <i>vr</i>.</p>
<i>select</i>	<p>LOGICAL.</p> <p>Array, DIMENSION at least max (1, n).</p> <p>Specifies which eigenvectors are to be computed.</p> <p><i>For real flavors:</i></p> <p>To obtain the real eigenvector corresponding to the real eigenvalue $wr(j)$, set <i>select(j)</i> to .TRUE.</p> <p>To select the complex eigenvector corresponding to the complex eigenvalue $(wr(j), wi(j))$ with complex conjugate $(wr(j+1), wi(j+1))$, set <i>select(j)</i> and/or <i>select(j+1)</i> to .TRUE.; the eigenvector corresponding to the first eigenvalue in the pair is computed.</p> <p><i>For complex flavors:</i></p> <p>To select the eigenvector corresponding to the eigenvalue $w(j)$, set <i>select(j)</i> to .TRUE.</p>
<i>n</i>	<p>INTEGER. The order of the matrix H ($n \geq 0$).</p>
<i>h, vl, vr, work</i>	<p>REAL for shsein DOUBLE PRECISION for dhsein COMPLEX for chsein DOUBLE COMPLEX for zhsein.</p>

Arrays:

h(ldh,)* The n by n upper Hessenberg matrix H .
The second dimension of h must be at least $\max(1, n)$.

vl(ldvl,)*

If *initv* = 'V' and *job* = 'L' or 'B', then *vl* must contain starting vectors for inverse iteration for the left eigenvectors. Each starting vector must be stored in the same column or columns as will be used to store the corresponding eigenvector.

If *initv* = 'N', then *vl* need not be set.

The second dimension of *vl* must be at least $\max(1, mm)$ if *job* = 'L' or 'B' and at least 1 if *job* = 'R'.

The array *vl* is not referenced if *job* = 'R'.

vr(ldvr,)*

If *initv* = 'V' and *job* = 'R' or 'B', then *vr* must contain starting vectors for inverse iteration for the right eigenvectors. Each starting vector must be stored in the same column or columns as will be used to store the corresponding eigenvector.

If *initv* = 'N', then *vr* need not be set.

The second dimension of *vr* must be at least $\max(1, mm)$ if *job* = 'R' or 'B' and at least 1 if *job* = 'L'.

The array *vr* is not referenced if *job* = 'L'.

work()* is a workspace array.

DIMENSION at least $\max(1, n*(n+2))$ for real flavors and at least $\max(1, n*n)$ for complex flavors.

ldh **INTEGER**. The first dimension of h ; at least $\max(1, n)$.

w **COMPLEX** for *chsein*

DOUBLE COMPLEX for *zhsein*.

Array, **DIMENSION** at least $\max(1, n)$.

Contains the eigenvalues of the matrix H .

If *eigsrc* = 'Q', the array must be exactly as returned by *?hseqr*.

<i>wr, wi</i>	<p>REAL for <i>shsein</i> DOUBLE PRECISION for <i>dhsein</i> Arrays, DIMENSION at least max (1, <i>n</i>) each. Contain the real and imaginary parts, respectively, of the eigenvalues of the matrix <i>H</i>. Complex conjugate pairs of values must be stored in consecutive elements of the arrays. If <i>eigsrc</i> = 'Q', the arrays must be exactly as returned by <i>?hseqr</i>.</p>
<i>ldvl</i>	<p>INTEGER. The first dimension of <i>vl</i>. If <i>job</i> = 'L' or 'B', <i>ldvl</i> ≥ max(1,<i>n</i>). If <i>job</i> = 'R', <i>ldvl</i> ≥ 1.</p>
<i>ldvr</i>	<p>INTEGER. The first dimension of <i>vr</i>. If <i>job</i> = 'R' or 'B', <i>ldvr</i> ≥ max(1,<i>n</i>). If <i>job</i> = 'L', <i>ldvr</i> ≥ 1.</p>
<i>mm</i>	<p>INTEGER. The number of columns in <i>vl</i> and/or <i>vr</i>. Must be at least <i>m</i>, the actual number of columns required (see <i>Output Parameters</i> below). For real flavors, <i>m</i> is obtained by counting 1 for each selected real eigenvector and 2 for each selected complex eigenvector (see <i>select</i>). For complex flavors, <i>m</i> is the number of selected eigenvectors (see <i>select</i>). Constraint: 0 ≤ <i>mm</i> ≤ <i>n</i>.</p>
<i>rwork</i>	<p>REAL for <i>chsein</i> DOUBLE PRECISION for <i>zhsein</i>. Array, DIMENSION at least max (1, <i>n</i>).</p>

Output Parameters

<i>select</i>	<p>Overwritten for real flavors only. If a complex eigenvector was selected as specified above, then <i>select</i>(<i>j</i>) is set to <i>.TRUE.</i> and <i>select</i>(<i>j</i>+1) to <i>.FALSE.</i></p>
<i>w</i>	<p>The real parts of some elements of <i>w</i> may be modified, as close eigenvalues are perturbed slightly in searching for independent eigenvectors.</p>

- wr* Some elements of *wr* may be modified, as close eigenvalues are perturbed slightly in searching for independent eigenvectors.
- v1, vr* If *job* = 'L' or 'B', *v1* contains the computed left eigenvectors (as specified by *select*).
If *job* = 'R' or 'B', *vr* contains the computed right eigenvectors (as specified by *select*).
The eigenvectors are stored consecutively in the columns of the array, in the same order as their eigenvalues.
For real flavors: a real eigenvector corresponding to a selected real eigenvalue occupies one column; a complex eigenvector corresponding to a selected complex eigenvalue occupies two columns: the first column holds the real part and the second column holds the imaginary part.
- m* **INTEGER.** *For real flavors*: the number of columns of *v1* and/or *vr* required to store the selected eigenvectors.
For complex flavors: the number of selected eigenvectors.
- ifaill, ifailr* **INTEGER.**
Arrays, **DIMENSION** at least $\max(1, mm)$ each.
ifaill(*i*) = 0 if the *i*th column of *v1* converged;
ifaill(*i*) = *j* > 0 if the eigenvector stored in the *i*th column of *v1* (corresponding to the *j*th eigenvalue) failed to converge.
ifailr(*i*) = 0 if the *i*th column of *vr* converged;
ifailr(*i*) = *j* > 0 if the eigenvector stored in the *i*th column of *vr* (corresponding to the *j*th eigenvalue) failed to converge.
For real flavors: if the *i*th and (*i*+1)th columns of *v1* contain a selected complex eigenvector, then *ifaill*(*i*) and *ifaill*(*i*+1) are set to the same value. A similar rule holds for *vr* and *ifailr*.
The array *ifaill* is not referenced if *job* = 'R'.
The array *ifailr* is not referenced if *job* = 'L'.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

If *info* > 0, then *i* eigenvectors (as indicated by the parameters *ifaill* and/or *ifailr* above) failed to converge. The corresponding columns of *vl* and/or *vr* contain no useful information.

Application Notes

Each computed right eigenvector x_i is the exact eigenvector of a nearby matrix $A + E_i$, such that $\|E_i\| < O(\epsilon)\|A\|$. Hence the residual is small: $\|Ax_i - \lambda_i x_i\| = O(\epsilon)\|A\|$.

However, eigenvectors corresponding to close or coincident eigenvalues may not accurately span the relevant subspaces.

Similar remarks apply to computed left eigenvectors.

?trevc

Computes selected eigenvectors of an upper (quasi-) triangular matrix computed by ?hseqr.

```
call strevc ( side, howmny, select, n, t, ldt, vl, ldvl, vr, ldvr,
             mm, m, work, info )
call dtrevc ( side, howmny, select, n, t, ldt, vl, ldvl, vr, ldvr,
             mm, m, work, info )
call ctrevc ( side, howmny, select, n, t, ldt, vl, ldvl, vr, ldvr,
             mm, m, work, rwork, info )
call ztrevc ( side, howmny, select, n, t, ldt, vl, ldvl, vr, ldvr,
             mm, m, work, rwork, info )
```

Discussion

This routine computes some or all of the right and/or left eigenvectors of an upper triangular matrix T (or, for real flavors, an upper quasi-triangular matrix T). Matrices of this type are produced by the Schur factorization of a general matrix: $A = QTQ^H$, as computed by ?hseqr (see [page 5-195](#)).

The right eigenvector x and the left eigenvector y of T corresponding to an eigenvalue w , are defined by:

$$Tx = wx, \quad y^H T = w y^H$$

where y^H denotes the conjugate transpose of y .

The eigenvalues are not input to this routine, but are read directly from the diagonal blocks of T .

This routine returns the matrices X and/or Y of right and left eigenvectors of T , or the products QX and/or QY , where Q is an input matrix.

If Q is the orthogonal/unitary factor that reduces a matrix A to Schur form T , then QX and QY are the matrices of right and left eigenvectors of A .

Input Parameters

<i>side</i>	<p>CHARACTER*1. Must be 'R' or 'L' or 'B'.</p> <p>If <i>side</i> = 'R', then only right eigenvectors are computed.</p> <p>If <i>side</i> = 'L', then only left eigenvectors are computed.</p> <p>If <i>side</i> = 'B', then all eigenvectors are computed.</p>
<i>howmny</i>	<p>CHARACTER*1. Must be 'A' or 'B' or 'S'.</p> <p>If <i>howmny</i> = 'A', then all eigenvectors (as specified by <i>side</i>) are computed.</p> <p>If <i>howmny</i> = 'B', then all eigenvectors (as specified by <i>side</i>) are computed and backtransformed by the matrices supplied in <i>vl</i> and <i>vr</i>.</p> <p>If <i>howmny</i> = 'S', then selected eigenvectors (as specified by <i>side</i> and <i>select</i>) are computed.</p>
<i>select</i>	<p>LOGICAL.</p> <p>Array, DIMENSION at least max (1, <i>n</i>).</p> <p>If <i>howmny</i> = 'S', <i>select</i> specifies which eigenvectors are to be computed.</p> <p>If <i>howmny</i> = 'A' or 'B', <i>select</i> is not referenced.</p> <p><i>For real flavors:</i></p> <p>If ω is a real eigenvalue, the corresponding real eigenvector is computed if <i>select</i>(<i>j</i>) is .TRUE..</p> <p>If ω and ω_{+I} are the real and imaginary parts of a complex eigenvalue, the corresponding complex eigenvector is computed if either <i>select</i>(<i>j</i>) or <i>select</i>(<i>j</i>+1) is .TRUE., and on exit <i>select</i>(<i>j</i>) is set to .TRUE. and <i>select</i>(<i>j</i>+1) is set to .FALSE..</p> <p><i>For complex flavors:</i></p> <p>The eigenvector corresponding to the <i>j</i>-th eigenvalue is computed if <i>select</i>(<i>j</i>) is .TRUE..</p>
<i>n</i>	<p>INTEGER. The order of the matrix T ($n \geq 0$).</p>
<i>t, vl, vr, work</i>	<p>REAL for <i>strevc</i></p> <p>DOUBLE PRECISION for <i>dtrevc</i></p> <p>COMPLEX for <i>ctrevc</i></p> <p>DOUBLE COMPLEX for <i>ztrevc</i>.</p> <p>Arrays:</p>

$t(ldt,*)$ contains the n by n matrix T in Schur canonical form.

The second dimension of t must be at least $\max(1, n)$.

$vl(ldvl,*)$

If $howmny = 'B'$ and $side = 'L'$ or $'B'$, then vl must contain an n by n matrix Q (usually the matrix of Schur vectors returned by `?hseqr`).

If $howmny = 'A'$ or $'S'$, then vl need not be set.

The second dimension of vl must be at least $\max(1, mm)$

if $side = 'L'$ or $'B'$ and at least 1 if $side = 'R'$.

The array vl is not referenced if $side = 'R'$.

$vr(ldvr,*)$

If $howmny = 'B'$ and $side = 'R'$ or $'B'$, then vr must contain an n by n matrix Q (usually the matrix of Schur vectors returned by `?hseqr`).

If $howmny = 'A'$ or $'S'$, then vr need not be set.

The second dimension of vr must be at least $\max(1, mm)$

if $side = 'R'$ or $'B'$ and at least 1 if $side = 'L'$.

The array vr is not referenced if $side = 'L'$.

$work(*)$ is a workspace array.

DIMENSION at least $\max(1, 3*n)$ for real flavors and at least $\max(1, 2*n)$ for complex flavors.

ldt **INTEGER.** The first dimension of t ; at least $\max(1, n)$.

$ldvl$ **INTEGER.** The first dimension of vl .
If $side = 'L'$ or $'B'$, $ldvl \geq \max(1, n)$.
If $side = 'R'$, $ldvl \geq 1$.

$ldvr$ **INTEGER.** The first dimension of vr .
If $side = 'R'$ or $'B'$, $ldvr \geq \max(1, n)$.
If $side = 'L'$, $ldvr \geq 1$.

mm **INTEGER.** The number of columns in the arrays vl and/or vr . Must be at least m (the precise number of columns required). If $howmny = 'A'$ or $'B'$, $m = n$.
If $howmny = 'S'$: for real flavors, m is obtained by counting 1 for each selected real eigenvector and 2 for each selected complex eigenvector;

for complex flavors, m is the number of selected eigenvectors (see *select*). Constraint: $0 \leq m \leq n$.

rwork REAL for *ctrevc*
DOUBLE PRECISION for *ztrevc*.
Workspace array, DIMENSION at least max (1, n).

Output Parameters

select If a complex eigenvector of a real matrix was selected as specified above, then *select*(j) is set to *.TRUE.* and *select*($j+1$) to *.FALSE.*

v1, vr If *side* = 'L' or 'B', *v1* contains the computed left eigenvectors (as specified by *howmny* and *select*). If *side* = 'R' or 'B', *vr* contains the computed right eigenvectors (as specified by *howmny* and *select*).

The eigenvectors are stored consecutively in the columns of the array, in the same order as their eigenvalues.

For real flavors: corresponding to each real eigenvalue is a real eigenvector, occupying one column; corresponding to each complex conjugate pair of eigenvalues is a complex eigenvector, occupying two columns; the first column holds the real part and the second column holds the imaginary part.

m INTEGER.
For complex flavors: the number of selected eigenvectors. If *howmny* = 'A' or 'B', m is set to n .
For real flavors: the number of columns of *v1* and/or *vr* actually used to store the selected eigenvectors.
If *howmny* = 'A' or 'B', m is set to n .

info INTEGER. If *info* = 0, the execution is successful. If *info* = $-i$, the i th parameter had an illegal value.

Application Notes

If x_i is an exact right eigenvector and y_i is the corresponding computed eigenvector, then the angle $\theta(y_i, x_i)$ between them is bounded as follows:
 $\theta(y_i, x_i) \leq (c(n)\epsilon \|T\|_2) / \text{sep}_i$ where sep_i is the reciprocal condition number of x_i . The condition number sep_i may be computed by calling `?trsna`.

?trsna

Estimates condition numbers for specified eigenvalues and right eigenvectors of an upper (quasi-) triangular matrix.

```
call strсна ( job, howmny, select, n, t, ldt, vl, ldvl, vr, ldvr,
             s, sep, mm, m, work, ldwork, iwork, info )
call dtrsna ( job, howmny, select, n, t, ldt, vl, ldvl, vr, ldvr,
             s, sep, mm, m, work, ldwork, iwork, info )
call ctrсна ( job, howmny, select, n, t, ldt, vl, ldvl, vr, ldvr,
             s, sep, mm, m, work, ldwork, rwork, info )
call ztrsna ( job, howmny, select, n, t, ldt, vl, ldvl, vr, ldvr,
             s, sep, mm, m, work, ldwork, rwork, info )
```

Discussion

This routine estimates condition numbers for specified eigenvalues and/or right eigenvectors of an upper triangular matrix T (or, for real flavors, upper quasi-triangular matrix T in canonical Schur form). These are the same as the condition numbers of the eigenvalues and right eigenvectors of an original matrix $A = ZTZ^H$ (with unitary or, for real flavors, orthogonal Z), from which T may have been derived.

The routine computes the reciprocal of the condition number of an eigenvalue λ_i as $s_i = |v^H u| / (|u| |E| + |v| |E|)$, where u and v are the right and left eigenvectors of T , respectively, corresponding to λ_i . This reciprocal condition number always lies between zero (ill-conditioned) and one (well-conditioned).

An approximate error estimate for a computed eigenvalue λ_i is then given by $\epsilon |T| / s_i$, where ϵ is the *machine precision*.

To estimate the reciprocal of the condition number of the right eigenvector corresponding to λ_i , the routine first calls `?trexc` (see [page 5-215](#)) to reorder the eigenvalues so that λ_i is in the leading position:

$$T = Q \begin{bmatrix} \lambda_i & C^H \\ 0 & T_{22} \end{bmatrix} Q^H$$

The reciprocal condition number of the eigenvector is then estimated as sep_i , the smallest singular value of the matrix $T_{22} - \lambda_i I$. This number ranges from zero (ill-conditioned) to very large (well-conditioned).

An approximate error estimate for a computed right eigenvector u corresponding to λ_i is then given by $\epsilon \|T\| / sep_i$.

Input Parameters

- job* CHARACTER*1. Must be 'E' or 'V' or 'B'.
 If *job* = 'E', then condition numbers for eigenvalues only are computed.
 If *job* = 'V', then condition numbers for eigenvectors only are computed.
 If *job* = 'B', then condition numbers for both eigenvalues and eigenvectors are computed.
- howmny* CHARACTER*1. Must be 'A' or 'S'.
 If *howmny* = 'A', then the condition numbers for all eigenpairs are computed.
 If *howmny* = 'S', then condition numbers for selected eigenpairs (as specified by *select*) are computed.
- select* LOGICAL.
 Array, DIMENSION at least max(1, *n*) if *howmny* = 'S' and at least 1 otherwise.
 Specifies the eigenpairs for which condition numbers are to be computed if *howmny* = 'S'.
 For real flavors:
 To select condition numbers for the eigenpair corresponding to the real eigenvalue λ_j , *select*(*j*) must be set .TRUE.; to select condition numbers for the

eigenpair corresponding to a complex conjugate pair of eigenvalues λ_j and λ_{j+1} , *select*(*j*) and/or *select*(*j*+1) must be set **.TRUE.**

For complex flavors:

To select condition numbers for the eigenpair corresponding to the eigenvalue λ_j , *select*(*j*) must be set **.TRUE.**

select is not referenced if *howmny* = 'A'.

n **INTEGER.** The order of the matrix *T* ($n \geq 0$).

t, vl, vr, work **REAL** for *strsna*
DOUBLE PRECISION for *dtrsna*
COMPLEX for *ctrsna*
DOUBLE COMPLEX for *ztrsna*.

Arrays:

t(*ldt*, *) contains the *n* by *n* matrix *T*.

The second dimension of *t* must be at least max(1, *n*).

vl(*ldvl*, *)

If *job* = 'E' or 'B', then *vl* must contain the left eigenvectors of *T* (or of any matrix QTQ^H with *Q* unitary or orthogonal) corresponding to the eigenpairs specified by *howmny* and *select*. The eigenvectors must be stored in consecutive columns of *vl*, as returned by ?*trevc* or ?*hsein*.

The second dimension of *vl* must be at least max(1, *mm*) if *job* = 'E' or 'B' and at least 1 if *job* = 'V'.

The array *vl* is not referenced if *job* = 'V'.

vr(*ldvr*, *)

If *job* = 'E' or 'B', then *vr* must contain the right eigenvectors of *T* (or of any matrix QTQ^H with *Q* unitary or orthogonal) corresponding to the eigenpairs specified by *howmny* and *select*. The eigenvectors must be stored in consecutive columns of *vr*, as returned by ?*trevc* or ?*hsein*.

The second dimension of *vr* must be at least max(1, *mm*) if *job* = 'E' or 'B' and at least 1 if *job* = 'V'.

The array *vr* is not referenced if *job* = 'V'.

work(ldwork,)* is a workspace array.
 The second dimension of *work* must be
 at least $\max(1, n+1)$ for complex flavors and
 at least $\max(1, n+6)$ for real flavors if *job*='V' or 'B';
 at least 1 if *job*='E'.
 The array *work* is not referenced if *job*='E'.

ldt INTEGER. The first dimension of *t*; at least $\max(1, n)$.

ldvl INTEGER. The first dimension of *vl*.
 If *job*='E' or 'B', $ldvl \geq \max(1, n)$.
 If *job*='V', $ldvl \geq 1$.

ldvr INTEGER. The first dimension of *vr*.
 If *job*='E' or 'B', $ldvr \geq \max(1, n)$.
 If *job*='R', $ldvr \geq 1$.

mm INTEGER. The number of elements in the arrays *s* and
sep, and the number of columns in *vl* and *vr* (if used).
 Must be at least *m* (the precise number required).
 If *howmny*='A', $m = n$;
 if *howmny*='S', for real flavors *m* is obtained by
 counting 1 for each selected real eigenvalue and 2 for
 each selected complex conjugate pair of eigenvalues.
 for complex flavors *m* is the number of selected
 eigenpairs (see *select*). Constraint: $0 \leq m \leq n$.

ldwork INTEGER. The first dimension of *work*.
 If *job*='V' or 'B', $ldwork \geq \max(1, n)$.
 If *job*='E', $ldwork \geq 1$.

rwork REAL for *ctrсна*, *ztrsна*.
 Array, DIMENSION at least $\max(1, n)$.

iwork INTEGER for *strсна*, *dtrsна*.
 Array, DIMENSION at least $\max(1, n)$.

Output Parameters

s REAL for single-precision flavors
 DOUBLE PRECISION for double-precision flavors.
 Array, DIMENSION at least $\max(1, mm)$ if *job*='E' or
 'B' and at least 1 if *job*='V'.

Contains the reciprocal condition numbers of the selected eigenvalues if *job* = 'E' or 'B', stored in consecutive elements of the array. Thus *s*(*j*), *sep*(*j*) and the *j*th columns of *v1* and *vr* all correspond to the same eigenpair (but not in general the *j*th eigenpair unless all eigenpairs have been selected). *For real flavors*: For a complex conjugate pair of eigenvalues, two consecutive elements of *S* are set to the same value.

The array *s* is not referenced if *job* = 'V'.

sep

REAL for single-precision flavors

DOUBLE PRECISION for double-precision flavors.

Array, DIMENSION at least max(1, *mm*)

if *job* = 'V' or 'B' and at least 1 if *job* = 'E'.

Contains the estimated reciprocal condition numbers of the selected right eigenvectors if *job* = 'V' or 'B', stored in consecutive elements of the array.

For real flavors: for a complex eigenvector, two consecutive elements of *sep* are set to the same value; if the eigenvalues cannot be reordered to compute *sep*(*j*), then *sep*(*j*) is set to zero; this can only occur when the true value would be very small anyway.

The array *sep* is not referenced if *job* = 'E'.

m

INTEGER.

For complex flavors: the number of selected eigenpairs.

If *howmny* = 'A', *m* is set to *n*.

For real flavors: the number of elements of *s* and/or *sep* actually used to store the estimated condition numbers.

If *howmny* = 'A', *m* is set to *n*.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The computed values *sep*_{*i*} may overestimate the true value, but seldom by a factor of more than 3.

?trexc

Reorders the Schur factorization of a general matrix.

```

call strexc ( compq, n, t, ldt, q, ldq, ifst, ilst, work, info )
call dtrexc ( compq, n, t, ldt, q, ldq, ifst, ilst, work, info )
call ctrexc ( compq, n, t, ldt, q, ldq, ifst, ilst, info )
call ztrexc ( compq, n, t, ldt, q, ldq, ifst, ilst, info )

```

Discussion

This routine reorders the Schur factorization of a general matrix $A = QTQ^H$, so that the diagonal element or block of T with row index *ifst* is moved to row *ilst*.

The reordered Schur form S is computed by an unitary (or, for real flavors, orthogonal) similarity transformation: $S = Z^H T Z$. Optionally the updated matrix P of Schur vectors is computed as $P = QZ$, giving $A = PSP^H$.

Input Parameters

compq CHARACTER*1. Must be 'V' or 'N'.
 If *compq* = 'V', then the Schur vectors (Q) are updated.
 If *compq* = 'N', then no Schur vectors are updated.

n INTEGER. The order of the matrix T ($n \geq 0$).

t, *q* REAL for *strexc*
 DOUBLE PRECISION for *dtrexc*
 COMPLEX for *ctrexc*
 DOUBLE COMPLEX for *ztrexc*.

Arrays:
t(*ldt*,*) contains the n by n matrix T .
 The second dimension of *t* must be at least $\max(1, n)$.

q(*ldq*,*)
 If *compq* = 'V', then *q* must contain Q (Schur vectors).
 If *compq* = 'N', then *q* is not referenced.

	The second dimension of q must be at least $\max(1, n)$ if $compq = 'V'$ and at least 1 if $compq = 'N'$.
ldt	INTEGER. The first dimension of t ; at least $\max(1, n)$.
ldq	INTEGER. The first dimension of q ; If $compq = 'N'$, then $ldq \geq 1$. If $compq = 'V'$, then $ldq \geq \max(1, n)$.
$ifst, ilst$	INTEGER. $1 \leq ifst \leq n$; $1 \leq ilst \leq n$. Must specify the reordering of the diagonal elements (or blocks, which is possible for real flavors) of the matrix T . The element (or block) with row index $ifst$ is moved to row $ilst$ by a sequence of exchanges between adjacent elements (or blocks).
$work$	REAL for <code>strex</code> DOUBLE PRECISION for <code>dtrex</code> . Array, DIMENSION at least $\max(1, n)$.

Output Parameters

t	Overwritten by the updated matrix S .
q	If $compq = 'V'$, q contains the updated matrix of Schur vectors.
$ifst, ilst$	Overwritten for real flavors only. If $ifst$ pointed to the second row of a 2 by 2 block on entry, it is changed to point to the first row; $ilst$ always points to the first row of the block in its final position (which may differ from its input value by ± 1).
$info$	INTEGER. If $info = 0$, the execution is successful. If $info = -i$, the i th parameter had an illegal value.

Application Notes

The computed matrix S is exactly similar to a matrix $T + E$, where $\|E\|_2 = O(\epsilon) \|T\|_2$, and ϵ is the machine precision.

Note that if a 2 by 2 diagonal block is involved in the re-ordering, its off-diagonal elements are in general changed; the diagonal elements and the eigenvalues of the block are unchanged unless the block is sufficiently ill-conditioned, in which case they may be noticeably altered. It is possible for a 2 by 2 block to break into two 1 by 1 blocks, that is, for a pair of complex eigenvalues to become purely real.

The values of eigenvalues however are never changed by the re-ordering.

The approximate number of floating-point operations is

for real flavors: $6n(\text{ifst-ilst})$ if *compq* = 'N';
 $12n(\text{ifst-ilst})$ if *compq* = 'V';

for complex flavors: $20n(\text{ifst-ilst})$ if *compq* = 'N';
 $40n(\text{ifst-ilst})$ if *compq* = 'V'.

?trsen

Reorders the Schur factorization of a matrix and (optionally) computes the reciprocal condition numbers and invariant subspace for the selected cluster of eigenvalues.

```
call strsen (job, compq, select, n, t, ldt, q, ldq, wr, wi, m, s,
            sep, work, lwork, iwork, liwork, info)
call dtrsen (job, compq, select, n, t, ldt, q, ldq, wr, wi, m, s,
            sep, work, lwork, iwork, liwork, info)
call ctrsen (job, compq, select, n, t, ldt, q, ldq, w, m, s,
            sep, work, lwork, info)
call ztrsen (job, compq, select, n, t, ldt, q, ldq, w, m, s,
            sep, work, lwork, info)
```

Discussion

This routine reorders the Schur factorization of a general matrix $A = QTQ^H$ so that a selected cluster of eigenvalues appears in the leading diagonal elements (or, for real flavors, diagonal blocks) of the Schur form.

The reordered Schur form R is computed by a unitary(orthogonal) similarity transformation: $R = Z^H T Z$. Optionally the updated matrix P of Schur vectors is computed as $P = QZ$, giving $A = PRP^H$.

Let

$$R = \begin{bmatrix} T_{11} & T_{12} \\ 0 & T_{13} \end{bmatrix}$$

where the selected eigenvalues are precisely the eigenvalues of the leading m by m submatrix T_{11} . Let P be correspondingly partitioned as $(Q_1 Q_2)$ where Q_1 consists of the first m columns of Q . Then $AQ_1 = Q_1 T_{11}$, and so the m columns of Q_1 form an orthonormal basis for the invariant subspace corresponding to the selected cluster of eigenvalues.

Optionally the routine also computes estimates of the reciprocal condition numbers of the average of the cluster of eigenvalues and of the invariant subspace.

Input Parameters

- job* CHARACTER*1. Must be 'N' or 'E' or 'V' or 'B'.
 If *job* = 'N', then no condition numbers are required.
 If *job* = 'E', then only the condition number for the cluster of eigenvalues is computed.
 If *job* = 'V', then only the condition number for the invariant subspace is computed.
 If *job* = 'B', then condition numbers for both the cluster and the invariant subspace are computed.
- compq* CHARACTER*1. Must be 'V' or 'N'.
 If *compq* = 'V', then Q of the Schur vectors is updated.
 If *compq* = 'N', then no Schur vectors are updated.
- select* LOGICAL.
 Array, DIMENSION at least max(1, n).
 Specifies the eigenvalues in the selected cluster.
 To select an eigenvalue λ_j , *select*(j) must be .TRUE..
 For real flavors: to select a complex conjugate pair of eigenvalues λ_j and λ_{j+1} (corresponding 2 by 2 diagonal

block), *select(j)* and/or *select(j+1)* must be `.TRUE.`; the complex conjugate λ_j and λ_{j+1} must be either both included in the cluster or both excluded.

n INTEGER. The order of the matrix T ($n \geq 0$).

t, *q*, *work* REAL for *strsen*
 DOUBLE PRECISION for *dtrsen*
 COMPLEX for *ctrsen*
 DOUBLE COMPLEX for *ztrsen*.

Arrays:
t (*ldt*,*) The n by n T .
 The second dimension of *t* must be at least $\max(1, n)$.

q (*ldq*,*)
 If *compq* = 'V', then *q* must contain Q of Schur vectors.
 If *compq* = 'N', then *q* is not referenced.
 The second dimension of *q* must be at least $\max(1, n)$ if *compq* = 'V' and at least 1 if *compq* = 'N'.

work (*lwork*) is a workspace array.
 For complex flavors: the array *work* is not referenced if *job* = 'N'.
 The actual amount of workspace required cannot exceed $n^2/4$ if *job* = 'E' or $n^2/2$ if *job* = 'V' or 'B'.

ldt INTEGER. The first dimension of *t*; at least $\max(1, n)$.

ldq INTEGER. The first dimension of *q*;
 If *compq* = 'N', then *ldq* ≥ 1 .
 If *compq* = 'V', then *ldq* $\geq \max(1, n)$.

lwork INTEGER. The dimension of the array *work*.
 If *job* = 'V' or 'B', *lwork* $\geq \max(1, 2m(n-m))$.
 If *job* = 'E', then *lwork* $\geq \max(1, m(n-m))$.
 If *job* = 'N', then *lwork* ≥ 1 for complex flavors and *lwork* $\geq \max(1, n)$ for real flavors.

iwork INTEGER.
iwork (*liwork*) is a workspace array.
 The array *iwork* is not referenced if *job* = 'N' or 'E'.
 The actual amount of workspace required cannot exceed $n^2/2$ if *job* = 'V' or 'B'.

liwork **INTEGER.**
 The dimension of the array *iwork*.
 If *job* = 'V' or 'B', *liwork* ≥ max(1, 2*m*(*n*-*m*)).
 If *job* = 'E' or 'E', *liwork* ≥ 1.

Output Parameters

t Overwritten by the updated matrix *R*.

q If *compq* = 'V', *q* contains the updated matrix of Schur vectors; the first *m* columns of the *Q* form an orthogonal basis for the specified invariant subspace.

w **COMPLEX** for *ctrsen*
DOUBLE COMPLEX for *ztrsen*.
 Array, **DIMENSION** at least max(1, *n*).
 The recorded eigenvalues of *R*. The eigenvalues are stored in the same order as on the diagonal of *R*.

wr, wi **REAL** for *strsen*
DOUBLE PRECISION for *dtrsen*
 Arrays, **DIMENSION** at least max(1, *n*).
 Contain the real and imaginary parts, respectively, of the reordered eigenvalues of *R*. The eigenvalues are stored in the same order as on the diagonal of *R*. Note that if a complex eigenvalue is sufficiently ill-conditioned, then its value may differ significantly from its value before reordering.

m **INTEGER.**
For complex flavors: the number of the specified invariant subspaces, which is the same as the number of selected eigenvalues (see *select*).
For real flavors: the dimension of the specified invariant subspace. The value of *m* is obtained by counting 1 for each selected real eigenvalue and 2 for each selected complex conjugate pair of eigenvalues (see *select*).
 Constraint: 0 ≤ *m* ≤ *n*.

<i>s</i>	<p>REAL for single-precision flavors DOUBLE PRECISION for double-precision flavors. If <i>job</i> = 'E' or 'B', <i>s</i> is a lower bound on the reciprocal condition number of the average of the selected cluster of eigenvalues. If <i>m</i> = 0 or <i>n</i>, then <i>s</i> = 1. For real flavors: if <i>info</i> = 1, then <i>s</i> is set to zero. <i>s</i> is not referenced if <i>job</i> = 'N' or 'V'.</p>
<i>sep</i>	<p>REAL for single-precision flavors DOUBLE PRECISION for double-precision flavors. If <i>job</i> = 'V' or 'B', <i>sep</i> is the estimated reciprocal condition number of the specified invariant subspace. If <i>m</i> = 0 or <i>n</i>, then <i>sep</i> = $\ T\$. For real flavors: if <i>info</i> = 1, then <i>sep</i> is set to zero. <i>sep</i> is not referenced if <i>job</i> = 'N' or 'E'.</p>
<i>work(1)</i>	<p>On exit, if <i>info</i> = 0, then <i>work(1)</i> returns the required minimal size of <i>lwork</i>.</p>
<i>iwork(1)</i>	<p>On exit, if <i>info</i> = 0, then <i>iwork(1)</i> returns the required minimal size of <i>liwork</i>.</p>
<i>info</i>	<p>INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = -<i>i</i>, the <i>i</i>th parameter had an illegal value.</p>

Application Notes

The computed matrix *R* is exactly similar to a matrix $T + E$, where $\|E\|_2 = O(\epsilon)\|T\|_2$, and ϵ is the machine precision.

The computed *s* cannot underestimate the true reciprocal condition number by more than a factor of $(\min(m, n-m))^{1/2}$; *sep* may differ from the true value by $(m^*n-m^2)^{1/2}$. The angle between the computed invariant subspace and the true subspace is $O(\epsilon)\|A\|_2/sep$.

Note that if a 2 by 2 diagonal block is involved in the re-ordering, its off-diagonal elements are in general changed; the diagonal elements and the eigenvalues of the block are unchanged unless the block is sufficiently ill-conditioned, in which case they may be noticeably altered. It is possible for a 2 by 2 block to break into two 1 by 1 blocks, that is, for a pair of complex eigenvalues to become purely real. The values of eigenvalues however are never changed by the re-ordering.

?trsyl

Solves Sylvester's equation for real quasi-triangular or complex triangular matrices.

```
call strsyl ( trana,tranb, isgn,m,n,a, lda,b,ldb,c, ldc, scale, info )
call dtrsyl ( trana,tranb, isgn,m,n,a, lda,b,ldb,c, ldc, scale, info )
call ctrsyl ( trana,tranb, isgn,m,n,a, lda,b,ldb,c, ldc, scale, info )
call ztrsyl ( trana,tranb, isgn,m,n,a, lda,b,ldb,c, ldc, scale, info )
```

Discussion

This routine solves the Sylvester matrix equation $\text{op}(A)X \pm X\text{op}(B) = \alpha C$, where $\text{op}(A) = A$ or A^H , and the matrices A and B are upper triangular (or, for real flavors, upper quasi-triangular in canonical Schur form); $\alpha \leq 1$ is a scale factor determined by the routine to avoid overflow in X ; A is m by m , B is n by n , and C and X are both m by n . The matrix X is obtained by a straightforward process of back substitution.

The equation has a unique solution if and only if $\alpha_i \pm \beta_i \neq 0$, where $\{\alpha_i\}$ and $\{\beta_i\}$ are the eigenvalues of A and B , respectively, and the sign (+ or -) is the same as that used in the equation to be solved.

Input Parameters

trana CHARACTER*1. Must be 'N' or 'T' or 'C'.
 If *trana* = 'N', then $\text{op}(A) = A$.
 If *trana* = 'T', then $\text{op}(A) = A^T$ (real flavors only).
 If *trana* = 'C' then $\text{op}(A) = A^H$.

tranb CHARACTER*1. Must be 'N' or 'T' or 'C'.
 If *tranb* = 'N', then $\text{op}(B) = B$.
 If *tranb* = 'T', then $\text{op}(B) = B^T$ (real flavors only).
 If *tranb* = 'C', then $\text{op}(B) = B^H$.

isgn INTEGER. Indicates the form of the Sylvester equation.
 If *isgn* = +1, $\text{op}(A)X + X\text{op}(B) = \alpha C$.
 If *isgn* = -1, $\text{op}(A)X - X\text{op}(B) = \alpha C$.

<i>m</i>	INTEGER. The order of <i>A</i> , and the number of rows in <i>X</i> and <i>C</i> ($m \geq 0$).
<i>n</i>	INTEGER. The order of <i>B</i> , and the number of columns in <i>X</i> and <i>C</i> ($n \geq 0$).
<i>a</i> , <i>b</i> , <i>c</i>	REAL for strsyl DOUBLE PRECISION for dtrsyl COMPLEX for ctrsyl DOUBLE COMPLEX for ztrsyl . Arrays: <i>a</i> (<i>lda</i> ,*) contains the matrix <i>A</i> . The second dimension of <i>a</i> must be at least $\max(1, m)$. <i>b</i> (<i>ldb</i> ,*) contains the matrix <i>B</i> . The second dimension of <i>b</i> must be at least $\max(1, n)$. <i>c</i> (<i>ldc</i> ,*) contains the matrix <i>C</i> . The second dimension of <i>c</i> must be at least $\max(1, n)$.
<i>lda</i>	INTEGER. The first dimension of <i>a</i> ; at least $\max(1, m)$.
<i>ldb</i>	INTEGER. The first dimension of <i>b</i> ; at least $\max(1, n)$.
<i>ldc</i>	INTEGER. The first dimension of <i>c</i> ; at least $\max(1, n)$.

Output Parameters

<i>c</i>	Overwritten by the solution matrix <i>X</i> .
<i>scale</i>	REAL for single-precision flavors DOUBLE PRECISION for double-precision flavors. The value of the scale factor α .
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = 1, <i>A</i> and <i>B</i> have common or close eigenvalues perturbed values were used to solve the equation.

Application Notes

Let *X* be the exact, *Y* the corresponding computed solution, and *R* the residual matrix: $R = C - (AY \pm YB)$. Then the residual is always small:

$$\|R\|_F = O(\epsilon) (\|A\|_F + \|B\|_F) \|Y\|_F.$$

However, Y is not necessarily the exact solution of a slightly perturbed equation; in other words, the solution is not backwards stable.

For the forward error, the following bound holds:

$$\|Y - X\|_F \leq \|R\|_F / \text{sep}(A, B)$$

but this may be a considerable overestimate. See [\[Golub96\]](#) for a definition of $\text{sep}(A, B)$.

The approximate number of floating-point operations for real flavors is $m^* n^* (m + n)$. For complex flavors it is 4 times greater.

Generalized Nonsymmetric Eigenvalue Problems

This section describes LAPACK routines for solving generalized nonsymmetric eigenvalue problems, reordering the generalized Schur factorization of a pair of matrices, as well as performing a number of related computational tasks.

A *generalized nonsymmetric eigenvalue problem* is as follows: given a pair of nonsymmetric (or non-Hermitian) n -by- n matrices A and B , find the *generalized eigenvalues* λ and the corresponding *generalized eigenvectors* x and y that satisfy the equations

$$Ax = \lambda Bx \quad (\text{right generalized eigenvectors } x)$$

and

$$y^H A = \lambda y^H B \quad (\text{left generalized eigenvectors } y).$$

[Table 5-6](#) lists LAPACK routines used to solve the generalized nonsymmetric eigenvalue problems and the generalized Sylvester equation.

Table 5-6 Computational Routines for Solving Generalized Nonsymmetric Eigenvalue Problems

Routine name	Operation performed
?gghrd	Reduces a pair of matrices to generalized upper Hessenberg form using orthogonal/unitary transformations.
?ggbal	Balances a pair of general real or complex matrices.
?ggbak	Forms the right or left eigenvectors of a generalized eigenvalue problem.
?hgeqz	Implements the QZ method for finding the generalized eigenvalues of the matrix pair (H,T).
?tgevc	Computes some or all of the right and/or left generalized eigenvectors of a pair of upper triangular matrices
?tgexc	Reorders the generalized Schur decomposition of a pair of matrices (A,B) so that one diagonal block of (A,B) moves to another row index.
?tgsen	Reorders the generalized Schur decomposition of a pair of matrices (A,B) so that a selected cluster of eigenvalues appears in the leading diagonal blocks of (A,B).
?tgsyl	Solves the generalized Sylvester equation.
?tgsna	Estimates reciprocal condition numbers for specified eigenvalues and/or eigenvectors of a pair of matrices in generalized real Schur canonical form.

?gghrd

Reduces a pair of matrices to generalized upper Hessenberg form using orthogonal/unitary transformations.

```
call sgghrd ( compq, compz, n, ilo, ihi, a, lda, b, ldb, q, ldq,
             z, ldz, info )
call dgghrd ( compq, compz, n, ilo, ihi, a, lda, b, ldb, q, ldq,
             z, ldz, info )
call cgghrd ( compq, compz, n, ilo, ihi, a, lda, b, ldb, q, ldq,
             z, ldz, info )
call zgghrd ( compq, compz, n, ilo, ihi, a, lda, b, ldb, q, ldq,
             z, ldz, info )
```

Discussion

This routine reduces a pair of real/complex matrices (A,B) to generalized upper Hessenberg form using orthogonal/unitary transformations, where A is a general matrix and B is upper triangular. The form of the generalized eigenvalue problem is $Ax = \lambda Bx$, and B is typically made upper triangular by computing its QR factorization and moving the orthogonal matrix Q to the left side of the equation.

This routine simultaneously reduces A to a Hessenberg matrix H:

$$Q^H A Z = H$$

and transforms B to another upper triangular matrix T:

$$Q^H B Z = T$$

in order to reduce the problem to its standard form $Hy = \lambda Ty$ where $y = Z^H x$.

The orthogonal/unitary matrices Q and Z are determined as products of Givens rotations. They may either be formed explicitly, or they may be postmultiplied into input matrices Q_1 and Z_1 , so that

$$Q_1 A Z_1^H = (Q_1 Q) H (Z_1 Z)^H$$

$$Q_1 B Z_1^H = (Q_1 Q) T (Z_1 Z)^H$$

If Q_1 is the orthogonal matrix from the QR factorization of B in the original equation $Ax = \lambda Bx$, then `?gghrd` reduces the original problem to generalized Hessenberg form.

Input Parameters

`compq` CHARACTER*1. Must be 'N', 'I', or 'V'.
 If `compq` = 'N', matrix Q is not computed.
 If `compq` = 'I', Q is initialized to the unit matrix, and the orthogonal/unitary matrix Q is returned;
 If `compq` = 'V', Q must contain an orthogonal/unitary matrix Q_1 on entry, and the product Q_1Q is returned.

`compz` CHARACTER*1. Must be 'N', 'I', or 'V'.
 If `compz` = 'N', matrix Z is not computed.
 If `compz` = 'I', Z is initialized to the unit matrix, and the orthogonal/unitary matrix Z is returned;
 If `compz` = 'V', Z must contain an orthogonal/unitary matrix Z_1 on entry, and the product Z_1Z is returned.

`n` INTEGER. The order of the matrices A and B ($n \geq 0$).

`ilo, ihi` INTEGER. `ilo` and `ihi` mark the rows and columns of A which are to be reduced. It is assumed that A is already upper triangular in rows and columns $1:\text{ilo}-1$ and $\text{ihi}+1:n$. Values of `ilo` and `ihi` are normally set by a previous call to `?ggbal`; otherwise they should be set to 1 and n respectively. Constraint:
 If $n > 0$, then $1 \leq \text{ilo} \leq \text{ihi} \leq n$;
 if $n = 0$, then `ilo` = 1 and `ihi` = 0.

`a, b, q, z` REAL for `sgghrd`
 DOUBLE PRECISION for `dgghrd`
 COMPLEX for `cgghrd`
 DOUBLE COMPLEX for `zgghrd`.
 Arrays:
`a(lda,*)` contains the n -by- n general matrix A .
 The second dimension of `a` must be at least $\max(1, n)$.
`b(ldb,*)` contains the n -by- n upper triangular matrix B .
 The second dimension of `b` must be at least $\max(1, n)$.

$q(ldq, *)$

If $compq = 'N'$, then q is not referenced.

If $compq = 'I'$, then, on entry, q need not be set.

If $compq = 'V'$, then q must contain the orthogonal/unitary matrix Q_1 , typically from the QR factorization of B .

The second dimension of q must be at least $\max(1, n)$.

$z(ldz, *)$

If $compq = 'N'$, then z is not referenced.

If $compq = 'I'$, then, on entry, z need not be set.

If $compq = 'V'$, then z must contain the orthogonal/unitary matrix Z_1 .

The second dimension of z must be at least $\max(1, n)$.

lda INTEGER. The first dimension of a ; at least $\max(1, n)$.

ldb INTEGER. The first dimension of b ; at least $\max(1, n)$.

ldq INTEGER. The first dimension of q ;
 If $compq = 'N'$, then $ldq \geq 1$.
 If $compq = 'I'$ or $'V'$, then $ldq \geq \max(1, n)$.

ldz INTEGER. The first dimension of z ;
 If $compq = 'N'$, then $ldz \geq 1$.
 If $compq = 'I'$ or $'V'$, then $ldz \geq \max(1, n)$.

Output Parameters

a On exit, the upper triangle and the first subdiagonal of A are overwritten with the upper Hessenberg matrix H , and the rest is set to zero.

b On exit, overwritten by the upper triangular matrix $T = Q^H B Z$. The elements below the diagonal are set to zero.

q If $compq = 'I'$, then q contains the orthogonal/unitary matrix Q , where Q^H is the product of the Givens transformations which are applied to A and B on the left;
 If $compq = 'V'$, then q is overwritten by the product $Q_1 Q$.

z If *compq* = 'I', then *z* contains the orthogonal/unitary matrix *Z*, which is the product of the Givens transformations which are applied to *A* and *B* on the right;
If *compq* = 'V', then *z* is overwritten by the product $Z_1 Z$.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.

?ggbal

Balances a pair of general real or complex matrices.

```
call sggbal ( job, n, a, lda, b, ldb, ilo, ihi, lscale, rscale,
             work, info )
call dggbal ( job, n, a, lda, b, ldb, ilo, ihi, lscale, rscale,
             work, info )
call cggbal ( job, n, a, lda, b, ldb, ilo, ihi, lscale, rscale,
             work, info )
call zggbal ( job, n, a, lda, b, ldb, ilo, ihi, lscale, rscale,
             work, info )
```

Discussion

This routine balances a pair of general real/complex matrices (A, B). This involves, first, permuting A and B by similarity transformations to isolate eigenvalues in the first 1 to $ilo-1$ and last $ihi+1$ to n elements on the diagonal; and second, applying a diagonal similarity transformation to rows and columns ilo to ihi to make the rows and columns as close in norm as possible. Both steps are optional.

Balancing may reduce the 1-norm of the matrices, and improve the accuracy of the computed eigenvalues and/or eigenvectors in the generalized eigenvalue problem $Ax = \lambda Bx$.

Input Parameters

job CHARACTER*1. Specifies the operations to be performed on A and B . Must be 'N' or 'P' or 'S' or 'B'.
 If $job = 'N'$, then no operations are done; simply set $ilo=1$, $ihi=n$, $lscale(i)=1.0$ and $rscale(i)=1.0$ for $i = 1, \dots, n$.
 If $job = 'P'$, then permute only.
 If $job = 'S'$, then scale only.
 If $job = 'B'$, then both permute and scale.

n INTEGER. The order of the matrices A and B ($n \geq 0$).

a, b REAL for `sggbal`
 DOUBLE PRECISION for `dggbal`
 COMPLEX for `cggbal`
 DOUBLE COMPLEX for `zggbal`.
 Arrays:
a(lda,)* contains the matrix *A*.
 The second dimension of *a* must be at least $\max(1, n)$.
b(ldb,)* contains the matrix *B*.
 The second dimension of *b* must be at least $\max(1, n)$.

lda INTEGER. The first dimension of *a*; at least $\max(1, n)$.

ldb INTEGER. The first dimension of *b*; at least $\max(1, n)$.

work REAL for single precision flavors
 DOUBLE PRECISION for double precision flavors.
 Workspace array, DIMENSION at least $\max(1, 6n)$.

Output Parameters

a, b Overwritten by the balanced matrices *A* and *B*, respectively. If *job*='N', *a* and *b* are not referenced.

ilo, ihi INTEGER. *ilo* and *ihi* are set to integers such that on exit $a(i, j)=0$ and $b(i, j)=0$ if $i > j$ and $j=1, \dots, ilo-1$ or $i=ihi+1, \dots, n$.
 If *job*='N' or 'S', then *ilo* = 1 and *ihi* = *n*.

lscale, rscale REAL for single precision flavors
 DOUBLE PRECISION for double precision flavors.
 Arrays, DIMENSION at least $\max(1, n)$.
lscale contains details of the permutations and scaling factors applied to the left side of *A* and *B*.
 If P_j is the index of the row interchanged with row *j*, and D_j is the scaling factor applied to row *j*, then

$$lscale(j) = P_j, \text{ for } j = 1, \dots, ilo-1$$

$$= D_j, \text{ for } j = ilo, \dots, ihi$$

$$= P_j, \text{ for } j = ihi+1, \dots, n.$$
rscale contains details of the permutations and scaling factors applied to the right side of *A* and *B*.

If P_j is the index of the column interchanged with column j , and D_j is the scaling factor applied to column j , then

$$\begin{aligned} rscale(j) &= P_j, \text{ for } j = 1, \dots, ilo-1 \\ &= D_j, \text{ for } j = ilo, \dots, ihi \\ &= P_j, \text{ for } j = ihi+1, \dots, n \end{aligned}$$

The order in which the interchanges are made is n to $ihi+1$, then 1 to $ilo-1$.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = $-i$, the i th parameter had an illegal value.

?ggbak

Forms the right or left eigenvectors of a generalized eigenvalue problem.

```
call sggbak(job, side, n, ilo, ihi, lscale, rscale, m, v, ldv, info)
call dggbak(job, side, n, ilo, ihi, lscale, rscale, m, v, ldv, info)
call cggbak(job, side, n, ilo, ihi, lscale, rscale, m, v, ldv, info)
call zggbak(job, side, n, ilo, ihi, lscale, rscale, m, v, ldv, info)
```

Discussion

This routine forms the right or left eigenvectors of a real/complex generalized eigenvalue problem

$$Ax = \lambda Bx$$

by backward transformation on the computed eigenvectors of the balanced pair of matrices output by [?ggbal](#).

Input Parameters

job CHARACTER*1. Specifies the type of backward transformation required. Must be 'N', 'P', 'S', or 'B'.
 If *job* = 'N', then no operations are done; return.
 If *job* = 'P', then do backward transformation for permutation only.
 If *job* = 'S', then do backward transformation for scaling only.
 If *job* = 'B', then do backward transformation for both permutation and scaling.
 This argument must be the same as the argument *job* supplied to [?ggbal](#).

side CHARACTER*1. Must be 'L' or 'R'.
 If *side* = 'L', then *v* contains left eigenvectors .
 If *side* = 'R', then *v* contains right eigenvectors .

n INTEGER. The number of rows of the matrix *V* ($n \geq 0$).

ilo, ihi **INTEGER**. The integers *ilo* and *ihi* determined by `?gebal`. Constraint:
 If $n > 0$, then $1 \leq ilo \leq ihi \leq n$;
 if $n = 0$, then $ilo = 1$ and $ihi = 0$.

lscale, rscale **REAL** for single precision flavors
DOUBLE PRECISION for double precision flavors.
 Arrays, **DIMENSION** at least $\max(1, n)$.
 The array *lscale* contains details of the permutations and/or scaling factors applied to the left side of *A* and *B*, as returned by `?ggbal`.
 The array *rscale* contains details of the permutations and/or scaling factors applied to the right side of *A* and *B*, as returned by `?ggbal`.

m **INTEGER**. The number of columns of the matrix *V* ($m \geq 0$).

v **REAL** for `sggbak`
DOUBLE PRECISION for `dggbak`
COMPLEX for `cggbak`
DOUBLE COMPLEX for `zggbak`.
 Array `v(ldv, *)`. Contains the matrix of right or left eigenvectors to be transformed, as returned by `?tgevc`.
 The second dimension of *v* must be at least $\max(1, m)$.

ldv **INTEGER**. The first dimension of *v*; at least $\max(1, n)$.

Output Parameters

v Overwritten by the transformed eigenvectors

info **INTEGER**.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

?hgeqz

Implements the *QZ* method for finding the generalized eigenvalues of the matrix pair (H,T) .

```
call shgeqz(job, compq, compz, n, ilo, ihi, h, ldh, t, ldt, alphas,
           alpha_i, beta, q, ldq, z, ldz, work, lwork, info )
call dhgeqz(job, compq, compz, n, ilo, ihi, h, ldh, t, ldt, alphas,
           alpha_i, beta, q, ldq, z, ldz, work, lwork, info )
call chgeqz(job, compq, compz, n, ilo, ihi, h, ldh, t, ldt, alpha,
           beta, q, ldq, z, ldz, work, lwork, rwork, info )
call zhgeqz(job, compq, compz, n, ilo, ihi, h, ldh, t, ldt, alpha,
           beta, q, ldq, z, ldz, work, lwork, rwork, info )
```

Discussion

This routine computes the eigenvalues of a real/complex matrix pair (H,T) , where H is an upper Hessenberg matrix and T is upper triangular, using the double-shift version (for real flavors) or single-shift version (for complex flavors) of the *QZ* method.

Matrix pairs of this type are produced by the reduction to generalized upper Hessenberg form of a real/complex matrix pair (A,B) :

$$A = Q_1 H Z_1^H, \quad B = Q_1 T Z_1^H,$$

as computed by `?gghrd`.

For real flavors:

If `job='S'`, then the Hessenberg-triangular pair (H,T) is also reduced to generalized Schur form,

$$H = Q S Z^T, \quad T = Q P Z^T,$$

where Q and Z are orthogonal matrices, P is an upper triangular matrix, and S is a quasi-triangular matrix with 1-by-1 and 2-by-2 diagonal blocks.

The 1-by-1 blocks correspond to real eigenvalues of the matrix pair (H,T) and the 2-by-2 blocks correspond to complex conjugate pairs of eigenvalues.

Additionally, the 2-by-2 upper triangular diagonal blocks of P

corresponding to 2-by-2 blocks of S are reduced to positive diagonal form, that is, if $S(j+1,j)$ is non-zero, then $P(j+1,j) = P(j,j+1) = 0$, $P(j,j) > 0$, and $P(j+1,j+1) > 0$.

For complex flavors:

If `job = 'S'`, then the Hessenberg-triangular pair (H,T) is also reduced to generalized Schur form,

$$H = Q S Z^H, \quad T = Q P Z^H,$$

where Q and Z are unitary matrices, and S and P are upper triangular.

For all function flavors:

Optionally, the orthogonal/unitary matrix Q from the generalized Schur factorization may be postmultiplied into an input matrix Q_I , and the orthogonal/unitary matrix Z may be postmultiplied into an input matrix Z_I . If Q_I and Z_I are the orthogonal/unitary matrices from `?gghrd` that reduced the matrix pair (A,B) to generalized upper Hessenberg form, then the output matrices $Q_I Q$ and $Z_I Z$ are the orthogonal/unitary factors from the generalized Schur factorization of (A,B) :

$$A = (Q_I Q) S (Z_I Z)^H, \quad B = (Q_I Q) P (Z_I Z)^H.$$

To avoid overflow, eigenvalues of the matrix pair (H,T) (equivalently, of (A,B)) are computed as a pair of values (α, β) . For `chgeqz/zhgeqz`, α and β are complex, and for `shgeqz/dhgeqz`, α is complex and β real. If β is nonzero, $\lambda = \alpha / \beta$ is an eigenvalue of the generalized nonsymmetric eigenvalue problem (GNEP)

$$Ax = \lambda Bx$$

and if α is nonzero, $\mu = \beta / \alpha$ is an eigenvalue of the alternate form of the GNEP

$$\mu Ay = By.$$

Real eigenvalues (for real flavors) or the values of α and β for the i -th eigenvalue (for complex flavors) can be read directly from the generalized Schur form:

$$\alpha = S(i,i), \quad \beta = P(i,i).$$

Input Parameters

- job* CHARACTER*1. Specifies the operations to be performed. Must be 'E' or 'S' .
 If *job* = 'E', then compute eigenvalues only;
 If *job* = 'S', then compute eigenvalues and the Schur form.
- compq* CHARACTER*1. Must be 'N', 'I', or 'V' .
 If *compq* = 'N', left Schur vectors (*q*) are not computed;
 If *compq* = 'I', *q* is initialized to the unit matrix and the matrix of left Schur vectors of (*H*,*T*) is returned;
 If *compq* = 'V', *q* must contain an orthogonal/unitary matrix Q_I on entry and the product $Q_I Q$ is returned.
- compz* CHARACTER*1. Must be 'N', 'I', or 'V' .
 If *compz* = 'N', left Schur vectors (*q*) are not computed;
 If *compz* = 'I', *z* is initialized to the unit matrix and the matrix of right Schur vectors of (*H*,*T*) is returned;
 If *compz* = 'V', *z* must contain an orthogonal/unitary matrix Z_I on entry and the product $Z_I Z$ is returned.
- n* INTEGER. The order of the matrices *H*, *T*, *Q*, and *Z* ($n \geq 0$).
- ilo, ihi* INTEGER. *ilo* and *ihi* mark the rows and columns of *H* which are in Hessenberg form. It is assumed that *H* is already upper triangular in rows and columns 1:*ilo*-1 and *ihi*+1:*n*. Constraint:
 If $n > 0$, then $1 \leq ilo \leq ihi \leq n$;
 if $n = 0$, then *ilo* = 1 and *ihi* = 0.
- h, t, q, z, work* REAL for shgeqz
 DOUBLE PRECISION for dhgeqz
 COMPLEX for chgeqz
 DOUBLE COMPLEX for zhgeqz.
 Arrays:
 On entry, *h*(*ldh*,*) contains the *n*-by-*n* upper Hessenberg matrix *H*.
 The second dimension of *h* must be at least $\max(1, n)$.

On entry, $t(ldt, *)$ contains the n -by- n upper triangular matrix T .

The second dimension of t must be at least $\max(1, n)$.

$q(ldq, *)$:

On entry, if $compq = 'V'$, this array contains the orthogonal/unitary matrix Q_I used in the reduction of (A, B) to generalized Hessenberg form.

If $compq = 'N'$, then q is not referenced.

The second dimension of q must be at least $\max(1, n)$.

$z(ldz, *)$:

On entry, if $compz = 'V'$, this array contains the orthogonal/unitary matrix Z_I used in the reduction of (A, B) to generalized Hessenberg form.

If $compz = 'N'$, then z is not referenced.

The second dimension of z must be at least $\max(1, n)$.

$work(lwork)$ is a workspace array.

ldh INTEGER. The first dimension of h ; at least $\max(1, n)$.

ldt INTEGER. The first dimension of t ; at least $\max(1, n)$.

ldq INTEGER. The first dimension of q ;
If $compq = 'N'$, then $ldq \geq 1$.
If $compq = 'I'$ or $'V'$, then $ldq \geq \max(1, n)$.

ldz INTEGER. The first dimension of z ;
If $compz = 'N'$, then $ldz \geq 1$.
If $compz = 'I'$ or $'V'$, then $ldz \geq \max(1, n)$.

$lwork$ INTEGER. The dimension of the array $work$.
 $lwork \geq \max(1, n)$.

$rwork$ REAL for $chgeqz$
DOUBLE PRECISION for $zhgeqz$.
Workspace array, DIMENSION at least $\max(1, n)$. Used in complex flavors only.

Output Parameters

- h* For real flavors:
 If *job* = 'S', then, on exit, *h* contains the upper quasi-triangular matrix *S* from the generalized Schur factorization; 2-by-2 diagonal blocks (corresponding to complex conjugate pairs of eigenvalues) are returned in standard form, with $h(i,i) = h(i+1, i+1)$ and $h(i+1, i) * h(i, i+1) < 0$.
 If *job* = 'E', then on exit the diagonal blocks of *h* match those of *S*, but the rest of *h* is unspecified.
- For complex flavors:
 If *job* = 'S', then, on exit, *h* contains the upper triangular matrix *S* from the generalized Schur factorization.
 If *job* = 'E', then on exit the diagonal of *h* matches that of *S*, but the rest of *h* is unspecified.
- t* If *job* = 'S', then, on exit, *t* contains the upper triangular matrix *P* from the generalized Schur factorization.
- For real flavors:
 2-by-2 diagonal blocks of *P* corresponding to 2-by-2 blocks of *S* are reduced to positive diagonal form, that is, if $h(j+1,j)$ is non-zero, then $t(j+1,j)=t(j,j+1)=0$ and $t(j,j)$ and $t(j+1,j+1)$ will be positive.
 If *job* = 'E', then on exit the diagonal blocks of *t* match those of *P*, but the rest of *t* is unspecified.
- For complex flavors:
 If *job* = 'E', then on exit the diagonal of *t* matches that of *P*, but the rest of *t* is unspecified.
- alphan, alphai* REAL for shgeqz;
 DOUBLE PRECISION for dhgeqz.
 Arrays, DIMENSION at least max(1,n).
 The real and imaginary parts, respectively, of each scalar *alpha* defining an eigenvalue of GNEP.

	<p>If $\mathit{alpha}(j)$ is zero, then the j-th eigenvalue is real; if positive, then the j-th and $(j+1)$-th eigenvalues are a complex conjugate pair, with $\mathit{alpha}(j+1) = -\mathit{alpha}(j)$.</p>
alpha	<p>COMPLEX for chgeqz; DOUBLE COMPLEX for zhgeqz. Array, DIMENSION at least $\max(1,n)$. The complex scalars alpha that define the eigenvalues of GNEP. $\mathit{alpha}(i) = S(i,i)$ in the generalized Schur factorization.</p>
beta	<p>REAL for shgeqz DOUBLE PRECISION for dhgeqz COMPLEX for chgeqz DOUBLE COMPLEX for zhgeqz. Array, DIMENSION at least $\max(1,n)$. For real flavors: The scalars beta that define the eigenvalues of GNEP. Together, the quantities $\mathit{alpha} = (\mathit{alphar}(j), \mathit{alpha}(j))$ and $\mathit{beta} = \mathit{beta}(j)$ represent the j-th eigenvalue of the matrix pair (A,B), in one of the forms $\lambda = \mathit{alpha}/\mathit{beta}$ or $\mu = \mathit{beta}/\mathit{alpha}$. Since either λ or μ may overflow, they should not, in general, be computed. For complex flavors: The real non-negative scalars beta that define the eigenvalues of GNEP. $\mathit{beta}(i) = P(i,i)$ in the generalized Schur factorization. Together, the quantities $\mathit{alpha} = \mathit{alpha}(j)$ and $\mathit{beta} = \mathit{beta}(j)$ represent the j-th eigenvalue of the matrix pair (A,B), in one of the forms $\lambda = \mathit{alpha}/\mathit{beta}$ or $\mu = \mathit{beta}/\mathit{alpha}$. Since either λ or μ may overflow, they should not, in general, be computed.</p>
q	<p>On exit, if $\mathit{compq} = 'I'$, q is overwritten by the orthogonal/unitary matrix of left Schur vectors of the pair (H,T), and if $\mathit{compq} = 'V'$, q is overwritten by the orthogonal/unitary matrix of left Schur vectors of (A,B).</p>

<i>z</i>	On exit, if <i>compz</i> = 'I', <i>z</i> is overwritten by the orthogonal/unitary matrix of right Schur vectors of the pair (<i>H</i> , <i>T</i>), and if <i>compz</i> = 'V', <i>z</i> is overwritten by the orthogonal/unitary matrix of right Schur vectors of (<i>A</i> , <i>B</i>).
<i>work(1)</i>	If <i>info</i> ≥ 0, on exit, <i>work(1)</i> contains the minimum value of <i>lwork</i> required for optimum performance. Use this <i>lwork</i> for subsequent runs.
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = 1,..., <i>n</i> , the <i>QZ</i> iteration did not converge. (<i>H</i> , <i>T</i>) is not in Schur form, but <i>alphar</i> (<i>i</i>), <i>alphai</i> (<i>i</i>) (for real flavors), <i>alpha</i> (<i>i</i>) (for complex flavors), and <i>beta</i> (<i>i</i>), <i>i</i> = <i>info</i> +1,..., <i>n</i> should be correct. If <i>info</i> = <i>n</i> +1,...,2 <i>n</i> , the shift calculation failed. (<i>H</i> , <i>T</i>) is not in Schur form, but <i>alphar</i> (<i>i</i>), <i>alphai</i> (<i>i</i>) (for real flavors), <i>alpha</i> (<i>i</i>) (for complex flavors), and <i>beta</i> (<i>i</i>), <i>i</i> = <i>info</i> - <i>n</i> +1,..., <i>n</i> should be correct.

?tgevc

Computes some or all of the right and/or left generalized eigenvectors of a pair of upper triangular matrices.

```
call stgevc ( side, howmny, select, n, s, lds, p, ldp, vl, ldvl, vr,
             ldvr, mm, m, work, info )
call dtgevc ( side, howmny, select, n, s, lds, p, ldp, vl, ldvl, vr,
             ldvr, mm, m, work, info )
call ctgevc ( side, howmny, select, n, s, lds, p, ldp, vl, ldvl, vr,
             ldvr, mm, m, work, rwork, info )
call ztgevc ( side, howmny, select, n, s, lds, p, ldp, vl, ldvl, vr,
             ldvr, mm, m, work, rwork, info )
```

Discussion

This routine computes some or all of the right and/or left eigenvectors of a pair of real/complex matrices (S,P) , where S is quasi-triangular (for real flavors) or upper triangular (for complex flavors) and P is upper triangular.

Matrix pairs of this type are produced by the generalized Schur factorization of a real/complex matrix pair (A,B) :

$$A = Q S Z^H, \quad B = Q P Z^H$$

as computed by ?gghrd plus ?hgeqz.

The right eigenvector x and the left eigenvector y of (S,P) corresponding to an eigenvalue w are defined by:

$$S x = w P x, \quad y^H S = w y^H P$$

The eigenvalues are not input to this routine, but are computed directly from the diagonal blocks or diagonal elements of S and P .

This routine returns the matrices X and/or Y of right and left eigenvectors of (S,P) , or the products ZX and/or QY , where Z and Q are input matrices. If Q and Z are the orthogonal/unitary factors from the generalized Schur factorization of a matrix pair (A,B) , then ZX and QY are the matrices of right and left eigenvectors of (A,B) .

Input Parameters

side CHARACTER*1. Must be 'R', 'L', or 'B'.
 If *side* = 'R', compute right eigenvectors only.
 If *side* = 'L', compute left eigenvectors only.
 If *side* = 'B', compute both right and left eigenvectors.

howmny CHARACTER*1. Must be 'A', 'B', or 'S'.
 If *howmny* = 'A', compute all right and/or left eigenvectors.
 If *howmny* = 'B', compute all right and/or left eigenvectors, backtransformed by the matrices in *vr* and/or *vl*.
 If *howmny* = 'S', compute selected right and/or left eigenvectors, specified by the logical array *select*.

select LOGICAL.
 Array, DIMENSION at least max (1, *n*).
 If *howmny* = 'S', *select* specifies the eigenvectors to be computed.
 If *howmny* = 'A' or 'B', *select* is not referenced.
 For real flavors:
 If ω_j is a real eigenvalue, the corresponding real eigenvector is computed if *select*(*j*) is .TRUE..
 If ω_j and ω_{j+1} are the real and imaginary parts of a complex eigenvalue, the corresponding complex eigenvector is computed if either *select*(*j*) or *select*(*j*+1) is .TRUE., and on exit *select*(*j*) is set to .TRUE. and *select*(*j*+1) is set to .FALSE..
 For complex flavors:
 The eigenvector corresponding to the *j*-th eigenvalue is computed if *select*(*j*) is .TRUE..

n INTEGER. The order of the matrices *A* and *B* ($n \geq 0$).

s,p,vl,vr,work REAL for *stgevc*
 DOUBLE PRECISION for *dtgevc*
 COMPLEX for *ctgevc*
 DOUBLE COMPLEX for *ztgevc*.
 Arrays:

s(lds,)* contains the matrix *S* from a generalized Schur factorization as computed by ?hgeqz. This matrix is upper quasi-triangular for real flavors, and upper triangular for complex flavors.

The second dimension of *s* must be at least $\max(1, n)$.

p(ldp,)* contains the upper triangular matrix *P* from a generalized Schur factorization as computed by ?hgeqz.

For real flavors, 2-by-2 diagonal blocks of *P* corresponding to 2-by-2 blocks of *S* must be in positive diagonal form.

For complex flavors, *P* must have real diagonal elements.

The second dimension of *p* must be at least $\max(1, n)$.

If *side* = 'L' or 'B' and *howmny* = 'B',

vl(ldvl,)* must contain an *n*-by-*n* matrix *Q* (usually the orthogonal/unitary matrix *Q* of left Schur vectors returned by ?hgeqz). The second dimension of *vl* must be at least $\max(1, mm)$. If *side* = 'R', *vl* is not referenced.

If *side* = 'R' or 'B' and *howmny* = 'B',

vr(ldvr,)* must contain an *n*-by-*n* matrix *Z* (usually the orthogonal/unitary matrix *Z* of right Schur vectors returned by ?hgeqz). The second dimension of *vr* must be at least $\max(1, mm)$. If *side* = 'L', *vr* is not referenced.

work()* is a workspace array.

DIMENSION at least $\max(1, 6*n)$ for real flavors and at least $\max(1, 2*n)$ for complex flavors.

lda **INTEGER.** The first dimension of *a*; at least $\max(1, n)$.

ldb **INTEGER.** The first dimension of *b*; at least $\max(1, n)$.

ldvl **INTEGER.** The first dimension of *vl*;

If *side* = 'L' or 'B', then *ldvl* $\geq \max(1, n)$.

If *side* = 'R', then *ldvl* ≥ 1 .

ldvr **INTEGER**. The first dimension of *vr*;
If *side* = 'R' or 'B', then *ldvr* \geq max(1,*n*).
If *side* = 'L', then *ldvr* \geq 1.

mm **INTEGER**. The number of columns in the arrays *v1*
and/or *vr* (*mm* \geq *m*).

rwork **REAL** for *ctgevc*
DOUBLE PRECISION for *ztgevc*.
Workspace array, **DIMENSION** at least max (1, 2 * *n*).
Used in complex flavors only.

Output Parameters

v1 On exit, if *side* = 'L' or 'B', *v1* contains:
if *howmny* = 'A', the matrix *Y* of left eigenvectors of
(*S*,*P*);
if *howmny* = 'B', the matrix *QY*;
if *howmny* = 'S', the left eigenvectors of (*S*,*P*) specified
by *select*, stored consecutively in the columns of *v1*,
in the same order as their eigenvalues.
For real flavors:
A complex eigenvector corresponding to a complex
eigenvalue is stored in two consecutive columns, the
first holding the real part, and the second the imaginary
part.

vr On exit, if *side* = 'R' or 'B', *vr* contains:
if *howmny* = 'A', the matrix *X* of right eigenvectors of
(*S*,*P*);
if *howmny* = 'B', the matrix *ZX*;
if *howmny* = 'S', the right eigenvectors of (*S*,*P*)
specified by *select*, stored consecutively in the
columns of *vr*, in the same order as their eigenvalues.
For real flavors:
A complex eigenvector corresponding to a complex
eigenvalue is stored in two consecutive columns, the
first holding the real part, and the second the imaginary
part.

m **INTEGER.** The number of columns in the arrays *vl* and/or *vr* actually used to store the eigenvectors. If *howmny* = 'A' or 'B', *m* is set to *n*.
For real flavors:
Each selected real eigenvector occupies one column and each selected complex eigenvector occupies two columns.
For complex flavors:
Each selected eigenvector occupies one column.

info **INTEGER.**
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
For real flavors:
If *info* = *i*>0, the 2-by-2 block (*i:i*+1) does not have a complex eigenvalue.

?tgexc

Reorders the generalized Schur decomposition of a pair of matrices (A,B) so that one diagonal block of (A,B) moves to another row index.

```
call stgexc ( wantq, wantz, n, a, lda, b, ldb, q, ldq, z, ldz,
             ifst, ilst, work, lwork, info )
call dtgexc ( wantq, wantz, n, a, lda, b, ldb, q, ldq, z, ldz,
             ifst, ilst, work, lwork, info )
call ctgexc ( wantq, wantz, n, a, lda, b, ldb, q, ldq, z, ldz,
             ifst, ilst, info )
call ztgexc ( wantq, wantz, n, a, lda, b, ldb, q, ldq, z, ldz,
             ifst, ilst, info )
```

Discussion

This routine reorders the generalized real-Schur/Schur decomposition of a real/complex matrix pair (A,B) using an orthogonal/unitary equivalence transformation

$$(A, B) = Q (A, B) Z^H,$$

so that the diagonal block of (A, B) with row index *ifst* is moved to row *ilst*.

Matrix pair (A, B) must be in generalized real-Schur/Schur canonical form (as returned by [?gges](#)), i.e. A is block upper triangular with 1-by-1 and 2-by-2 diagonal blocks and B is upper triangular.

Optionally, the matrices Q and Z of generalized Schur vectors are updated.

$$Q(\text{in}) * A(\text{in}) * Z(\text{in})' = Q(\text{out}) * A(\text{out}) * Z(\text{out})'$$

$$Q(\text{in}) * B(\text{in}) * Z(\text{in})' = Q(\text{out}) * B(\text{out}) * Z(\text{out})'.$$

Input Parameters

wantq, wantz LOGICAL.

If *wantq* = **.TRUE.**, update the left transformation matrix Q ;

If `wantq = .FALSE.`, do not update Q ;
 If `wantz = .TRUE.`, update the right transformation matrix Z ;
 If `wantz = .FALSE.`, do not update Z .

`n` **INTEGER**. The order of the matrices A and B ($n \geq 0$).

`a, b, q, z` **REAL** for `stgexc`
 DOUBLE PRECISION for `dtgexc`
 COMPLEX for `ctgexc`
 DOUBLE COMPLEX for `ztgexc`.

Arrays:
`a(lda,*)` contains the matrix A .
 The second dimension of `a` must be at least $\max(1, n)$.
`b(ldb,*)` contains the matrix B .
 The second dimension of `b` must be at least $\max(1, n)$.
`q(ldq,*)`
 If `wantq = .FALSE.`, then `q` is not referenced.
 If `wantq = .TRUE.`, then `q` must contain the orthogonal/unitary matrix Q .
 The second dimension of `q` must be at least $\max(1, n)$.
`z(ldz,*)`
 If `wantz = .FALSE.`, then `z` is not referenced.
 If `wantz = .TRUE.`, then `z` must contain the orthogonal/unitary matrix Z .
 The second dimension of `z` must be at least $\max(1, n)$.

`lda` **INTEGER**. The first dimension of `a`; at least $\max(1, n)$.
`ldb` **INTEGER**. The first dimension of `b`; at least $\max(1, n)$.
`ldq` **INTEGER**. The first dimension of `q`;
 If `wantq = .FALSE.`, then `ldq` ≥ 1 .
 If `wantq = .TRUE.`, then `ldq` $\geq \max(1, n)$.
`ldz` **INTEGER**. The first dimension of `z`;
 If `wantz = .FALSE.`, then `ldz` ≥ 1 .
 If `wantz = .TRUE.`, then `ldz` $\geq \max(1, n)$.

ifst, ilst **INTEGER**. Specify the reordering of the diagonal blocks of (A, B) . The block with row index *ifst* is moved to row *ilst*, by a sequence of swapping between adjacent blocks. Constraint: $1 \leq \textit{ifst}, \textit{ilst} \leq n$.

work **REAL** for *stgexc*;
DOUBLE PRECISION for *dtgexc*.
Workspace array, **DIMENSION** (*lwork*). Used in real flavors only.

lwork **INTEGER**. The dimension of *work*; must be at least $4n + 16$.

Output Parameters

a, b Overwritten by the updated matrices A and B .

ifst, ilst Overwritten for real flavors only.
If *ifst* pointed to the second row of a 2 by 2 block on entry, it is changed to point to the first row; *ilst* always points to the first row of the block in its final position (which may differ from its input value by ± 1).

info **INTEGER**.
If *info* = 0, the execution is successful.
If *info* = $-i$, the *i*th parameter had an illegal value.
If *info* = 1, the transformed matrix pair (A, B) would be too far from generalized Schur form; the problem is ill-conditioned. (A, B) may have been partially reordered, and *ilst* points to the first row of the current position of the block being moved.

?tgsen

Reorders the generalized Schur decomposition of a pair of matrices (A,B) so that a selected cluster of eigenvalues appears in the leading diagonal blocks of (A,B).

```
call stgsen ( ijob, wantq, wantz, select, n, a, lda, b, ldb, alphas,
             alphas, beta, q, ldq, z, ldz, m, pl, pr, dif, work,
             lwork, iwork, liwork, info )
call dtgsen ( ijob, wantq, wantz, select, n, a, lda, b, ldb, alphas,
             alphas, beta, q, ldq, z, ldz, m, pl, pr, dif, work,
             lwork, iwork, liwork, info )
call ctgsen ( ijob, wantq, wantz, select, n, a, lda, b, ldb, alpha,
             beta, q, ldq, z, ldz, m, pl, pr, dif, work,
             lwork, iwork, liwork, info )
call ztgsen ( ijob, wantq, wantz, select, n, a, lda, b, ldb, alpha,
             beta, q, ldq, z, ldz, m, pl, pr, dif, work,
             lwork, iwork, liwork, info )
```

Discussion

This routine reorders the generalized real-Schur/Schur decomposition of a real/complex matrix pair (A, B) (in terms of an orthogonal/unitary equivalence transformation $Q' * (A, B) * Z$), so that a selected cluster of eigenvalues appears in the leading diagonal blocks of the pair (A, B) . The leading columns of Q and Z form orthonormal/unitary bases of the corresponding left and right eigenspaces (deflating subspaces). (A, B) must be in generalized real-Schur/Schur canonical form (as returned by [?gges](#)), that is, A and B are both upper triangular.

[?tgsen](#) also computes the generalized eigenvalues

$\omega = (\text{alphar}(j) + \text{alphai}(j)*i)/\text{beta}(j)$ (for real flavors)
 $\omega = \text{alpha}(j)/\text{beta}(j)$ (for complex flavors)
of the reordered matrix pair (A, B) .

Optionally, the routine computes the estimates of reciprocal condition numbers for eigenvalues and eigenspaces. These are $\text{Difu}[(A_{11}, B_{11}), (A_{22}, B_{22})]$ and $\text{Difl}[(A_{11}, B_{11}), (A_{22}, B_{22})]$, that is, the separation(s) between the matrix pairs (A_{11}, B_{11}) and (A_{22}, B_{22}) that

correspond to the selected cluster and the eigenvalues outside the cluster, respectively, and norms of "projections" onto left and right eigenspaces with respect to the selected cluster in the (1,1)-block.

Input Parameters

ijob **INTEGER**. Specifies whether condition numbers are required for the cluster of eigenvalues (*p1* and *pr*) or the deflating subspaces *Difu* and *Difl*.
 If *ijob*=0, only reorder with respect to *select*;
 If *ijob*=1, reciprocal of norms of "projections" onto left and right eigenspaces with respect to the selected cluster (*p1* and *pr*);
 If *ijob*=2, compute upper bounds on *Difu* and *Difl*, using F-norm-based estimate (*dif* (1:2));
 If *ijob*=3, compute estimate of *Difu* and *Difl*, using 1-norm-based estimate (*dif* (1:2)). This option is about 5 times as expensive as *ijob*=2;
 If *ijob*=4, compute *p1*, *pr* and *dif* (i.e., options 0, 1 and 2 above). This is an economic version to get it all;
 If *ijob*=5, compute *p1*, *pr* and *dif* (i.e., options 0, 1 and 3 above).

wantq, wantz **LOGICAL**.
 If *wantq* = **.TRUE.**, update the left transformation matrix *Q*;
 If *wantq* = **.FALSE.**, do not update *Q*;
 If *wantz* = **.TRUE.**, update the right transformation matrix *Z*;
 If *wantz* = **.FALSE.**, do not update *Z*.

select **LOGICAL**.
 Array, **DIMENSION** at least $\max(1, n)$.
 Specifies the eigenvalues in the selected cluster.
 To select an eigenvalue α_j , *select*(*j*) must be **.TRUE.**
 For real flavors: to select a complex conjugate pair of eigenvalues α_j and α_{j+1} (corresponding 2 by 2 diagonal

block), *select(j)* and/or *select(j+1)* must be set to `.TRUE.`; the complex conjugate α_j and α_{j+1} must be either both included in the cluster or both excluded.

n **INTEGER.** The order of the matrices *A* and *B* ($n \geq 0$).

a, b, q, z, work **REAL** for *stgsen*
DOUBLE PRECISION for *dtgsen*
COMPLEX for *ctgsen*
DOUBLE COMPLEX for *ztgsen*.

Arrays:

a(lda,)* contains the matrix *A*.
For real flavors: *A* is upper quasi-triangular, with (*A, B*) in generalized real Schur canonical form.
For complex flavors: *A* is upper triangular, in generalized Schur canonical form.
The second dimension of *a* must be at least $\max(1, n)$.

b(ldb,)* contains the matrix *B*.
For real flavors: *B* is upper triangular, with (*A, B*) in generalized real Schur canonical form.
For complex flavors: *B* is upper triangular, in generalized Schur canonical form.
The second dimension of *b* must be at least $\max(1, n)$.

q(ldq,)*
If *wantq* = `.TRUE.`, then *q* is an *n*-by-*n* matrix;
If *wantq* = `.FALSE.`, then *q* is not referenced.
The second dimension of *q* must be at least $\max(1, n)$.

z(ldz,)*
If *wantz* = `.TRUE.`, then *z* is an *n*-by-*n* matrix;
If *wantz* = `.FALSE.`, then *z* is not referenced.
The second dimension of *z* must be at least $\max(1, n)$.

work(lwork) is a workspace array. If *ijob*=0, *work* is not referenced.

lda **INTEGER.** The first dimension of *a*; at least $\max(1, n)$.

ldb **INTEGER.** The first dimension of *b*; at least $\max(1, n)$.

<i>ldq</i>	INTEGER . The first dimension of <i>q</i> ; $ldq \geq 1$. If <i>wantq</i> = .TRUE. , then $ldq \geq \max(1, n)$.
<i>ldz</i>	INTEGER . The first dimension of <i>z</i> ; $ldz \geq 1$. If <i>wantz</i> = .TRUE. , then $ldz \geq \max(1, n)$.
<i>lwork</i>	INTEGER . The dimension of the array <i>work</i> . For real flavors: If <i>ijob</i> = 1, 2, or 4, $lwork \geq \max(4n+16, 2m(n-m))$. If <i>ijob</i> = 3 or 5, $lwork \geq \max(4n+16, 4m(n-m))$. For complex flavors: If <i>ijob</i> = 1, 2, or 4, $lwork \geq \max(1, 2m(n-m))$. If <i>ijob</i> = 3 or 5, $lwork \geq \max(1, 4m(n-m))$.
<i>iwork</i>	INTEGER . Workspace array, DIMENSION (<i>liwork</i>). If <i>ijob</i> =0, <i>iwork</i> is not referenced.
<i>liwork</i>	INTEGER . The dimension of the array <i>iwork</i> . For real flavors: If <i>ijob</i> = 1, 2, or 4, $liwork \geq n+6$. If <i>ijob</i> = 3 or 5, $liwork \geq \max(n+6, 2m(n-m))$. For complex flavors: If <i>ijob</i> = 1, 2, or 4, $liwork \geq n+2$. If <i>ijob</i> = 3 or 5, $liwork \geq \max(n+2, 2m(n-m))$.

Output Parameters

<i>a, b</i>	Overwritten by the reordered matrices <i>A</i> and <i>B</i> , respectively.
<i>alphar, alphai</i>	REAL for <i>stgsen</i> ; DOUBLE PRECISION for <i>dtgsen</i> . Arrays, DIMENSION at least $\max(1, n)$. Contain values that form generalized eigenvalues in real flavors. See <i>beta</i> .
<i>alpha</i>	COMPLEX for <i>ctgsen</i> ; DOUBLE COMPLEX for <i>ztgsen</i> . Array, DIMENSION at least $\max(1, n)$. Contain values that form generalized eigenvalues in complex flavors. See <i>beta</i> .

beta REAL for *stgsen*
DOUBLE PRECISION for *dtgsen*
COMPLEX for *ctgsen*
DOUBLE COMPLEX for *ztgsen*.
Array, DIMENSION at least $\max(1,n)$.
For real flavors:
On exit, $(\text{alphar}(j) + \text{alphai}(j)*i)/\text{beta}(j)$, $j=1,\dots,n$,
will be the generalized eigenvalues.
 $\text{alphar}(j) + \text{alphai}(j)*i$ and $\text{beta}(j)$, $j=1,\dots,n$ are the
diagonals of the complex Schur form (S,T) that would
result if the 2-by-2 diagonal blocks of the real
generalized Schur form of (A,B) were further reduced to
triangular form using complex unitary transformations.
If $\text{alphai}(j)$ is zero, then the j -th eigenvalue is real; if
positive, then the j -th and $(j+1)$ -st eigenvalues are a
complex conjugate pair, with $\text{alphai}(j+1)$ negative.
For complex flavors:
The diagonal elements of A and B , respectively, when
the pair (A,B) has been reduced to generalized Schur
form. $\text{alpha}(i)/\text{beta}(i)$, $i=1,\dots,n$ are the generalized
eigenvalues.

q If *wantq* = .TRUE., then, on exit, Q has been
postmultiplied by the left orthogonal transformation
matrix which reorder (A, B) . The leading m columns of
 Q form orthonormal bases for the specified pair of left
eigenspaces (deflating subspaces).

z If *wantz* = .TRUE., then, on exit, Z has been
postmultiplied by the left orthogonal transformation
matrix which reorder (A, B) . The leading m columns of Z
form orthonormal bases for the specified pair of left
eigenspaces (deflating subspaces).

m INTEGER. The dimension of the specified pair of left
and right eigen-spaces (deflating subspaces); $0 \leq m \leq n$.

pl, pr REAL for single precision flavors;
DOUBLE PRECISION for double precision flavors.
If *ijob* = 1, 4, or 5, *pl* and *pr* are lower bounds on the

reciprocal of the norm of "projections" onto left and right eigenspaces with respect to the selected cluster.
 $0 < pl, pr \leq 1$. If $m = 0$ or $m = n$, $pl = pr = 1$.
If $ijob = 0, 2$ or 3 , pl and pr are not referenced

dif **REAL** for single precision flavors;
 DOUBLE PRECISION for double precision flavors.
 Array, **DIMENSION** (2).
 If $ijob \geq 2$, *dif*(1:2) store the estimates of Difu and Difl.
 If $ijob = 2$ or 4 , *dif*(1:2) are F-norm-based upper bounds on Difu and Difl.
 If $ijob = 3$ or 5 , *dif*(1:2) are 1-norm-based estimates of Difu and Difl. If $m = 0$ or n ,
 dif(1:2) = F-norm([A, B]).
 If $ijob = 0$ or 1 , *dif* is not referenced.

work(1) If $ijob$ is not 0 and $info = 0$, on exit, *work*(1) contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

iwork(1) If $ijob$ is not 0 and $info = 0$, on exit, *iwork*(1) contains the minimum value of *liwork* required for optimum performance. Use this *liwork* for subsequent runs.

info **INTEGER**.
 If $info = 0$, the execution is successful.
 If $info = -i$, the i th parameter had an illegal value.
 If $info = 1$, Reordering of (A, B) failed because the transformed matrix pair (A, B) would be too far from generalized Schur form; the problem is very ill-conditioned. (A, B) may have been partially reordered. If requested, 0 is returned in *dif*(*), pl and pr .

?tgsyl

Solves the generalized Sylvester equation.

```
call stgsyl ( trans, ijob, m, n, a, lda, b, ldb, c, ldc, d, ldd, e,
             lde, f, ldf, scale, dif, work, lwork, iwork, info )
call dtgsyl ( trans, ijob, m, n, a, lda, b, ldb, c, ldc, d, ldd, e,
             lde, f, ldf, scale, dif, work, lwork, iwork, info )
call ctgsyl ( trans, ijob, m, n, a, lda, b, ldb, c, ldc, d, ldd, e,
             lde, f, ldf, scale, dif, work, lwork, iwork, info )
call ztgsyl ( trans, ijob, m, n, a, lda, b, ldb, c, ldc, d, ldd, e,
             lde, f, ldf, scale, dif, work, lwork, iwork, info )
```

Discussion

This routine solves the generalized Sylvester equation:

$$A R - L B = scale * C$$

$$D R - L E = scale * F$$

where R and L are unknown m -by- n matrices, (A, D) , (B, E) and (C, F) are given matrix pairs of size m -by- m , n -by- n and m -by- n , respectively, with real/complex entries. (A, D) and (B, E) must be in generalized real-Schur/Schur canonical form, that is, A, B are upper quasi-triangular/triangular and D, E are upper triangular.

The solution (R, L) overwrites (C, F) . The factor *scale*, $0 \leq scale \leq 1$, is an output scaling factor chosen to avoid overflow.

In matrix notation the above equation is equivalent to the following:

solve $Zx = scale * b$, where Z is defined as

$$Z = \begin{pmatrix} kron(I_n, A) & -kron(B', I_m) \\ kron(I_n, D) & -kron(E', I_m) \end{pmatrix}$$

Here I_k is the identity matrix of size k and X' is the transpose/conjugate-transpose of X . $kron(X, Y)$ is the Kronecker product between the matrices X and Y .

If $trans = 'T'$ (for real flavors), or $trans = 'C'$ (for complex flavors), the routine `?tgstyl` solves the transposed/conjugate-transposed system $Z' y = scale * b$, which is equivalent to solve for R and L in

$$\begin{aligned} A' R + D' L &= scale * C \\ R B' + L E' &= scale * (-F) \end{aligned}$$

This case ($trans = 'T'$ for `stgsyl/dtgsyl` or $trans = 'C'$ for `ctgsyl/ztgsyl`) is used to compute an one-norm-based estimate of $Dif[(A,D), (B,E)]$, the separation between the matrix pairs (A,D) and (B,E) , using `slacon/clacon`.

If $ijob \geq 1$, `?tgstyl` computes a Frobenius norm-based estimate of $Dif[(A,D), (B,E)]$. That is, the reciprocal of a lower bound on the reciprocal of the smallest singular value of Z . This is a level 3 BLAS algorithm.

Input Parameters

trans CHARACTER*1. Must be 'N', 'T', or 'C'.
 If $trans = 'N'$, solve the generalized Sylvester equation.
 If $trans = 'T'$, solve the 'transposed' system (for real flavors only).
 If $trans = 'C'$, solve the 'conjugate transposed' system (for complex flavors only).

ijob INTEGER. Specifies what kind of functionality to be performed:
 If $ijob = 0$, solve the generalized Sylvester equation only ;
 If $ijob = 1$, perform the functionality of $ijob = 0$ and $ijob = 3$;
 If $ijob = 2$, perform the functionality of $ijob = 0$ and $ijob = 4$;
 If $ijob = 3$, only an estimate of $Dif[(A,D), (B,E)]$ is computed (look ahead strategy is used);

If *ijob*=4, only an estimate of $\text{Dif}[(A,D), (B,E)]$ is computed (*?gecon* on sub-systems is used).

If *trans* = 'T' or 'C', *ijob* is not referenced.

m INTEGER.

The order of the matrices *A* and *D*, and the row dimension of the matrices *C*, *F*, *R* and *L*.

n INTEGER.

The order of the matrices *B* and *E*, and the column dimension of the matrices *C*, *F*, *R* and *L*.

a,b,c,d,e,f,work REAL for *stgsyl*
 DOUBLE PRECISION for *dtgsyl*
 COMPLEX for *ctgsyl*
 DOUBLE COMPLEX for *ztgsyl*.

Arrays:

a(lda,)* contains the upper quasi-triangular (for real flavors) or upper triangular (for complex flavors) matrix *A*.

The second dimension of *a* must be at least $\max(1, m)$.

b(ldb,)* contains the upper quasi-triangular (for real flavors) or upper triangular (for complex flavors) matrix *B*.

The second dimension of *b* must be at least $\max(1, n)$.

c ldc,)* contains the right-hand-side of the first matrix equation in the generalized Sylvester equation (as defined by *trans*)

The second dimension of *c* must be at least $\max(1, n)$.

d(ldd,)* contains the upper triangular matrix *D*.

The second dimension of *d* must be at least $\max(1, m)$.

e(lde,)* contains the upper triangular matrix *E*.

The second dimension of *e* must be at least $\max(1, n)$.

f(ldf,)* contains the right-hand-side of the second matrix equation in the generalized Sylvester equation (as defined by *trans*)

The second dimension of *f* must be at least $\max(1, n)$.

work(lwork) is a workspace array. If *ijob*=0, *work* is not referenced.

<i>lda</i>	INTEGER . The first dimension of <i>a</i> ; at least $\max(1, m)$.
<i>ldb</i>	INTEGER . The first dimension of <i>b</i> ; at least $\max(1, n)$.
<i>ldc</i>	INTEGER . The first dimension of <i>c</i> ; at least $\max(1, m)$.
<i>ldd</i>	INTEGER . The first dimension of <i>d</i> ; at least $\max(1, m)$.
<i>lde</i>	INTEGER . The first dimension of <i>e</i> ; at least $\max(1, n)$.
<i>ldf</i>	INTEGER . The first dimension of <i>f</i> ; at least $\max(1, m)$.
<i>lwork</i>	INTEGER . The dimension of the array <i>work</i> . $lwork \geq 1$. If <i>ijob</i> = 1 or 2 and <i>trans</i> = 'N', $lwork \geq 2mn$.
<i>iwork</i>	INTEGER . Workspace array, DIMENSION at least $(m+n+6)$ for real flavors, and at least $(m+n+2)$ for complex flavors. If <i>ijob</i> =0, <i>iwork</i> is not referenced.

Output Parameters

<i>c</i>	If <i>ijob</i> =0, 1, or 2, overwritten by the solution <i>R</i> . If <i>ijob</i> =3 or 4 and <i>trans</i> = 'N', <i>c</i> holds <i>R</i> , the solution achieved during the computation of the Dif-estimate.
<i>f</i>	If <i>ijob</i> =0, 1, or 2, overwritten by the solution <i>L</i> . If <i>ijob</i> =3 or 4 and <i>trans</i> = 'N', <i>f</i> holds <i>L</i> , the solution achieved during the computation of the Dif-estimate.
<i>dif</i>	REAL for single-precision flavors DOUBLE PRECISION for double-precision flavors. On exit, <i>dif</i> is the reciprocal of a lower bound of the reciprocal of the Dif-function, i.e. <i>dif</i> is an upper bound of $\text{Dif}[(A,D), (B,E)] = \sigma_{\min}(Z)$, where <i>Z</i> as in (2). If <i>ijob</i> = 0, or <i>trans</i> = 'T' (for real flavors), or <i>trans</i> = 'C' (for complex flavors), <i>dif</i> is not touched.

scale **REAL** for single-precision flavors
DOUBLE PRECISION for double-precision flavors.
 On exit, *scale* is the scaling factor in the generalized Sylvester equation. If $0 < scale < 1$, *c* and *f* hold the solutions *R* and *L*, respectively, to a slightly perturbed system but the input matrices *A*, *B*, *D* and *E* have not been changed. If *scale* = 0, *c* and *f* hold the solutions *R* and *L*, respectively, to the homogeneous system with $C = F = 0$. Normally, *scale* = 1.

work(1) If *ijob* is not 0 and *info* = 0, on exit, *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

info **INTEGER**.
 If *info* = 0, the execution is successful.
 If *info* = *-i*, the *i*th parameter had an illegal value.
 If *info* > 0, (*A*, *D*) and (*B*, *E*) have common or close eigenvalues.

?tgsna

Estimates reciprocal condition numbers for specified eigenvalues and/or eigenvectors of a pair of matrices in generalized real Schur canonical form.

```
call stgsna ( job, howmny, select, n, a, lda, b, ldb, vl, ldvl, vr,
             ldvr, s, dif, mm, m, work, lwork, iwork, info )
call dtgsna ( job, howmny, select, n, a, lda, b, ldb, vl, ldvl, vr,
             ldvr, s, dif, mm, m, work, lwork, iwork, info )
call ctgsna ( job, howmny, select, n, a, lda, b, ldb, vl, ldvl, vr,
             ldvr, s, dif, mm, m, work, lwork, iwork, info )
call ztgsna ( job, howmny, select, n, a, lda, b, ldb, vl, ldvl, vr,
             ldvr, s, dif, mm, m, work, lwork, iwork, info )
```

Discussion

The real flavors `stgsna/dtgsna` of this routine estimate reciprocal condition numbers for specified eigenvalues and/or eigenvectors of a matrix pair (A, B) in generalized real Schur canonical form (or of any matrix pair $(QA Z^T, QB Z^T)$ with orthogonal matrices Q and Z).

(A, B) must be in generalized real Schur form (as returned by `sgges/dgges`), that is, A is block upper triangular with 1-by-1 and 2-by-2 diagonal blocks. B is upper triangular.

The complex flavors `ctgsna/ztgsna` estimate reciprocal condition numbers for specified eigenvalues and/or eigenvectors of a matrix pair (A, B) . (A, B) must be in generalized Schur canonical form, that is, A and B are both upper triangular.

Input Parameters

`job` CHARACTER*1. Specifies whether condition numbers are required for eigenvalues or eigenvectors. Must be 'E' or 'V' or 'B'. If `job='E'`, for eigenvalues only (compute `s`).

If *job* = 'V', for eigenvectors only (compute *dif*).
 If *job* = 'B', for both eigenvalues and eigenvectors (compute both *s* and *dif*).

howmny CHARACTER*1. Must be 'A' or 'S'.
 If *howmny* = 'A', compute condition numbers for all eigenpairs.
 If *howmny* = 'S', compute condition numbers for selected eigenpairs specified by the logical array *select*.

select LOGICAL.
 Array, DIMENSION at least max (1, *n*).
 If *howmny* = 'S', *select* specifies the eigenpairs for which condition numbers are required.
 If *howmny* = 'A', *select* is not referenced.
 For real flavors:
 To select condition numbers for the eigenpair corresponding to a real eigenvalue ω_j , *select*(*j*) must be set to .TRUE.; to select condition numbers corresponding to a complex conjugate pair of eigenvalues ω_j and ω_{j+1} , either *select*(*j*) or *select*(*j*+1) must be set to .TRUE.
 For complex flavors:
 To select condition numbers for the corresponding *j*-th eigenvalue and/or eigenvector, *select*(*j*) must be set to .TRUE..

n INTEGER. The order of the square matrix pair (*A*, *B*) (*n* ≥ 0).

a, *b*, *vl*, *vr*, *work* REAL for *stgsna*
 DOUBLE PRECISION for *dtgsna*
 COMPLEX for *ctgsna*
 DOUBLE COMPLEX for *ztgsna*.
 Arrays:
a(*lda*, *) contains the upper quasi-triangular (for real flavors) or upper triangular (for complex flavors) matrix *A* in the pair (*A*, *B*) .
 The second dimension of *a* must be at least max(1, *n*).

$b(ldb, *)$ contains the upper triangular matrix B in the pair (A, B) .

The second dimension of b must be at least $\max(1, n)$.

If $job = 'E'$ or $'B'$,

$v1(ldv1, *)$ must contain left eigenvectors of (A, B) , corresponding to the eigenpairs specified by $howmny$ and $select$. The eigenvectors must be stored in consecutive columns of $v1$, as returned by $?tgevc$.

If $job = 'V'$, $v1$ is not referenced.

The second dimension of $v1$ must be at least $\max(1, m)$.

If $job = 'E'$ or $'B'$,

$vr(ldvr, *)$ must contain right eigenvectors of (A, B) , corresponding to the eigenpairs specified by $howmny$ and $select$. The eigenvectors must be stored in consecutive columns of vr , as returned by $?tgevc$.

If $job = 'V'$, vr is not referenced.

The second dimension of vr must be at least $\max(1, m)$.

$work(lwork)$ is a workspace array. If $job = 'E'$, $work$ is not referenced.

lda	INTEGER. The first dimension of a ; at least $\max(1, n)$.
ldb	INTEGER. The first dimension of b ; at least $\max(1, n)$.
$ldv1$	INTEGER. The first dimension of $v1$; $ldv1 \geq 1$. If $job = 'E'$ or $'B'$, then $ldv1 \geq \max(1, n)$.
$ldvr$	INTEGER. The first dimension of vr ; $ldvr \geq 1$. If $job = 'E'$ or $'B'$, then $ldvr \geq \max(1, n)$.
mm	INTEGER. The number of elements in the arrays s and dif ($mm \geq m$).
$lwork$	INTEGER. The dimension of the array $work$.

For real flavors:

$lwork \geq n$.

If $job = 'V'$ or $'B'$, $lwork \geq 2n(n+2)+16$.

For complex flavors:

$lwork \geq 1$.

If $job = 'V'$ or $'B'$, $lwork \geq 2n^2$.

iwork **INTEGER**. Workspace array, **DIMENSION** at least $(n+6)$ for real flavors, and at least $(n+2)$ for complex flavors. If *ijob* = 'E', *iwork* is not referenced.

Output Parameters

s **REAL** for single-precision flavors
DOUBLE PRECISION for double-precision flavors.
 Array, **DIMENSION** (*mm*).
 If *job* = 'E' or 'B', contains the reciprocal condition numbers of the selected eigenvalues, stored in consecutive elements of the array.
 If *job* = 'V', *s* is not referenced.
For real flavors:
 For a complex conjugate pair of eigenvalues two consecutive elements of *s* are set to the same value. Thus, *s*(*j*), *dif*(*j*), and the *j*-th columns of *vl* and *vr* all correspond to the same eigenpair (but not in general the *j*-th eigenpair, unless all eigenpairs are selected).

dif **REAL** for single-precision flavors
DOUBLE PRECISION for double-precision flavors.
 Array, **DIMENSION** (*mm*).
 If *job* = 'V' or 'B', contains the estimated reciprocal condition numbers of the selected eigenvectors, stored in consecutive elements of the array. If the eigenvalues cannot be reordered to compute *dif*(*j*), *dif*(*j*) is set to 0; this can only occur when the true value would be very small anyway.
 If *job* = 'E', *dif* is not referenced.
For real flavors:
 For a complex eigenvector, two consecutive elements of *dif* are set to the same value.
For complex flavors:
 For each eigenvalue/vector specified by *select*, *dif* stores a Frobenius norm-based estimate of Difl.

m **INTEGER.** The number of elements in the arrays *s* and *dif* used to store the specified condition numbers; for each selected eigenvalue one element is used. If *howmny* = 'A', *m* is set to *n*.

work(1) **work(1)** If *job* is not 'E' and *info* = 0, on exit, *work(1)* contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.

info **INTEGER.**
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.

Generalized Singular Value Decomposition

This section describes LAPACK computational routines used for finding the generalized singular value decomposition (GSVD) of two matrices A and B as

$$U^H A Q = D_1 * (0 \ R),$$

$$V^H B Q = D_2 * (0 \ R),$$

where U , V , and Q are orthogonal/unitary matrices, R is a nonsingular upper triangular matrix, and D_1 , D_2 are “diagonal” matrices of the structure detailed in the routines description section.

Table 5-7 Computational Routines for Generalized Singular Value Decomposition

Routine name	Operation performed
?ggsvp	Computes the preprocessing decomposition for the generalized SVD
?tgsja	Computes the generalized SVD of two upper triangular or trapezoidal matrices

You can use routines listed in the above table as well as the driver routine [?ggsvd](#) to find the GSVD of a pair of general rectangular matrices.

?ggsvp

Computes the preprocessing decomposition for the generalized SVD.

```
call sggsvp ( jobu, jobv, jobq, m, p, n, a, lda, b, ldb, tola, tolb,
              k, l, u, ldu, v, ldv, q, ldq, iwork, tau, work, info )
call dggsvp ( jobu, jobv, jobq, m, p, n, a, lda, b, ldb, tola, tolb,
              k, l, u, ldu, v, ldv, q, ldq, iwork, tau, work, info )
call cggsvp ( jobu, jobv, jobq, m, p, n, a, lda, b, ldb, tola, tolb,
              k, l, u, ldu, v, ldv, q, ldq, iwork, rwork, tau, work, info )
call zggsvp ( jobu, jobv, jobq, m, p, n, a, lda, b, ldb, tola, tolb,
              k, l, u, ldu, v, ldv, q, ldq, iwork, rwork, tau, work, info )
```

Discussion

This routine computes orthogonal matrices U , V and Q such that

$$U^H A Q = \begin{matrix} & n-k-l & k & l \\ & k & & \\ & l & & \\ m-k-l & & & \end{matrix} \begin{pmatrix} 0 & A_{12} & A_{13} \\ 0 & 0 & A_{23} \\ 0 & 0 & 0 \end{pmatrix}, \text{ if } m-k-l \geq 0$$

$$= \begin{matrix} & n-k-l & k & l \\ & k & & \\ m-k & & & \end{matrix} \begin{pmatrix} 0 & A_{12} & A_{13} \\ 0 & 0 & A_{23} \end{pmatrix}, \text{ if } m-k-l < 0$$

$$V^H B Q = \begin{matrix} & n-k-l & k & l \\ & l & & \\ p-l & & & \end{matrix} \begin{pmatrix} 0 & 0 & B_{13} \\ 0 & 0 & 0 \end{pmatrix}$$

where the k -by- k matrix A_{12} and l -by- l matrix B_{13} are nonsingular upper triangular; A_{23} is l -by- l upper triangular if $m-k-l \geq 0$, otherwise A_{23} is $(m-k)$ -by- l upper trapezoidal. The sum $k+l$ is equal to the effective numerical rank of the $(m+p)$ -by- n matrix $(A^H, B^H)^H$.

This decomposition is the preprocessing step for computing the Generalized Singular Value Decomposition (GSVD), see subroutine [?ggsvd](#).

Input Parameters

jobu CHARACTER*1. Must be 'U' or 'N'.
If *jobu* = 'U', orthogonal/unitary matrix U is computed.
If *jobu* = 'N', U is not computed.

jobv CHARACTER*1. Must be 'V' or 'N'.
If *jobv* = 'V', orthogonal/unitary matrix V is computed.
If *jobv* = 'N', V is not computed.

jobq CHARACTER*1. Must be 'Q' or 'N'.
If *jobq* = 'Q', orthogonal/unitary matrix Q is computed.
If *jobq* = 'N', Q is not computed.

m INTEGER. The number of rows of the matrix A ($m \geq 0$).

p INTEGER. The number of rows of the matrix B ($p \geq 0$).

n INTEGER. The number of columns of the matrices A and B ($n \geq 0$).

a, b, tau, work REAL for `sggsvp`
DOUBLE PRECISION for `dggsvp`
COMPLEX for `cggsvp`
DOUBLE COMPLEX for `zggsvp`.
Arrays:
a(lda,)* contains the m -by- n matrix A .
The second dimension of *a* must be at least $\max(1, n)$.
b(ldb,)* contains the p -by- n matrix B .
The second dimension of *b* must be at least $\max(1, n)$.
tau()* is a workspace array. The dimension of *tau* must be at least $\max(1, n)$.
work()* is a workspace array. The dimension of *work* must be at least $\max(1, 3n, m, p)$.

<i>lda</i>	INTEGER . The first dimension of <i>a</i> ; at least $\max(1, m)$.
<i>ldb</i>	INTEGER . The first dimension of <i>b</i> ; at least $\max(1, p)$.
<i>tola, tolB</i>	REAL for single-precision flavors DOUBLE PRECISION for double-precision flavors. <i>tolA</i> and <i>tolB</i> are the thresholds to determine the effective numerical rank of matrix <i>B</i> and a subblock of <i>A</i> . Generally, they are set to $tolA = \max(m, n) * A * \text{MACHEPS}$, $tolB = \max(p, n) * B * \text{MACHEPS}$. The size of <i>tolA</i> and <i>tolB</i> may affect the size of backward errors of the decomposition.
<i>ldu</i>	INTEGER . The first dimension of the output array <i>u</i> . $ldu \geq \max(1, m)$ if <i>jobu</i> = 'U'; $ldu \geq 1$ otherwise.
<i>ldv</i>	INTEGER . The first dimension of the output array <i>v</i> . $ldv \geq \max(1, p)$ if <i>jobv</i> = 'V'; $ldv \geq 1$ otherwise.
<i>ldq</i>	INTEGER . The first dimension of the output array <i>q</i> . $ldq \geq \max(1, n)$ if <i>jobq</i> = 'Q'; $ldq \geq 1$ otherwise.
<i>iwork</i>	INTEGER . Workspace array, DIMENSION at least $\max(1, n)$.
<i>rwork</i>	REAL for <i>cggsvp</i> DOUBLE PRECISION for <i>zggsvp</i> . Workspace array, DIMENSION at least $\max(1, 2n)$. Used in complex flavors only.

Output Parameters

<i>a</i>	Overwritten by the triangular (or trapezoidal) matrix described in the <i>Discussion</i> section.
<i>b</i>	Overwritten by the triangular matrix described in the <i>Discussion</i> section.
<i>k, l</i>	INTEGER . On exit, <i>k</i> and <i>l</i> specify the dimension of subblocks. The sum $k + l$ is equal to effective numerical rank of $(A^H, B^H)^H$.

u, v, q REAL for *sggsvp*
 DOUBLE PRECISION for *dggsvp*
 COMPLEX for *cggsvp*
 DOUBLE COMPLEX for *zggsvp*.
 Arrays:
 If *jobu* = 'U', *u(ldu,*)* contains the orthogonal/unitary matrix *U*.
 The second dimension of *u* must be at least $\max(1, m)$.
 If *jobu* = 'N', *u* is not referenced.
 If *jobv* = 'V', *v(ldv,*)* contains the orthogonal/unitary matrix *V*.
 The second dimension of *v* must be at least $\max(1, m)$.
 If *jobv* = 'N', *v* is not referenced.
 If *jobq* = 'Q', *q(ldq,*)* contains the orthogonal/unitary matrix *Q*.
 The second dimension of *q* must be at least $\max(1, n)$.
 If *jobq* = 'N', *q* is not referenced.

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = *-i*, the *i*th parameter had an illegal value.

?tgsja

Computes the generalized SVD of two upper triangular or trapezoidal matrices.

```
call stgsja ( jobu, jobv, jobq, m, p, n, k, l, a, lda, b, ldb, tola,
             tolB, alpha, beta, u, ldu, v, ldv, q, ldq, work, ncycle, info )
call dtgsja ( jobu, jobv, jobq, m, p, n, k, l, a, lda, b, ldb, tola,
             tolB, alpha, beta, u, ldu, v, ldv, q, ldq, work, ncycle, info )
call ctgsja ( jobu, jobv, jobq, m, p, n, k, l, a, lda, b, ldb, tola,
             tolB, alpha, beta, u, ldu, v, ldv, q, ldq, work, ncycle, info )
call ztgsja ( jobu, jobv, jobq, m, p, n, k, l, a, lda, b, ldb, tola,
             tolB, alpha, beta, u, ldu, v, ldv, q, ldq, work, ncycle, info )
```

Discussion

This routine computes the generalized singular value decomposition (GSVD) of two real/complex upper triangular (or trapezoidal) matrices A and B . On entry, it is assumed that matrices A and B have the following forms, which may be obtained by the preprocessing subroutine [?ggsvp](#) from a general m -by- n matrix A and p -by- n matrix B :

$$A = \begin{matrix} & n-k-l & k & l \\ & k & & \\ & l & & \\ m-k-l & & & \end{matrix} \begin{pmatrix} 0 & A_{12} & A_{13} \\ 0 & 0 & A_{23} \\ 0 & 0 & 0 \end{pmatrix}, \quad \text{if } m-k-l \geq 0$$

$$= \begin{matrix} & n-k-l & k & l \\ & k & & \\ m-k & & & \end{matrix} \begin{pmatrix} 0 & A_{12} & A_{13} \\ 0 & 0 & A_{23} \end{pmatrix}, \quad \text{if } m-k-l < 0$$

$$B = \begin{matrix} & n-k-l & k & l \\ & l \left(\begin{matrix} 0 & 0 & B_{13} \\ 0 & 0 & 0 \end{matrix} \right) \\ p-l & & & \end{matrix}$$

where the k -by- k matrix A_{12} and l -by- l matrix B_{13} are nonsingular upper triangular; A_{23} is l -by- l upper triangular if $m-k-l \geq 0$, otherwise A_{23} is $(m-k)$ -by- l upper trapezoidal.

On exit,

$$U^H A Q = D_1 * (0 \ R), \quad V^H B Q = D_2 * (0 \ R),$$

where U , V and Q are orthogonal/unitary matrices, R is a nonsingular upper triangular matrix, and D_1 and D_2 are "diagonal" matrices, which are of the following structures:

If $m-k-l \geq 0$,

$$D_1 = \begin{matrix} & k & l \\ & l \left(\begin{matrix} I & 0 \\ 0 & C \\ 0 & 0 \end{matrix} \right) \\ m-k-l & & \end{matrix}$$

$$D_2 = \begin{matrix} & k & l \\ & l \left(\begin{matrix} 0 & S \\ 0 & 0 \end{matrix} \right) \\ p-l & & \end{matrix}$$

$$(0 \ R) = \begin{matrix} & n-k-l & k & l \\ & k \left(\begin{matrix} 0 & R_{11} & R_{12} \\ 0 & 0 & R_{22} \end{matrix} \right) \\ l & & & \end{matrix}$$

where

$$C = \text{diag}(\alpha(k+1), \dots, \alpha(k+1))$$

$$S = \text{diag}(\beta(k+1), \dots, \beta(k+1))$$

$$C^2 + S^2 = I$$

R is stored in $a(1:k+1, n-k-l+1:n)$ on exit.

If $m-k-l < 0$,

$$D_1 = \begin{matrix} & k & m-k & k+l-m \\ m-k & \begin{pmatrix} I & 0 & 0 \\ 0 & C & 0 \end{pmatrix} \end{matrix}$$

$$D_2 = \begin{matrix} & k & m-k & k+l-m \\ m-k & \begin{pmatrix} 0 & S & 0 \\ 0 & 0 & I \\ 0 & 0 & 0 \end{pmatrix} \\ k+l-m & & & \\ p-l & & & \end{matrix}$$

$$\begin{pmatrix} 0 & R \end{pmatrix} = \begin{matrix} & n-k-l & k & m-k & k+l-m \\ m-k & \begin{pmatrix} 0 & R_{11} & R_{12} & R_{13} \\ 0 & 0 & R_{22} & R_{23} \\ 0 & 0 & 0 & R_{33} \end{pmatrix} \\ k+l-m & & & & \end{matrix}$$

where

$$C = \text{diag}(\alpha(k+1), \dots, \alpha(m)),$$

$$S = \text{diag}(\beta(k+1), \dots, \beta(m)),$$

$$C^2 + S^2 = I$$

On exit, $\begin{pmatrix} R_{11} & R_{12} & R_{13} \\ 0 & R_{22} & R_{23} \end{pmatrix}$ is stored in $\mathbf{a}(1:m, n-k-l+1:n)$ and R_{33} is stored

in $\mathbf{b}(m-k+1:l, n+m-k-l+1:n)$.

The computation of the orthogonal/unitary transformation matrices U , V or Q is optional. These matrices may either be formed explicitly, or they may be postmultiplied into input matrices U_1 , V_1 , or Q_1 .

Input Parameters

<i>jobu</i>	<p>CHARACTER*1. Must be 'U', 'I', or 'N'.</p> <p>If <i>jobu</i> = 'U', <i>u</i> must contain an orthogonal/unitary matrix U_l on entry.</p> <p>If <i>jobu</i> = 'I', <i>u</i> is initialized to the unit matrix.</p> <p>If <i>jobu</i> = 'N', <i>u</i> is not computed.</p>
<i>jobv</i>	<p>CHARACTER*1. Must be 'V', 'I', or 'N'.</p> <p>If <i>jobv</i> = 'V', <i>v</i> must contain an orthogonal/unitary matrix V_l on entry.</p> <p>If <i>jobv</i> = 'I', <i>v</i> is initialized to the unit matrix.</p> <p>If <i>jobv</i> = 'N', <i>v</i> is not computed.</p>
<i>jobq</i>	<p>CHARACTER*1. Must be 'Q', 'I', or 'N'.</p> <p>If <i>jobq</i> = 'Q', <i>q</i> must contain an orthogonal/unitary matrix Q_l on entry.</p> <p>If <i>jobq</i> = 'I', <i>q</i> is initialized to the unit matrix.</p> <p>If <i>jobq</i> = 'N', <i>q</i> is not computed.</p>
<i>m</i>	INTEGER. The number of rows of the matrix <i>A</i> ($m \geq 0$).
<i>p</i>	INTEGER. The number of rows of the matrix <i>B</i> ($p \geq 0$).
<i>n</i>	INTEGER. The number of columns of the matrices <i>A</i> and <i>B</i> ($n \geq 0$).
<i>k, l</i>	INTEGER. Specify the subblocks in the input matrices <i>A</i> and <i>B</i> , whose GSVD is going to be computed by ?tgsja.
<i>a, b, u, v, q, work</i>	<p>REAL for stgsja</p> <p>DOUBLE PRECISION for dtgsja</p> <p>COMPLEX for ctgsja</p> <p>DOUBLE COMPLEX for ztgsja.</p> <p>Arrays:</p> <p><i>a</i>(<i>lda</i>,*) contains the <i>m</i>-by-<i>n</i> matrix <i>A</i>. The second dimension of <i>a</i> must be at least max(1, <i>n</i>).</p> <p><i>b</i>(<i>ldb</i>,*) contains the <i>p</i>-by-<i>n</i> matrix <i>B</i>. The second dimension of <i>b</i> must be at least max(1, <i>n</i>).</p>

If *jobu* = 'U', *u(ldu,*)* must contain a matrix U_I (usually the orthogonal/unitary matrix returned by ?ggsvp).

The second dimension of *u* must be at least $\max(1, m)$.

If *jobv* = 'V', *v(ldv,*)* must contain a matrix V_I (usually the orthogonal/unitary matrix returned by ?ggsvp).

The second dimension of *v* must be at least $\max(1, p)$.

If *jobq* = 'Q', *q(ldq,*)* must contain a matrix Q_I (usually the orthogonal/unitary matrix returned by ?ggsvp).

The second dimension of *q* must be at least $\max(1, n)$.

work()* is a workspace array. The dimension of *work* must be at least $\max(1, 2n)$.

lda INTEGER. The first dimension of *a*; at least $\max(1, m)$.

ldb INTEGER. The first dimension of *b*; at least $\max(1, p)$.

ldu INTEGER. The first dimension of the array *u*.
ldu $\geq \max(1, m)$ if *jobu* = 'U'; *ldu* ≥ 1 otherwise.

ldv INTEGER. The first dimension of the array *v*.
ldv $\geq \max(1, p)$ if *jobv* = 'V'; *ldv* ≥ 1 otherwise.

ldq INTEGER. The first dimension of the array *q*.
ldq $\geq \max(1, n)$ if *jobq* = 'Q'; *ldq* ≥ 1 otherwise.

tola, tolb REAL for single-precision flavors
 DOUBLE PRECISION for double-precision flavors.
tola and *tolb* are the convergence criteria for the Jacobi-Kogbetliantz iteration procedure. Generally, they are the same as used in ?ggsvp :

tola = $\max(m, n) * |A| * \text{MACHEPS}$,

tolb = $\max(p, n) * |B| * \text{MACHEPS}$.

Output Parameters

a On exit, *a*(*n-k+1:n*, 1: $\min(k+1, m)$) contains the triangular matrix *R* or part of *R*.

<i>b</i>	On exit, if necessary, $b(m-k+1: l, n+m-k-l+1: n)$ contains a part of R .
<i>alpha, beta</i>	<p>REAL for single-precision flavors DOUBLE PRECISION for double-precision flavors. Arrays, DIMENSION at least $\max(1, n)$. Contain the generalized singular value pairs of A and B:</p> <p>$alpha(1:k) = 1,$ $beta(1:k) = 0,$</p> <p>and if $m-k-l \geq 0,$ $alpha(k+1:k+l) = \text{diag}(C),$ $beta(k+1:k+l) = \text{diag}(S),$</p> <p>or if $m-k-l < 0,$ $alpha(k+1:m) = C, alpha(m+1:k+l) = 0$ $beta(k+1:m) = S, beta(m+1:k+l) = 1.$</p> <p>Furthermore, if $k+l < n,$ $alpha(k+l+1:n) = 0$ and $beta(k+l+1:n) = 0.$</p>
<i>u</i>	<p>If $jobu = 'I',$ u contains the orthogonal/unitary matrix U. If $jobu = 'U',$ u contains the product $U_I U$. If $jobu = 'N',$ u is not referenced.</p>
<i>v</i>	<p>If $jobv = 'I',$ v contains the orthogonal/unitary matrix U. If $jobv = 'V',$ v contains the product $V_I V$. If $jobv = 'N',$ v is not referenced.</p>
<i>q</i>	<p>If $jobq = 'I',$ q contains the orthogonal/unitary matrix U. If $jobq = 'Q',$ q contains the product $Q_I Q$. If $jobq = 'N',$ q is not referenced.</p>
<i>ncycle</i>	INTEGER. The number of cycles required for convergence.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

If *info* = 1, the procedure does not converge after
MAXIT cycles.

Driver Routines

Each of the LAPACK driver routines solves a complete problem. To arrive at the solution, driver routines typically call a sequence of appropriate [computational routines](#).

Driver routines are described in the following sections:

[Linear Least Squares \(LLS\) Problems](#)

[Generalized LLS Problems](#)

[Symmetric Eigenproblems](#)

[Nonsymmetric Eigenproblems](#)

[Singular Value Decomposition](#)

[Generalized Symmetric Definite Eigenproblems](#)

[Generalized Nonsymmetric Eigenproblems](#)

Linear Least Squares (LLS) Problems

This section describes LAPACK driver routines used for solving linear least-squares problems. [Table 5-8](#) lists routines described in more detail below.

Table 5-8 **Driver Routines for Solving LLS Problems**

Routine Name	Operation performed
?gels	Uses QR or LQ factorization to solve a overdetermined or underdetermined linear system with full rank matrix.
?gelsy	Computes the minimum-norm solution to a linear least squares problem using a complete orthogonal factorization of A.
?gelss	Computes the minimum-norm solution to a linear least squares problem using the singular value decomposition of A.
?gelsd	Computes the minimum-norm solution to a linear least squares problem using the singular value decomposition of A and a divide and conquer method.

?gels

Uses *QR* or *LQ* factorization to solve a overdetermined or underdetermined linear system with full rank matrix.

```
call sgels ( trans, m, n, nrhs, a, lda, b, ldb, work, lwork, info )
call dgels ( trans, m, n, nrhs, a, lda, b, ldb, work, lwork, info )
call cgels ( trans, m, n, nrhs, a, lda, b, ldb, work, lwork, info )
call zgels ( trans, m, n, nrhs, a, lda, b, ldb, work, lwork, info )
```

Discussion

This routine solves overdetermined or underdetermined real/ complex linear systems involving an *m*-by-*n* matrix *A*, or its transpose/ conjugate-transpose, using a *QR* or *LQ* factorization of *A*. It is assumed that *A* has full rank.

The following options are provided:

1. If *trans* = 'N' and *m* ≥ *n*: find the least squares solution of an overdetermined system, that is, solve the least squares problem

$$\text{minimize } \| b - A x \|_2$$
2. If *trans* = 'N' and *m* < *n*: find the minimum norm solution of an underdetermined system $A X = B$.
3. If *trans* = 'T' or 'C' and *m* ≥ *n*: find the minimum norm solution of an underdetermined system $A^H X = B$.
4. If *trans* = 'T' or 'C' and *m* < *n*: find the least squares solution of an overdetermined system, that is, solve the least squares problem

$$\text{minimize } \| b - A^H x \|_2$$

Several right hand side vectors *b* and solution vectors *x* can be handled in a single call; they are stored as the columns of the *m*-by-*nrhs* right hand side matrix *B* and the *n*-by-*nrh* solution matrix *X*.

Input Parameters

<i>trans</i>	<p>CHARACTER*1. Must be 'N', 'T', or 'C'.</p> <p>If <i>trans</i> = 'N', the linear system involves matrix A;</p> <p>If <i>trans</i> = 'T', the linear system involves the transposed matrix A^T (for real flavors only);</p> <p>If <i>trans</i> = 'C', the linear system involves the conjugate-transposed matrix A^H (for complex flavors only).</p>
<i>m</i>	INTEGER. The number of rows of the matrix A ($m \geq 0$).
<i>n</i>	INTEGER. The number of columns of the matrix A ($n \geq 0$).
<i>nrhs</i>	INTEGER. The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).
<i>a, b, work</i>	<p>REAL for <i>sgels</i></p> <p>DOUBLE PRECISION for <i>dgels</i></p> <p>COMPLEX for <i>cgels</i></p> <p>DOUBLE COMPLEX for <i>zgels</i>.</p> <p>Arrays:</p> <p><i>a</i>(<i>lda</i>,*) contains the m-by-n matrix A. The second dimension of <i>a</i> must be at least $\max(1, n)$.</p> <p><i>b</i>(<i>ldb</i>,*) contains the matrix B of right hand side vectors, stored columnwise; B is m-by-$nrhs$ if <i>trans</i> = 'N', or n-by-$nrhs$ if <i>trans</i> = 'T' or 'C'. The second dimension of <i>b</i> must be at least $\max(1, nrhs)$.</p> <p><i>work</i>(<i>lwork</i>) is a workspace array.</p>
<i>lda</i>	INTEGER. The first dimension of <i>a</i> ; at least $\max(1, m)$.
<i>ldb</i>	INTEGER. The first dimension of <i>b</i> ; must be at least $\max(1, m, n)$.
<i>lwork</i>	<p>INTEGER. The size of the <i>work</i> array; must be at least $\min(m, n) + \max(1, m, n, nrhs)$.</p> <p>See <i>Application notes</i> for the suggested value of <i>lwork</i>.</p>

Output Parameters

- a*** On exit, overwritten by the factorization data as follows:
 if $m \geq n$, array ***a*** contains the details of the *QR* factorization of the matrix *A* as returned by [?geqrf](#);
 if $m < n$, array ***a*** contains the details of the *LQ* factorization of the matrix *A* as returned by [?gelqf](#).
- b*** Overwritten by the solution vectors, stored columnwise:
 If *trans* = 'N' and $m \geq n$, rows 1 to *n* of ***b*** contain the least squares solution vectors; the residual sum of squares for the solution in each column is given by the sum of squares of elements *n*+1 to *m* in that column;
 If *trans* = 'N' and $m < n$, rows 1 to *n* of ***b*** contain the minimum norm solution vectors;
 if *trans* = 'T' or 'C' and $m \geq n$, rows 1 to *m* of ***b*** contain the minimum norm solution vectors;
 if *trans* = 'T' or 'C' and $m < n$, rows 1 to *m* of ***b*** contain the least squares solution vectors; the residual sum of squares for the solution in each column is given by the sum of squares of elements *m*+1 to *n* in that column.
- work(1)*** If *info* = 0, on exit ***work(1)*** contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.
- info*** INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

For better performance, try using

$lwork = \min(m, n) + \max(1, m, n, nrhs) * blocksize$, where *blocksize* is a machine-dependent value (typically, 16 to 64) required for optimum performance of the *blocked algorithm*.

If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine ***work(1)*** and use this value for subsequent runs.

?gelsy

Computes the minimum-norm solution to a linear least squares problem using a complete orthogonal factorization of A.

```
call sgelsy ( m, n, nrhs, a, lda, b, ldb, jpvt, rcond, rank, work,
             lwork, info )
call dgelsy ( m, n, nrhs, a, lda, b, ldb, jpvt, rcond, rank, work,
             lwork, info )
call cgelsy ( m, n, nrhs, a, lda, b, ldb, jpvt, rcond, rank, work,
             lwork, rwork, info )
call zgelsy ( m, n, nrhs, a, lda, b, ldb, jpvt, rcond, rank, work,
             lwork, rwork, info )
```

Discussion

This routine computes the minimum-norm solution to a real/complex linear least squares problem:

$$\text{minimize } \| b - A x \|_2$$

using a complete orthogonal factorization of A . A is an m -by- n matrix which may be rank-deficient.

Several right hand side vectors b and solution vectors x can be handled in a single call; they are stored as the columns of the m -by- $nrhs$ right hand side matrix B and the n -by- $nrhs$ solution matrix X .

The routine first computes a QR factorization with column pivoting:

$$AP = Q \begin{pmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{pmatrix}$$

with R_{11} defined as the largest leading submatrix whose estimated condition number is less than $1/rcond$. The order of R_{11} , $rank$, is the effective rank of A .

Then, R_{22} is considered to be negligible, and R_{12} is annihilated by orthogonal/unitary transformations from the right, arriving at the complete orthogonal factorization:

$$AP = Q \begin{pmatrix} T_{11} & 0 \\ 0 & 0 \end{pmatrix} Z$$

The minimum-norm solution is then

$$x = PZ^H \begin{pmatrix} T_{11}^{-1} Q_1^H b \\ 0 \end{pmatrix}$$

where Q_1 consists of the first *rank* columns of Q . This routine is basically identical to the original `?gelsx` except three differences:

- The call to the subroutine `?geqpf` has been substituted by the call to the subroutine `?geqp3`. This subroutine is a BLAS-3 version of the QR factorization with column pivoting.
- Matrix B (the right hand side) is updated with BLAS-3.
- The permutation of matrix B (the right hand side) is faster and more simple.

Input Parameters

- m* **INTEGER**. The number of rows of the matrix A ($m \geq 0$).
- n* **INTEGER**. The number of columns of the matrix A ($n \geq 0$).
- nrhs* **INTEGER**. The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).
- a, b, work* **REAL** for `sgelsy`
DOUBLE PRECISION for `dgelsy`
COMPLEX for `cgelsy`
DOUBLE COMPLEX for `zgelsy`.
- Arrays:
a(*lda*,*) contains the m -by- n matrix A .
The second dimension of *a* must be at least $\max(1, n)$.

	<p>$b(ldb, *)$ contains the m-by-$nrhs$ right hand side matrix B.</p> <p>The second dimension of b must be at least $\max(1, nrhs)$.</p> <p>$work(lwork)$ is a workspace array.</p>
lda	INTEGER. The first dimension of a ; at least $\max(1, m)$.
ldb	INTEGER. The first dimension of b ; must be at least $\max(1, m, n)$.
$jpvt$	<p>INTEGER. Array, DIMENSION at least $\max(1, n)$.</p> <p>On entry, if $jpvt(i) \neq 0$, the ith column of A is permuted to the front of AP, otherwise the ith column of A is a free column.</p>
$rcond$	<p>REAL for single-precision flavors</p> <p>DOUBLE PRECISION for double-precision flavors.</p> <p>$rcond$ is used to determine the effective rank of A, which is defined as the order of the largest leading triangular submatrix R_{11} in the QR factorization with pivoting of A, whose estimated condition number $< 1/rcond$.</p>
$lwork$	INTEGER. The size of the $work$ array. See <i>Application notes</i> for the suggested value of $lwork$.
$rwork$	<p>REAL for <code>cgelsy</code></p> <p>DOUBLE PRECISION for <code>zgelsy</code>.</p> <p>Workspace array, DIMENSION at least $\max(1, 2n)$. Used in complex flavors only.</p>

Output Parameters

a	On exit, overwritten by the details of the complete orthogonal factorization of A .
b	Overwritten by the n -by- $nrhs$ solution matrix X .
$jpvt$	On exit, if $jpvt(i) = k$, then the i th column of AP was the k th column of A .

<i>rank</i>	INTEGER. The effective rank of A , that is, the order of the submatrix R_{11} . This is the same as the order of the submatrix T_{11} in the complete orthogonal factorization of A .
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = $-i$, the i th parameter had an illegal value.

Application Notes

For real flavors:

The unblocked strategy requires that:

$$lwork \geq \max(mn+3n+1, 2*mn + nrhs),$$

where $mn = \min(m, n)$.

The block algorithm requires that:

$$lwork \geq \max(mn+2n+nb*(n+1), 2*mn+nb*nrhs),$$

where nb is an upper bound on the blocksize returned by *ilaenv* for the routines *sgeqp3/dgeqp3*, *stzrzf/dtzrzf*, *stzrqf/dtzrqf*, *sormqr/dormqr*, and *sormrz/dormrz*.

For complex flavors:

The unblocked strategy requires that:

$$lwork \geq mn + \max(2*mn, n+1, mn + nrhs),$$

where $mn = \min(m, n)$.

The block algorithm requires that:

$$lwork \geq mn + \max(2*mn, nb*(n+1), mn+mn*nb, mn+nb*nrhs),$$

where nb is an upper bound on the blocksize returned by *ilaenv* for the routines *cgeqp3/zgeqp3*, *ctzrzf/ztzrzf*, *ctzrqf/ztzrqf*, *cunmqr/zunmqr*, and *cunmrz/zunmrz*.

?gelss

Computes the minimum-norm solution to a linear least squares problem using the singular value decomposition of A .

```
call sgelss ( m, n, nrhs, a, lda, b, ldb, s, rcond, rank, work,
             lwork, info )
call dgelss ( m, n, nrhs, a, lda, b, ldb, s, rcond, rank, work,
             lwork, info )
call cgelss ( m, n, nrhs, a, lda, b, ldb, s, rcond, rank, work,
             lwork, rwork, info )
call zgelss ( m, n, nrhs, a, lda, b, ldb, s, rcond, rank, work,
             lwork, rwork, info )
```

Discussion

This routine computes the minimum norm solution to a real linear least squares problem:

$$\text{minimize } \| b - A x \|_2$$

using the singular value decomposition (SVD) of A . A is an m -by- n matrix which may be rank-deficient.

Several right hand side vectors b and solution vectors x can be handled in a single call; they are stored as the columns of the m -by- $nrhs$ right hand side matrix B and the n -by- $nrhs$ solution matrix X .

The effective rank of A is determined by treating as zero those singular values which are less than $rcond$ times the largest singular value.

Input Parameters

m **INTEGER**. The number of rows of the matrix A ($m \geq 0$).

n **INTEGER**. The number of columns of the matrix A ($n \geq 0$).

nrhs **INTEGER**. The number of right-hand sides; the number of columns in B ($nrhs \geq 0$).

<i>a, b, work</i>	<p>REAL for <i>sgelss</i> DOUBLE PRECISION for <i>dgelss</i> COMPLEX for <i>cgelss</i> DOUBLE COMPLEX for <i>zgelss</i>.</p> <p>Arrays: <i>a(lda,*)</i> contains the <i>m</i>-by-<i>n</i> matrix <i>A</i>. The second dimension of <i>a</i> must be at least $\max(1, n)$. <i>b(ldb,*)</i> contains the <i>m</i>-by-<i>nrhs</i> right hand side matrix <i>B</i>. The second dimension of <i>b</i> must be at least $\max(1, nrhs)$. <i>work(lwork)</i> is a workspace array.</p>
<i>lda</i>	INTEGER. The first dimension of <i>a</i> ; at least $\max(1, m)$.
<i>ldb</i>	INTEGER. The first dimension of <i>b</i> ; must be at least $\max(1, m, n)$.
<i>rcond</i>	<p>REAL for single-precision flavors DOUBLE PRECISION for double-precision flavors.</p> <p><i>rcond</i> is used to determine the effective rank of <i>A</i>. Singular values $s(i) \leq rcond * s(1)$ are treated as zero. If <i>rcond</i> < 0, machine precision is used instead.</p>
<i>lwork</i>	INTEGER. The size of the <i>work</i> array; <i>lwork</i> ≥ 1. See <i>Application notes</i> for the suggested value of <i>lwork</i> .
<i>rwork</i>	<p>REAL for <i>cgelss</i> DOUBLE PRECISION for <i>zgelss</i>.</p> <p>Workspace array used in complex flavors only. DIMENSION at least $\max(1, 5 * \min(m, n))$.</p>

Output Parameters

<i>a</i>	On exit, the first $\min(m, n)$ rows of <i>A</i> are overwritten with its right singular vectors, stored row-wise.
<i>b</i>	<p>Overwritten by the <i>n</i>-by-<i>nrhs</i> solution matrix <i>X</i>.</p> <p>If $m \geq n$ and <i>rank</i> = <i>n</i>, the residual sum-of-squares for the solution in the <i>i</i>-th column is given by the sum of squares of elements <i>n</i>+1:<i>m</i> in that column.</p>

<i>s</i>	<p>REAL for single precision flavors</p> <p>DOUBLE PRECISION for double precision flavors.</p> <p>Array, DIMENSION at least $\max(1, \min(m, n))$. The singular values of <i>A</i> in decreasing order. The condition number of <i>A</i> in the 2-norm is</p> $k_2(A) = s(1) / s(\min(m, n)) .$
<i>rank</i>	<p>INTEGER.</p> <p>The effective rank of <i>A</i>, that is, the number of singular values which are greater than <i>rcond</i> * <i>s</i>(1).</p>
<i>work(1)</i>	<p>If <i>info</i> = 0, on exit, <i>work(1)</i> contains the minimum value of <i>lwork</i> required for optimum performance. Use this <i>lwork</i> for subsequent runs.</p>
<i>info</i>	<p>INTEGER.</p> <p>If <i>info</i> = 0, the execution is successful.</p> <p>If <i>info</i> = -<i>i</i>, the <i>i</i>th parameter had an illegal value.</p> <p>If <i>info</i> = <i>i</i>, then the algorithm for computing the SVD failed to converge; <i>i</i> indicates the number of off-diagonal elements of an intermediate bidiagonal form which did not converge to zero.</p>

Application Notes

For real flavors:

$$lwork \geq 3 * \min(m, n) + \max(2 * \min(m, n), \max(m, n), nrhs)$$

For complex flavors:

$$lwork \geq 2 * \min(m, n) + \max(m, n , nrhs)$$

For good performance, *lwork* should generally be larger. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

?gelsd

Computes the minimum-norm solution to a linear least squares problem using the singular value decomposition of A and a divide and conquer method.

```
call sgelsd ( m, n, nrhs, a, lda, b, ldb, s, rcond, rank, work,
             lwork, iwork, info )
call dgelsd ( m, n, nrhs, a, lda, b, ldb, s, rcond, rank, work,
             lwork, iwork, info )
call cgelsd ( m, n, nrhs, a, lda, b, ldb, s, rcond, rank, work,
             lwork, rwork, iwork, info )
call zgelsd ( m, n, nrhs, a, lda, b, ldb, s, rcond, rank, work,
             lwork, rwork, iwork, info )
```

Discussion

This routine computes the minimum-norm solution to a real linear least squares problem:

$$\text{minimize } \| b - A x \|_2$$

using the singular value decomposition (SVD) of A . A is an m -by- n matrix which may be rank-deficient.

Several right hand side vectors b and solution vectors x can be handled in a single call; they are stored as the columns of the m -by- $nrhs$ right hand side matrix B and the n -by- $nrhs$ solution matrix X .

The problem is solved in three steps:

1. Reduce the coefficient matrix A to bidiagonal form with Householder transformations, reducing the original problem into a "bidiagonal least squares problem" (BLS).
2. Solve the BLS using a divide and conquer approach.
3. Apply back all the Householder transformations to solve the original least squares problem.

The effective rank of A is determined by treating as zero those singular values which are less than $rcond$ times the largest singular value.

Input Parameters

<i>m</i>	INTEGER. The number of rows of the matrix <i>A</i> ($m \geq 0$).
<i>n</i>	INTEGER. The number of columns of the matrix <i>A</i> ($n \geq 0$).
<i>nrhs</i>	INTEGER. The number of right-hand sides; the number of columns in <i>B</i> ($nrhs \geq 0$).
<i>a, b, work</i>	REAL for <i>sgelsd</i> DOUBLE PRECISION for <i>dgelsd</i> COMPLEX for <i>cgelsd</i> DOUBLE COMPLEX for <i>zgelsd</i> . Arrays: <i>a(lda,*)</i> contains the <i>m</i> -by- <i>n</i> matrix <i>A</i> . The second dimension of <i>a</i> must be at least $\max(1, n)$. <i>b(ldb,*)</i> contains the <i>m</i> -by- <i>nrhs</i> right hand side matrix <i>B</i> . The second dimension of <i>b</i> must be at least $\max(1, nrhs)$. <i>work(lwork)</i> is a workspace array.
<i>lda</i>	INTEGER. The first dimension of <i>a</i> ; at least $\max(1, m)$.
<i>ldb</i>	INTEGER. The first dimension of <i>b</i> ; must be at least $\max(1, m, n)$.
<i>rcond</i>	REAL for single-precision flavors DOUBLE PRECISION for double-precision flavors. <i>rcond</i> is used to determine the effective rank of <i>A</i> . Singular values $s(i) \leq rcond * s(1)$ are treated as zero. If <i>rcond</i> < 0, machine precision is used instead.
<i>lwork</i>	INTEGER. The size of the <i>work</i> array; $lwork \geq 1$. See <i>Application notes</i> for the suggested value of <i>lwork</i> .
<i>iwork</i>	INTEGER. Workspace array. See <i>Application notes</i> for the suggested dimension of <i>iwork</i> .
<i>rwork</i>	REAL for <i>cgelsd</i> DOUBLE PRECISION for <i>zgelsd</i> . Workspace array, used in complex flavors only. See

Application notes for the suggested dimension of *rwork*.

Output Parameters

- a* On exit, *A* has been overwritten.
- b* Overwritten by the *n*-by-*nrhs* solution matrix *X*.
If $m \geq n$ and *rank* = *n*, the residual sum-of-squares for the solution in the *i*-th column is given by the sum of squares of elements *n*+1:*m* in that column.
- s* **REAL** for single precision flavors
DOUBLE PRECISION for double precision flavors.
Array, **DIMENSION** at least $\max(1, \min(m, n))$. The singular values of *A* in decreasing order. The condition number of *A* in the 2-norm is
$$k_2(A) = s(1) / s(\min(m, n)) .$$
- rank* **INTEGER**.
The effective rank of *A*, that is, the number of singular values which are greater than *rcond* * *s*(1).
- work*(1) If *info* = 0, on exit, *work*(1) contains the minimum value of *lwork* required for optimum performance. Use this *lwork* for subsequent runs.
- info* **INTEGER**.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i*, then the algorithm for computing the SVD failed to converge; *i* indicates the number of off-diagonal elements of an intermediate bidiagonal form which did not converge to zero.

Application Notes

The divide and conquer algorithm makes very mild assumptions about floating point arithmetic. It will work on machines with a guard digit in add/subtract. It could conceivably fail on hexadecimal or decimal machines without guard digits, but we know of none.

The exact minimum amount of workspace needed depends on m , n and $nrhs$. The size $lwork$ of the workspace array $work$ must be as given below.

For real flavors:

If $m \geq n$,

$$lwork \geq 12n + 2n*smlsiz + 8n*nlvl + n*nrhs + (smlsiz+1)^2;$$

If $m < n$,

$$lwork \geq 12m + 2m*smlsiz + 8m*nlvl + m*nrhs + (smlsiz+1)^2;$$

For complex flavors:

If $m \geq n$,

$$lwork \geq 2n + n*nrhs ;$$

If $m < n$,

$$lwork \geq 2m + m*nrhs ;$$

where $smlsiz$ is returned by `ilaenv` and is equal to the maximum size of the subproblems at the bottom of the computation tree (usually about 25), and $nlvl = \text{INT}(\log_2(\min(m, n)/(smlsiz+1))) + 1$.

For good performance, $lwork$ should generally be larger. If you are in doubt how much workspace to supply, use a generous value of $lwork$ for the first run. On exit, examine `work(1)` and use this value for subsequent runs.

The dimension of the workspace array $iwork$ must be at least $3 * \min(m, n) * nlvl + 11 * \min(m, n)$.

The dimension $lrwork$ of the workspace array $rwork$ (for complex flavors) must be at least:

If $m \geq n$,

$$lrwork \geq 10n + 2n*smlsiz + 8n*nlvl + 3*smlsiz*nrhs + (smlsiz+1)^2;$$

If $m < n$,

$$lrwork \geq 10m + 2m*smlsiz + 8m*nlvl + 3*smlsiz*nrhs + (smlsiz+1)^2.$$

Generalized LLS Problems

This section describes LAPACK driver routines used for solving generalized linear least-squares problems. [Table 5-9](#) lists routines described in more detail below.

Table 5-9 Driver Routines for Solving Generalized LLS Problems

Routine Name	Operation performed
?gglse	Solves the linear equality-constrained least squares problem using a generalized RQ factorization.
?ggglm	Solves a general Gauss-Markov linear model problem using a generalized QR factorization.

?gglse

Solves the linear equality-constrained least squares problem using a generalized RQ factorization.

```
call sggls ( m, n, p, a, lda, b, ldb, c, d, x, work, lwork, info )
call dggls ( m, n, p, a, lda, b, ldb, c, d, x, work, lwork, info )
call cggls ( m, n, p, a, lda, b, ldb, c, d, x, work, lwork, info )
call zggls ( m, n, p, a, lda, b, ldb, c, d, x, work, lwork, info )
```

Discussion

This routine solves the linear equality-constrained least squares (LSE) problem:

$$\text{minimize } \|c - Ax\|_2 \quad \text{subject to } Bx = d$$

where A is an m -by- n matrix, B is a p -by- n matrix, c is a given m -vector, and d is a given p -vector.

It is assumed that $p \leq n \leq m+p$, and

$$\text{rank}(B) = p \quad \text{and} \quad \text{rank} \begin{pmatrix} A \\ B \end{pmatrix} = n .$$

These conditions ensure that the LSE problem has a unique solution, which is obtained using a generalized RQ factorization of the matrices B and A .

Input Parameters

m **INTEGER**. The number of rows of the matrix A ($m \geq 0$).

n **INTEGER**. The number of columns of the matrices A and B ($n \geq 0$).

p **INTEGER**. The number of rows of the matrix B ($0 \leq p \leq n \leq m+p$).

$a, b, c, d, work$ **REAL** for `sgglse`
DOUBLE PRECISION for `dggls`
COMPLEX for `cggls`
DOUBLE COMPLEX for `zggls`.

Arrays:

$a(lda, *)$ contains the m -by- n matrix A .

The second dimension of a must be at least $\max(1, n)$.

$b(l db, *)$ contains the p -by- n matrix B .

The second dimension of b must be at least $\max(1, n)$.

$c(*)$, dimension at least $\max(1, m)$, contains the right hand side vector for the least squares part of the LSE problem.

$d(*)$, dimension at least $\max(1, p)$, contains the right hand side vector for the constrained equation.

$work(lwork)$ is a workspace array.

lda **INTEGER**. The first dimension of a ; at least $\max(1, m)$.

ldb **INTEGER**. The first dimension of b ; at least $\max(1, p)$.

$lwork$ **INTEGER**. The size of the $work$ array;
 $lwork \geq \max(1, m+n+p)$. See *Application notes* for the suggested value of $lwork$.

Output Parameters

<i>x</i>	REAL for <i>sgglse</i> DOUBLE PRECISION for <i>dgglse</i> COMPLEX for <i>cgglse</i> DOUBLE COMPLEX for <i>zgglse</i> . Array, DIMENSION at least $\max(1, n)$. On exit, contains the solution of the LSE problem.
<i>a, b, d</i>	On exit, these arrays are overwritten.
<i>c</i>	On exit, the residual sum-of-squares for the solution is given by the sum of squares of elements $n-p+1$ to m of vector <i>c</i> .
<i>work(1)</i>	If <i>info</i> = 0, on exit, <i>work(1)</i> contains the minimum value of <i>lwork</i> required for optimum performance. Use this <i>lwork</i> for subsequent runs.
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = $-i$, the <i>i</i> th parameter had an illegal value.

Application Notes

For optimum performance use

$$lwork \geq p + \min(m, n) + \max(m, n) * nb,$$

where *nb* is an upper bound for the optimal blocksizes for *?geqrf*,
?gerqf, *?ormqr/?unmqr* and *?ormrq/?unmrq*.

?ggglm

Solves a general Gauss-Markov linear model problem using a generalized QR factorization.

```
call sggglm ( n, m, p, a, lda, b, ldb, d, x, y, work, lwork, info )
call dggglm ( n, m, p, a, lda, b, ldb, d, x, y, work, lwork, info )
call cggglm ( n, m, p, a, lda, b, ldb, d, x, y, work, lwork, info )
call zggglm ( n, m, p, a, lda, b, ldb, d, x, y, work, lwork, info )
```

Discussion

This routine solves a general Gauss-Markov linear model (GLM) problem:

$$\text{minimize}_x \|y\|_2 \quad \text{subject to } d = Ax + By$$

where A is an n -by- m matrix, B is an n -by- p matrix, and d is a given n -vector.

It is assumed that $m \leq n \leq m+p$, and

$$\text{rank}(A) = m \quad \text{and} \quad \text{rank}(A \ B) = n .$$

Under these assumptions, the constrained equation is always consistent, and there is a unique solution x and a minimal 2-norm solution y , which is obtained using a generalized QR factorization of A and B .

In particular, if matrix B is square nonsingular, then the problem GLM is equivalent to the following weighted linear least squares problem

$$\text{minimize}_x \|B^{-1}(d - Ax)\|_2 .$$

Input Parameters

n **INTEGER.** The number of rows of the matrices A and B ($n \geq 0$).

m **INTEGER.** The number of columns in A ($m \geq 0$).

p **INTEGER.** The number of columns in B ($p \geq n - m$).

$a, b, d, work$ **REAL** for sggglm
 DOUBLE PRECISION for dggglm
 COMPLEX for cggglm
 DOUBLE COMPLEX for zggglm.

Arrays:

$a(lda, *)$ contains the n -by- m matrix A .

The second dimension of a must be at least $\max(1, m)$.

$b(ldb, *)$ contains the n -by- p matrix B .

The second dimension of b must be at least $\max(1, p)$.

$d(*)$, dimension at least $\max(1, n)$, contains the left hand side of the GLM equation.

$work(lwork)$ is a workspace array.

lda INTEGER. The first dimension of a ; at least $\max(1, n)$.

ldb INTEGER. The first dimension of b ; at least $\max(1, n)$.

$lwork$ INTEGER. The size of the $work$ array;
 $lwork \geq \max(1, n+m+p)$. See *Application notes* for the suggested value of $lwork$.

Output Parameters

x, y REAL for `sggglm`
 DOUBLE PRECISION for `dggglm`
 COMPLEX for `cggglm`
 DOUBLE COMPLEX for `zggglm`.
 Arrays $x(*)$, $y(*)$. DIMENSION at least $\max(1, m)$ for x
 and at least $\max(1, p)$ for y .
 On exit, x and y are the solutions of the GLM problem.

a, b, d On exit, these arrays are overwritten.

$work(1)$ If $info = 0$, on exit, $work(1)$ contains the minimum value of $lwork$ required for optimum performance.

$info$ INTEGER.
 If $info = 0$, the execution is successful.
 If $info = -i$, the i th parameter had an illegal value.

Application Notes

For optimum performance use

$$lwork \geq m + \min(n, p) + \max(n, p) * nb,$$

where nb is an upper bound for the optimal blocksizes for `?geqrf`,
`?gerqf`, `?ormqr`/`?unmqr` and `?ormrq`/`?unmrq`.

Symmetric Eigenproblems

This section describes LAPACK driver routines used for solving symmetric eigenvalue problems. See also [computational routines](#) that can be called to solve these problems.

[Table 5-10](#) lists routines described in more detail below.

Table 5-10 Driver Routines for Solving Symmetric Eigenproblems

Routine Name	Operation performed
?syev/?heev	Computes all eigenvalues and, optionally, eigenvectors of a real symmetric / Hermitian matrix.
?syevd/?heevd	Computes all eigenvalues and (optionally) all eigenvectors of a real symmetric / Hermitian matrix using divide and conquer algorithm.
?syevx/?heevx	Computes selected eigenvalues and, optionally, eigenvectors of a symmetric / Hermitian matrix.
?syevr/?heevr	Computes selected eigenvalues and, optionally, eigenvectors of a real symmetric / Hermitian matrix using the Relatively Robust Representations.
?spev/?hpev	Computes all eigenvalues and, optionally, eigenvectors of a real symmetric / Hermitian matrix in packed storage.
?spevd/?hpevd	Uses divide and conquer algorithm to compute all eigenvalues and (optionally) all eigenvectors of a real symmetric / Hermitian matrix held in packed storage.
?spevx/?hpevx	Computes selected eigenvalues and, optionally, eigenvectors of a real symmetric / Hermitian matrix in packed storage.
?sbev/?hbev	Computes all eigenvalues and, optionally, eigenvectors of a real symmetric / Hermitian band matrix.
?sbevd/?hbevd	Computes all eigenvalues and (optionally) all eigenvectors of a real symmetric / Hermitian band matrix using divide and conquer algorithm.
?sbevz/?hbevz	Computes selected eigenvalues and, optionally, eigenvectors of a real symmetric / Hermitian band matrix.
?stev	Computes all eigenvalues and, optionally, eigenvectors of a real symmetric tridiagonal matrix.
?stevd	Computes all eigenvalues and (optionally) all eigenvectors of a real symmetric tridiagonal matrix using divide and conquer algorithm.
?stevx	Computes selected eigenvalues and eigenvectors of a real symmetric tridiagonal matrix.
?stevr	Computes selected eigenvalues and, optionally, eigenvectors of a real symmetric tridiagonal matrix using the Relatively Robust Representations.

?syev

Computes all eigenvalues and, optionally, eigenvectors of a real symmetric matrix.

```
call ssyev ( jobz, uplo, n, a, lda, w, work, lwork, info )
call dsyev ( jobz, uplo, n, a, lda, w, work, lwork, info )
```

Discussion

This routine computes all eigenvalues and, optionally, eigenvectors of a real symmetric matrix A .

Input Parameters

jobz CHARACTER*1. Must be 'N' or 'V'.
 If *jobz* = 'N', then only eigenvalues are computed.
 If *jobz* = 'V', then eigenvalues and eigenvectors are computed.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', *a* stores the upper triangular part of A .
 If *uplo* = 'L', *a* stores the lower triangular part of A .

n INTEGER. The order of the matrix A ($n \geq 0$).

a, work REAL for *ssyev*
 DOUBLE PRECISION for *dsyev*
 Arrays:
a(*lda*,*) is an array containing either upper or lower triangular part of the symmetric matrix A , as specified by *uplo*.
 The second dimension of *a* must be at least $\max(1, n)$.
work(*lwork*) is a workspace array.

lda INTEGER. The first dimension of the array *a*.
 Must be at least $\max(1, n)$.

lwork **INTEGER**. The dimension of the array *work*.
 Constraint: $lwork \geq \max(1, 3n-1)$. See *Application notes* for the suggested value of *lwork*.

Output Parameters

a On exit, if *jobz* = 'V', then if *info* = 0, array *a* contains the orthonormal eigenvectors of the matrix A. If *jobz* = 'N', then on exit the lower triangle (if *uplo* = 'L') or the upper triangle (if *uplo* = 'U') of A, including the diagonal, is overwritten.

w **REAL** for *ssyev*
DOUBLE PRECISION for *dsyev*
 Array, **DIMENSION** at least $\max(1, n)$.
 If *info* = 0, contains the eigenvalues of the matrix A in ascending order.

work(1) On exit, if *lwork* > 0, then *work(1)* returns the required minimal size of *lwork*.

info **INTEGER**.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.
 If *info* = *i*, then the algorithm failed to converge; *i* indicates the number of elements of an intermediate tridiagonal form which did not converge to zero.

Application Notes

For optimum performance use

$$lwork \geq (nb+2)*n,$$

where *nb* is the blocksize for *?sytrd* returned by *ilaenv*.

If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

?heev

Computes all eigenvalues and, optionally, eigenvectors of a Hermitian matrix.

```
call cheev ( jobz, uplo, n, a, lda, w, work, lwork, rwork, info )
call zheev ( jobz, uplo, n, a, lda, w, work, lwork, rwork, info )
```

Discussion

This routine computes all eigenvalues and, optionally, eigenvectors of a complex Hermitian matrix A .

Input Parameters

jobz CHARACTER*1. Must be 'N' or 'V'.
 If *jobz* = 'N', then only eigenvalues are computed.
 If *jobz* = 'V', then eigenvalues and eigenvectors are computed.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', *a* stores the upper triangular part of A .
 If *uplo* = 'L', *a* stores the lower triangular part of A .

n INTEGER. The order of the matrix A ($n \geq 0$).

a, work COMPLEX for *cheev*
 DOUBLE COMPLEX for *zheev*
 Arrays:
a(*lda*,*) is an array containing either upper or lower triangular part of the Hermitian matrix A , as specified by *uplo*.
 The second dimension of *a* must be at least $\max(1, n)$.
work(*lwork*) is a workspace array.

lda INTEGER. The first dimension of the array *a*.
 Must be at least $\max(1, n)$.

<i>lwork</i>	INTEGER. The dimension of the array <i>work</i> . Constraint: $lwork \geq \max(1, 2n-1)$. See <i>Application notes</i> for the suggested value of <i>lwork</i> .
<i>rwork</i>	REAL for <i>cheev</i> DOUBLE PRECISION for <i>zheev</i> . Workspace array, DIMENSION at least $\max(1, 3n-2)$.

Output Parameters

<i>a</i>	On exit, if <i>jobz</i> = 'V', then if <i>info</i> = 0, array <i>a</i> contains the orthonormal eigenvectors of the matrix A. If <i>jobz</i> = 'N', then on exit the lower triangle (if <i>uplo</i> = 'L') or the upper triangle (if <i>uplo</i> = 'U') of A, including the diagonal, is overwritten.
<i>w</i>	REAL for <i>cheev</i> DOUBLE PRECISION for <i>zheev</i> Array, DIMENSION at least $\max(1, n)$. If <i>info</i> = 0, contains the eigenvalues of the matrix A in ascending order.
<i>work(1)</i>	On exit, if <i>lwork</i> > 0, then <i>work(1)</i> returns the required minimal size of <i>lwork</i> .
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = <i>i</i> , then the algorithm failed to converge; <i>i</i> indicates the number of elements of an intermediate tridiagonal form which did not converge to zero.

Application Notes

For optimum performance use

$$lwork \geq (nb+1)*n,$$

where *nb* is the blocksize for *?hetrd* returned by *ilaenv*.

If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

?syevd

Computes all eigenvalues and (optionally) all eigenvectors of a real symmetric matrix using divide and conquer algorithm.

```
call ssyevd (job,uplo,n,a,lda,w,work,lwork,iwork,liwork,info)
call dsyevd (job,uplo,n,a,lda,w,work,lwork,iwork,liwork,info)
```

Discussion

This routine computes all the eigenvalues, and optionally all the eigenvectors, of a real symmetric matrix A . In other words, it can compute the spectral factorization of A as: $A = Z\Lambda Z^T$.

Here Λ is a diagonal matrix whose diagonal elements are the eigenvalues λ_i , and Z is the orthogonal matrix whose columns are the eigenvectors z_i . Thus,

$$Az_i = \lambda_i z_i \text{ for } i = 1, 2, \dots, n.$$

If the eigenvectors are requested, then this routine uses a divide and conquer algorithm to compute eigenvalues and eigenvectors. However, if only eigenvalues are required, then it uses the Pal-Walker-Kahan variant of the QL or QR algorithm.

Input Parameters

job CHARACTER*1. Must be 'N' or 'V'.
 If *job* = 'N', then only eigenvalues are computed.
 If *job* = 'V', then eigenvalues and eigenvectors are computed.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', *a* stores the upper triangular part of A .
 If *uplo* = 'L', *a* stores the lower triangular part of A .

n INTEGER. The order of the matrix A ($n \geq 0$).

a REAL for *ssyevd*
 DOUBLE PRECISION for *dsyevd*
 Array, DIMENSION (*lda*, *).

$a(lda, *)$ is an array containing either upper or lower triangular part of the symmetric matrix A , as specified by $uplo$.
The second dimension of a must be at least $\max(1, n)$.

lda	INTEGER. The first dimension of the array a . Must be at least $\max(1, n)$.
$work$	REAL for <code>ssyevd</code> DOUBLE PRECISION for <code>dsyevd</code> . Workspace array, DIMENSION at least $lwork$.
$lwork$	INTEGER. The dimension of the array $work$. Constraints: if $n \leq 1$, then $lwork \geq 1$; if $job = 'N'$ and $n > 1$, then $lwork \geq 2n+1$; if $job = 'V'$ and $n > 1$, then $lwork \geq 3n^2 + (5+2k) * n + 1$, where k is the smallest integer which satisfies $2^k \geq n$.
$iwork$	INTEGER. Workspace array, DIMENSION at least $liwork$.
$liwork$	INTEGER. The dimension of the array $iwork$. Constraints: if $n \leq 1$, then $liwork \geq 1$; if $job = 'N'$ and $n > 1$, then $liwork \geq 1$; if $job = 'V'$ and $n > 1$, then $liwork \geq 5n+2$.

Output Parameters

w	REAL for <code>ssyevd</code> DOUBLE PRECISION for <code>dsyevd</code> Array, DIMENSION at least $\max(1, n)$. If $info = 0$, contains the eigenvalues of the matrix A in ascending order. See also $info$.
a	If $job = 'V'$, then on exit this array is overwritten by the orthogonal matrix Z which contains the eigenvectors of A .

work(1) On exit, if *lwork* > 0, then *work(1)* returns the required minimal size of *lwork*.

iwork(1) On exit, if *liwork* > 0, then *iwork(1)* returns the required minimal size of *liwork*.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = *i*, then the algorithm failed to converge; *i* indicates the number of elements of an intermediate tridiagonal form which did not converge to zero.
If *info* = *-i*, the *i*th parameter had an illegal value.

Application Notes

The computed eigenvalues and eigenvectors are exact for a matrix $T + E$ such that $\|E\|_2 = O(\epsilon) \|T\|_2$, where ϵ is the machine precision.

The complex analogue of this routine is [?heevd](#).

?heevd

Computes all eigenvalues and (optionally) all eigenvectors of a complex Hermitian matrix using divide and conquer algorithm.

```
call cheevd (job, uplo, n, a, lda, w, work, lwork, rwork, lrwork,
            iwork, liwork, info)
call zheevd (job, uplo, n, a, lda, w, work, lwork, rwork, lrwork,
            iwork, liwork, info)
```

Discussion

This routine computes all the eigenvalues, and optionally all the eigenvectors, of a complex Hermitian matrix A . In other words, it can compute the spectral factorization of A as: $A = Z\Lambda Z^H$.

Here Λ is a real diagonal matrix whose diagonal elements are the eigenvalues λ_i , and Z is the (complex) unitary matrix whose columns are the eigenvectors z_i . Thus,

$$Az_i = \lambda_i z_i \text{ for } i = 1, 2, \dots, n.$$

If the eigenvectors are requested, then this routine uses a divide and conquer algorithm to compute eigenvalues and eigenvectors. However, if only eigenvalues are required, then it uses the Pal-Walker-Kahan variant of the QL or QR algorithm.

Input Parameters

<i>job</i>	CHARACTER*1. Must be 'N' or 'V'. If <i>job</i> = 'N', then only eigenvalues are computed. If <i>job</i> = 'V', then eigenvalues and eigenvectors are computed.
<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. If <i>uplo</i> = 'U', <i>a</i> stores the upper triangular part of A . If <i>uplo</i> = 'L', <i>a</i> stores the lower triangular part of A .
<i>n</i>	INTEGER. The order of the matrix A ($n \geq 0$).

<i>a</i>	<p>COMPLEX for <i>cheevd</i> DOUBLE COMPLEX for <i>zheevd</i> Array, DIMENSION (<i>lda</i>, *). <i>a</i>(<i>lda</i>, *) is an array containing either upper or lower triangular part of the Hermitian matrix <i>A</i>, as specified by <i>uplo</i>. The second dimension of <i>a</i> must be at least $\max(1, n)$.</p>
<i>lda</i>	<p>INTEGER. The first dimension of the array <i>a</i>. Must be at least $\max(1, n)$.</p>
<i>work</i>	<p>COMPLEX for <i>cheevd</i> DOUBLE COMPLEX for <i>zheevd</i>. Workspace array, DIMENSION at least <i>lwork</i>.</p>
<i>lwork</i>	<p>INTEGER. The dimension of the array <i>work</i>. Constraints: if $n \leq 1$, then $lwork \geq 1$; if <i>job</i> = 'N' and $n > 1$, then $lwork \geq n+1$; if <i>job</i> = 'V' and $n > 1$, then $lwork \geq n^2+2n$</p>
<i>rwork</i>	<p>REAL for <i>cheevd</i> DOUBLE PRECISION for <i>zheevd</i> Workspace array, DIMENSION at least <i>lrwork</i>.</p>
<i>lrwork</i>	<p>INTEGER. The dimension of the array <i>rwork</i>. Constraints: if $n \leq 1$, then $lrwork \geq 1$; if <i>job</i> = 'N' and $n > 1$, then $lrwork \geq n$; if <i>job</i> = 'V' and $n > 1$, then $lrwork \geq 3n^2 + (4+2k) * n + 1$, where <i>k</i> is the smallest integer which satisfies $2^k \geq n$.</p>
<i>iwork</i>	<p>INTEGER. Workspace array, DIMENSION at least <i>liwork</i>.</p>
<i>liwork</i>	<p>INTEGER. The dimension of the array <i>iwork</i>. Constraints: if $n \leq 1$, then $liwork \geq 1$; if <i>job</i> = 'N' and $n > 1$, then $liwork \geq 1$; if <i>job</i> = 'V' and $n > 1$, then $liwork \geq 5n+2$.</p>

Output Parameters

<i>w</i>	<p>REAL for <code>cheevd</code> DOUBLE PRECISION for <code>zheevd</code> Array, DIMENSION at least $\max(1, n)$. If <i>info</i> = 0, contains the eigenvalues of the matrix <i>A</i> in ascending order. See also <i>info</i>.</p>
<i>a</i>	<p>If <i>job</i> = 'V', then on exit this array is overwritten by the unitary matrix <i>Z</i> which contains the eigenvectors of <i>A</i>.</p>
<i>work(1)</i>	<p>On exit, if <i>lwork</i> > 0, then the real part of <i>work(1)</i> returns the required minimal size of <i>lwork</i>.</p>
<i>rwork(1)</i>	<p>On exit, if <i>lrwork</i> > 0, then <i>rwork(1)</i> returns the required minimal size of <i>lrwork</i>.</p>
<i>iwork(1)</i>	<p>On exit, if <i>liwork</i> > 0, then <i>iwork(1)</i> returns the required minimal size of <i>liwork</i>.</p>
<i>info</i>	<p>INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = <i>i</i>, then the algorithm failed to converge; <i>i</i> indicates the number of elements of an intermediate tridiagonal form which did not converge to zero. If <i>info</i> = -<i>i</i>, the <i>i</i>th parameter had an illegal value.</p>

Application Notes

The computed eigenvalues and eigenvectors are exact for a matrix $A + E$ such that $\|E\|_2 = O(\epsilon) \|A\|_2$, where ϵ is the machine precision.

The real analogue of this routine is [?syevd](#).

See also [?hpevd](#) for matrices held in packed storage, and [?hbevd](#) for banded matrices.

?syevx

Computes selected eigenvalues and, optionally, eigenvectors of a symmetric matrix.

```
call ssyevx (jobz, range, uplo, n, a, lda, vl, vu, il, iu, abstol,
            m, w, z, ldz, work, lwork, iwork, ifail, info)
call dsyevx (jobz, range, uplo, n, a, lda, vl, vu, il, iu, abstol,
            m, w, z, ldz, work, lwork, iwork, ifail, info)
```

Discussion

This routine computes selected eigenvalues and, optionally, eigenvectors of a real symmetric matrix A . Eigenvalues and eigenvectors can be selected by specifying either a range of values or a range of indices for the desired eigenvalues.

Input Parameters

jobz CHARACTER*1. Must be 'N' or 'V'.
If *jobz* = 'N', then only eigenvalues are computed.
If *jobz* = 'V', then eigenvalues and eigenvectors are computed.

range CHARACTER*1. Must be 'A', 'V', or 'I'.
If *range* = 'A', all eigenvalues will be found.
If *range* = 'V', all eigenvalues in the half-open interval $(vl, vu]$ will be found.
If *range* = 'I', the eigenvalues with indices *il* through *iu* will be found.

uplo CHARACTER*1. Must be 'U' or 'L'.
If *uplo* = 'U', *a* stores the upper triangular part of A .
If *uplo* = 'L', *a* stores the lower triangular part of A .

n INTEGER. The order of the matrix A ($n \geq 0$).

<i>a, work</i>	<p>REAL for <i>ssyevx</i> DOUBLE PRECISION for <i>dsyevx</i>. Arrays: <i>a(lda,*)</i> is an array containing either upper or lower triangular part of the symmetric matrix <i>A</i>, as specified by <i>uplo</i>. The second dimension of <i>a</i> must be at least $\max(1, n)$. <i>work(lwork)</i> is a workspace array.</p>
<i>lda</i>	<p>INTEGER. The first dimension of the array <i>a</i>. Must be at least $\max(1, n)$.</p>
<i>vl, vu</i>	<p>REAL for <i>ssyevx</i> DOUBLE PRECISION for <i>dsyevx</i>. If <i>range</i> = 'V', the lower and upper bounds of the interval to be searched for eigenvalues; $vl \leq vu$. Not referenced if <i>range</i> = 'A' or 'I'.</p>
<i>il, iu</i>	<p>INTEGER. If <i>range</i> = 'I', the indices of the smallest and largest eigenvalues to be returned. Constraints: $1 \leq il \leq iu \leq n$, if $n > 0$; $il = 1$ and $iu = 0$, if $n = 0$. Not referenced if <i>range</i> = 'A' or 'V'.</p>
<i>abstol</i>	<p>REAL for <i>ssyevx</i> DOUBLE PRECISION for <i>dsyevx</i>. The absolute error tolerance for the eigenvalues. See <i>Application notes</i> for more information.</p>
<i>ldz</i>	<p>INTEGER. The first dimension of the output array <i>z</i>; $ldz \geq 1$. If <i>jobz</i> = 'V', then $ldz \geq \max(1, n)$.</p>
<i>lwork</i>	<p>INTEGER. The dimension of the array <i>work</i>. Constraint: $lwork \geq \max(1, 8n)$. See <i>Application notes</i> for the suggested value of <i>lwork</i>.</p>
<i>iwork</i>	<p>INTEGER. Workspace array, DIMENSION at least $\max(1, 5n)$.</p>

Output Parameters

- a* On exit, the lower triangle (if *uplo* = 'L') or the upper triangle (if *uplo* = 'U') of *A*, including the diagonal, is overwritten.
- m* **INTEGER**. The total number of eigenvalues found; $0 \leq m \leq n$. If *range* = 'A', $m = n$, and if *range* = 'I', $m = iu - il + 1$.
- w* **REAL** for *ssyevx*
DOUBLE PRECISION for *dsyevx*
 Array, **DIMENSION** at least $\max(1, n)$.
 The first *m* elements contain the selected eigenvalues of the matrix *A* in ascending order.
- z* **REAL** for *ssyevx*
DOUBLE PRECISION for *dsyevx*.
 Array *z*(*ldz*, *) contains eigenvectors.
 The second dimension of *z* must be at least $\max(1, m)$.
 If *jobz* = 'V', then if *info* = 0, the first *m* columns of *z* contain the orthonormal eigenvectors of the matrix *A* corresponding to the selected eigenvalues, with the *i*-th column of *z* holding the eigenvector associated with *w*(*i*). If an eigenvector fails to converge, then that column of *z* contains the latest approximation to the eigenvector, and the index of the eigenvector is returned in *ifail*.
 If *jobz* = 'N', then *z* is not referenced.
 Note: you must ensure that at least $\max(1, m)$ columns are supplied in the array *z*; if *range* = 'V', the exact value of *m* is not known in advance and an upper bound must be used.
- work(1)* On exit, if *lwork* > 0, then *work(1)* returns the required minimal size of *lwork*.
- ifail* **INTEGER**. Array, **DIMENSION** at least $\max(1, n)$.
 If *jobz* = 'V', then if *info* = 0, the first *m* elements of *ifail* are zero; if *info* > 0, then *ifail* contains the indices of the eigenvectors that failed to converge.
 If *jobz* = 'V', then *ifail* is not referenced.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

If *info* = *i*, then *i* eigenvectors failed to converge; their indices are stored in the array *ifail*.

Application Notes

For optimum performance use $lwork \geq (nb+3)*n$, where *nb* is the maximum of the blocksize for *?sytrd* and *?ormtr* returned by *ilaenv*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

An approximate eigenvalue is accepted as converged when it is determined to lie in an interval [a,b] of width less than or equal to $abstol + \epsilon * \max(|a|,|b|)$, where ϵ is the machine precision. If *abstol* is less than or equal to zero, then $\epsilon * |T|$ will be used in its place, where $|T|$ is the 1-norm of the tridiagonal matrix obtained by reducing *A* to tridiagonal form.

Eigenvalues will be computed most accurately when *abstol* is set to twice the underflow threshold $2 * slamch('S')$, not zero. If this routine returns with *info* > 0, indicating that some eigenvectors did not converge, try setting *abstol* to $2 * slamch('S')$.

?heevx

Computes selected eigenvalues and, optionally, eigenvectors of a Hermitian matrix.

```
call cheevx (jobz, range, uplo, n, a, lda, vl, vu, il, iu, abstol,
            m, w, z, ldz, work, lwork, rwork, iwork, ifail, info)
call zheevx (jobz, range, uplo, n, a, lda, vl, vu, il, iu, abstol,
            m, w, z, ldz, work, lwork, rwork, iwork, ifail, info)
```

Discussion

This routine computes selected eigenvalues and, optionally, eigenvectors of a complex Hermitian matrix A . Eigenvalues and eigenvectors can be selected by specifying either a range of values or a range of indices for the desired eigenvalues.

Input Parameters

jobz CHARACTER*1. Must be 'N' or 'V'.
If *jobz* = 'N', then only eigenvalues are computed.
If *jobz* = 'V', then eigenvalues and eigenvectors are computed.

range CHARACTER*1. Must be 'A', 'V', or 'I'.
If *range* = 'A', all eigenvalues will be found.
If *range* = 'V', all eigenvalues in the half-open interval $(vl, vu]$ will be found.
If *range* = 'I', the eigenvalues with indices *il* through *iu* will be found.

uplo CHARACTER*1. Must be 'U' or 'L'.
If *uplo* = 'U', *a* stores the upper triangular part of A .
If *uplo* = 'L', *a* stores the lower triangular part of A .

n INTEGER. The order of the matrix A ($n \geq 0$).

<i>a, work</i>	<p>COMPLEX for <i>cheevx</i> DOUBLE COMPLEX for <i>zheevx</i>. Arrays: <i>a(lda,*)</i> is an array containing either upper or lower triangular part of the Hermitian matrix <i>A</i>, as specified by <i>uplo</i>. The second dimension of <i>a</i> must be at least $\max(1, n)$. <i>work(lwork)</i> is a workspace array.</p>
<i>lda</i>	<p>INTEGER. The first dimension of the array <i>a</i>. Must be at least $\max(1, n)$.</p>
<i>vl, vu</i>	<p>REAL for <i>cheevx</i> DOUBLE PRECISION for <i>zheevx</i>. If <i>range</i> = 'V', the lower and upper bounds of the interval to be searched for eigenvalues; $vl \leq vu$. Not referenced if <i>range</i> = 'A' or 'I'.</p>
<i>il, iu</i>	<p>INTEGER. If <i>range</i> = 'I', the indices of the smallest and largest eigenvalues to be returned. Constraints: $1 \leq il \leq iu \leq n$, if $n > 0$; $il = 1$ and $iu = 0$, if $n = 0$. Not referenced if <i>range</i> = 'A' or 'V'.</p>
<i>abstol</i>	<p>REAL for <i>cheevx</i> DOUBLE PRECISION for <i>zheevx</i>. The absolute error tolerance for the eigenvalues. See <i>Application notes</i> for more information.</p>
<i>ldz</i>	<p>INTEGER. The first dimension of the output array <i>z</i>; $ldz \geq 1$. If <i>jobz</i> = 'V', then $ldz \geq \max(1, n)$.</p>
<i>lwork</i>	<p>INTEGER. The dimension of the array <i>work</i>. Constraint: $lwork \geq \max(1, 2n-1)$. See <i>Application notes</i> for the suggested value of <i>lwork</i>.</p>
<i>rwork</i>	<p>REAL for <i>cheevx</i> DOUBLE PRECISION for <i>zheevx</i>. Workspace array, DIMENSION at least $\max(1, 7n)$.</p>
<i>iwork</i>	<p>INTEGER. Workspace array, DIMENSION at least $\max(1, 5n)$.</p>

Output Parameters

- a* On exit, the lower triangle (if *uplo* = 'L') or the upper triangle (if *uplo* = 'U') of *A*, including the diagonal, is overwritten.
- m* **INTEGER**. The total number of eigenvalues found; $0 \leq m \leq n$. If *range* = 'A', $m = n$, and if *range* = 'I', $m = iu - il + 1$.
- w* **REAL** for *cheevx*
DOUBLE PRECISION for *zheevx*
 Array, **DIMENSION** at least $\max(1, n)$.
 The first *m* elements contain the selected eigenvalues of the matrix *A* in ascending order.
- z* **COMPLEX** for *cheevx*
DOUBLE COMPLEX for *zheevx*.
 Array *z*(*ldz*, *) contains eigenvectors.
 The second dimension of *z* must be at least $\max(1, m)$.
 If *jobz* = 'V', then if *info* = 0, the first *m* columns of *z* contain the orthonormal eigenvectors of the matrix *A* corresponding to the selected eigenvalues, with the *i*-th column of *z* holding the eigenvector associated with *w*(*i*). If an eigenvector fails to converge, then that column of *z* contains the latest approximation to the eigenvector, and the index of the eigenvector is returned in *ifail*.
 If *jobz* = 'N', then *z* is not referenced.
 Note: you must ensure that at least $\max(1, m)$ columns are supplied in the array *z*; if *range* = 'V', the exact value of *m* is not known in advance and an upper bound must be used.
- work(1)* On exit, if *lwork* > 0, then *work(1)* returns the required minimal size of *lwork*.
- ifail* **INTEGER**. Array, **DIMENSION** at least $\max(1, n)$.
 If *jobz* = 'V', then if *info* = 0, the first *m* elements of *ifail* are zero; if *info* > 0, then *ifail* contains the indices of the eigenvectors that failed to converge.
 If *jobz* = 'V', then *ifail* is not referenced.

info

INTEGER.

If *info* = 0, the execution is successful.If *info* = -*i*, the *i*th parameter had an illegal value.If *info* = *i*, then *i* eigenvectors failed to converge; their indices are stored in the array *ifail*.

Application Notes

For optimum performance use $lwork \geq (nb+1)*n$, where *nb* is the maximum of the blocksize for `?hetrd` and `?unmtr` returned by `ilaenv`. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine `work(1)` and use this value for subsequent runs.

An approximate eigenvalue is accepted as converged when it is determined to lie in an interval [a,b] of width less than or equal to $abstol + \epsilon * \max(|a|,|b|)$, where ϵ is the machine precision. If *abstol* is less than or equal to zero, then $\epsilon * |T|$ will be used in its place, where $|T|$ is the 1-norm of the tridiagonal matrix obtained by reducing *A* to tridiagonal form.

Eigenvalues will be computed most accurately when *abstol* is set to twice the underflow threshold $2 * slamch('S')$, not zero. If this routine returns with *info* > 0, indicating that some eigenvectors did not converge, try setting *abstol* to $2 * slamch('S')$.

?syevr

Computes selected eigenvalues and, optionally, eigenvectors of a real symmetric matrix using the Relatively Robust Representations.

```
call ssyevr (jobz, range, uplo, n, a, lda, vl, vu, il, iu, abstol,
            m, w, z, ldz, isuppz, work, lwork, iwork, liwork, info)
call dsyevr (jobz, range, uplo, n, a, lda, vl, vu, il, iu, abstol,
            m, w, z, ldz, isuppz, work, lwork, iwork, liwork, info)
```

Discussion

This routine computes selected eigenvalues and, optionally, eigenvectors of a real symmetric matrix T . Eigenvalues and eigenvectors can be selected by specifying either a range of values or a range of indices for the desired eigenvalues.

Whenever possible, `?syevr` calls `sstegr/dstegr` to compute the eigenspectrum using Relatively Robust Representations. `?stegr` computes eigenvalues by the *dqds* algorithm, while orthogonal eigenvectors are computed from various “good” LDL^T representations (also known as Relatively Robust Representations). Gram-Schmidt orthogonalization is avoided as far as possible. More specifically, the various steps of the algorithm are as follows. For the i -th unreduced block of T ,

- (a) Compute $T - \mathfrak{Q}_i = L_i D_i L_i^T$, such that $L_i D_i L_i^T$ is a relatively robust representation;
- (b) Compute the eigenvalues, λ_j , of $L_i D_i L_i^T$ to high relative accuracy by the *dqds* algorithm;
- (c) If there is a cluster of close eigenvalues, “choose” \mathfrak{Q}_i close to the cluster, and go to step (a);
- (d) Given the approximate eigenvalue λ_j of $L_i D_i L_i^T$, compute the corresponding eigenvector by forming a rank-revealing twisted factorization.

The desired accuracy of the output can be specified by the input parameter *abstol*.

The routine `?syevr` calls `sstegr/dstegr` when the full spectrum is requested on machines which conform to the IEEE-754 floating point standard. `?syevr` calls `sstebz/dstebz` and `sstein/dstein` on non-IEEE machines and when partial spectrum requests are made.

Input Parameters

<i>jobz</i>	<p>CHARACTER*1. Must be 'N' or 'V'.</p> <p>If <i>jobz</i> = 'N', then only eigenvalues are computed.</p> <p>If <i>jobz</i> = 'V', then eigenvalues and eigenvectors are computed.</p>
<i>range</i>	<p>CHARACTER*1. Must be 'A' or 'V' or 'I'.</p> <p>If <i>range</i> = 'A', the routine computes all eigenvalues.</p> <p>If <i>range</i> = 'V', the routine computes eigenvalues λ_i in the half-open interval: $vl < \lambda_i \leq vu$.</p> <p>If <i>range</i> = 'I', the routine computes eigenvalues with indices <i>il</i> to <i>iu</i>.</p> <p>For <i>range</i> = 'V' or 'I' and $iu - il < n - 1$, <code>sstebz/dstebz</code> and <code>sstein/dstein</code> are called.</p>
<i>n</i>	INTEGER. The order of the matrix <i>A</i> ($n \geq 0$).
<i>a, work</i>	<p>REAL for <code>ssyevr</code></p> <p>DOUBLE PRECISION for <code>dsyevr</code>.</p> <p>Arrays:</p> <p><i>a</i>(<i>lda</i>,*) is an array containing either upper or lower triangular part of the symmetric matrix <i>A</i>, as specified by <i>uplo</i>.</p> <p>The second dimension of <i>a</i> must be at least $\max(1, n)$.</p> <p><i>work</i>(<i>lwork</i>) is a workspace array.</p>
<i>lda</i>	<p>INTEGER. The first dimension of the array <i>a</i>.</p> <p>Must be at least $\max(1, n)$.</p>

<i>vl, vu</i>	<p>REAL for <i>ssyevr</i> DOUBLE PRECISION for <i>dsyevr</i>.</p> <p>If <i>range</i> = 'V', the lower and upper bounds of the interval to be searched for eigenvalues. Constraint: $vl < vu$.</p> <p>If <i>range</i> = 'A' or 'I', <i>vl</i> and <i>vu</i> are not referenced.</p>
<i>il, iu</i>	<p>INTEGER.</p> <p>If <i>range</i> = 'I', the indices in ascending order of the smallest and largest eigenvalues to be returned. Constraint: $1 \leq il \leq iu \leq n$, if $n > 0$; $il=1$ and $iu=0$ if $n = 0$.</p> <p>If <i>range</i> = 'A' or 'V', <i>il</i> and <i>iu</i> are not referenced.</p>
<i>abstol</i>	<p>REAL for <i>ssyevr</i> DOUBLE PRECISION for <i>dsyevr</i>.</p> <p>The absolute error tolerance to which each eigenvalue/eigenvector is required.</p> <p>If <i>jobz</i> = 'V', the eigenvalues and eigenvectors output have residual norms bounded by <i>abstol</i>, and the dot products between different eigenvectors are bounded by <i>abstol</i>. If $abstol < n\epsilon$, then $n\epsilon$ will be used in its place, where ϵ is the machine precision. The eigenvalues are computed to an accuracy of ϵ irrespective of <i>abstol</i>. If high relative accuracy is important, set <i>abstol</i> to <i>?lamch</i>('S').</p>
<i>ldz</i>	<p>INTEGER. The leading dimension of the output array <i>z</i>. Constraints: $ldz \geq 1$ if <i>jobz</i> = 'N'; $ldz \geq \max(1, n)$ if <i>jobz</i> = 'V'.</p>
<i>lwork</i>	<p>INTEGER. The dimension of the array <i>work</i>. Constraint: $lwork \geq \max(1, 26n)$. See <i>Application notes</i> for the suggested value of <i>lwork</i>.</p>
<i>iwork</i>	<p>INTEGER. Workspace array, DIMENSION (<i>liwork</i>).</p>
<i>liwork</i>	<p>INTEGER. The dimension of the array <i>iwork</i>, $lwork \geq \max(1, 10n)$.</p>

Output Parameters

<i>a</i>	On exit, the lower triangle (if <i>uplo</i> = 'L') or the upper triangle (if <i>uplo</i> = 'U') of A, including the diagonal, is overwritten.
<i>m</i>	INTEGER . The total number of eigenvalues found, $0 \leq m \leq n$. If <i>range</i> = 'A', $m = n$, and if <i>range</i> = 'I', $m = iu - il + 1$.
<i>w, z</i>	REAL for <i>ssyevr</i> DOUBLE PRECISION for <i>dsyevr</i> . Arrays: <i>w</i> (*), DIMENSION at least $\max(1, n)$, contains the selected eigenvalues in ascending order, stored in <i>w</i> (1) to <i>w</i> (<i>m</i>); <i>z</i> (<i>ldz</i> , *), the second dimension of <i>z</i> must be at least $\max(1, m)$. If <i>jobz</i> = 'V', then if <i>info</i> = 0, the first <i>m</i> columns of <i>z</i> contain the orthonormal eigenvectors of the matrix <i>T</i> corresponding to the selected eigenvalues, with the <i>i</i> -th column of <i>z</i> holding the eigenvector associated with <i>w</i> (<i>i</i>). If <i>jobz</i> = 'N', then <i>z</i> is not referenced. Note: you must ensure that at least $\max(1, m)$ columns are supplied in the array <i>z</i> ; if <i>range</i> = 'V', the exact value of <i>m</i> is not known in advance and an upper bound must be used.
<i>isuppz</i>	INTEGER . Array, DIMENSION at least $2 * \max(1, m)$. The support of the eigenvectors in <i>z</i> , i.e., the indices indicating the nonzero elements in <i>z</i> . The <i>i</i> -th eigenvector is nonzero only in elements <i>isuppz</i> (2 <i>i</i> -1) through <i>isuppz</i> (2 <i>i</i>). Implemented only for <i>range</i> = 'A' or 'I' and $iu - il = n - 1$.
<i>work</i> (1)	On exit, if <i>info</i> = 0, then <i>work</i> (1) returns the required minimal size of <i>lwork</i> .

lwork(1) On exit, if *info* = 0, then *lwork(1)* returns the required minimal size of *liwork*.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i*, an internal error has occurred.

Application Notes

For optimum performance use $lwork \geq (nb+6)*n$, where *nb* is the maximum of the blocksize for *?sytrd* and *?ormtr* returned by *ilaenv*. If you are in doubt how much workspace to supply, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

Normal execution of *?stegr* may create NaNs and infinities and hence may abort due to a floating point exception in environments which do not handle NaNs and infinities in the IEEE standard default manner.

?heevr

Computes selected eigenvalues and, optionally, eigenvectors of a Hermitian matrix using the Relatively Robust Representations.

```
call cheevr ( jobz, range, uplo, n, a, lda, vl, vu, il, iu, abstol,
             m, w, z, ldz, isuppz, work, lwork, rwork, lrwork,
             iwork, liwork, info)
call zheevr ( jobz, range, uplo, n, a, lda, vl, vu, il, iu, abstol,
             m, w, z, ldz, isuppz, work, lwork, rwork, lrwork,
             iwork, liwork, info)
```

Discussion

This routine computes selected eigenvalues and, optionally, eigenvectors of a complex Hermitian matrix T . Eigenvalues and eigenvectors can be selected by specifying either a range of values or a range of indices for the desired eigenvalues.

Whenever possible, `?heevr` calls `cstegr/zstegr` to compute the eigenspectrum using Relatively Robust Representations. `?stegr` computes eigenvalues by the *dqds* algorithm, while orthogonal eigenvectors are computed from various “good” LDL^T representations (also known as Relatively Robust Representations). Gram-Schmidt orthogonalization is avoided as far as possible. More specifically, the various steps of the algorithm are as follows. For the i -th unreduced block of T ,

- (a) Compute $T - \mathfrak{Q} = L_i D_i L_i^T$, such that $L_i D_i L_i^T$ is a relatively robust representation;
- (b) Compute the eigenvalues, λ_j , of $L_i D_i L_i^T$ to high relative accuracy by the *dqds* algorithm;
- (c) If there is a cluster of close eigenvalues, “choose” \mathfrak{Q} close to the cluster, and go to step (a);
- (d) Given the approximate eigenvalue λ_j of $L_i D_i L_i^T$, compute the corresponding eigenvector by forming a rank-revealing twisted factorization.

The desired accuracy of the output can be specified by the input parameter *abstol*.

The routine `?heevr` calls `cstegr/zstegr` when the full spectrum is requested on machines which conform to the IEEE-754 floating point standard. `?heevr` calls `sstebz/dstebz` and `cstein/zstein` on non-IEEE machines and when partial spectrum requests are made.

Input Parameters

jobz CHARACTER*1. Must be 'N' or 'V'.
 If *jobz* = 'N', then only eigenvalues are computed.
 If *jobz* = 'V', then eigenvalues and eigenvectors are computed.

range CHARACTER*1. Must be 'A' or 'V' or 'I'.
 If *range* = 'A', the routine computes all eigenvalues.
 If *range* = 'V', the routine computes eigenvalues λ_i in the half-open interval: $v_l < \lambda_i \leq v_u$.
 If *range* = 'I', the routine computes eigenvalues with indices *il* to *iu*.
 For *range* = 'V' or 'I', `sstebz/dstebz` and `cstein/zstein` are called.

n INTEGER. The order of the matrix *A* ($n \geq 0$).

a, work COMPLEX for `cheevr`
 DOUBLE COMPLEX for `zheevr`.
 Arrays:
a(*lda*,*) is an array containing either upper or lower triangular part of the Hermitian matrix *A*, as specified by *uplo*.
 The second dimension of *a* must be at least $\max(1, n)$.
work(*lwork*) is a workspace array.

lda INTEGER. The first dimension of the array *a*.
 Must be at least $\max(1, n)$.

<i>vl, vu</i>	<p>REAL for <code>cheevr</code> DOUBLE PRECISION for <code>zheevr</code>.</p> <p>If <i>range</i> = 'V', the lower and upper bounds of the interval to be searched for eigenvalues. Constraint: <i>vl</i> < <i>vu</i>.</p> <p>If <i>range</i> = 'A' or 'I', <i>vl</i> and <i>vu</i> are not referenced.</p>
<i>il, iu</i>	<p>INTEGER.</p> <p>If <i>range</i> = 'I', the indices in ascending order of the smallest and largest eigenvalues to be returned. Constraint: $1 \leq il \leq iu \leq n$, if $n > 0$; <i>il</i>=1 and <i>iu</i>=0 if $n = 0$.</p> <p>If <i>range</i> = 'A' or 'V', <i>il</i> and <i>iu</i> are not referenced.</p>
<i>abstol</i>	<p>REAL for <code>cheevr</code> DOUBLE PRECISION for <code>zheevr</code>.</p> <p>The absolute error tolerance to which each eigenvalue/eigenvector is required.</p> <p>If <i>jobz</i> = 'V', the eigenvalues and eigenvectors output have residual norms bounded by <i>abstol</i>, and the dot products between different eigenvectors are bounded by <i>abstol</i>. If $abstol < n\epsilon$, then $n\epsilon$ will be used in its place, where ϵ is the machine precision. The eigenvalues are computed to an accuracy of ϵ irrespective of <i>abstol</i>. If high relative accuracy is important, set <i>abstol</i> to <code>?lamch('S')</code>.</p>
<i>ldz</i>	<p>INTEGER. The leading dimension of the output array <i>z</i>. Constraints: <i>ldz</i> ≥ 1 if <i>jobz</i> = 'N'; <i>ldz</i> $\geq \max(1, n)$ if <i>jobz</i> = 'V'.</p>
<i>lwork</i>	<p>INTEGER. The dimension of the array <i>work</i>. Constraint: <i>lwork</i> $\geq \max(1, 2n)$. See <i>Application notes</i> for the suggested value of <i>lwork</i>.</p>
<i>rwork</i>	<p>REAL for <code>cheevr</code> DOUBLE PRECISION for <code>zheevr</code>.</p> <p>Workspace array, DIMENSION (<i>lrwork</i>).</p>

lwork **INTEGER**. The dimension of the array *rwork*;
lwork $\geq \max(1, 24n)$. .

iwork **INTEGER**.
 Workspace array, **DIMENSION** (*liwork*).

liwork **INTEGER**. The dimension of the array *iwork*,
lwork $\geq \max(1, 10n)$.

Output Parameters

a On exit, the lower triangle (if *uplo* = 'L') or the upper triangle (if *uplo* = 'U') of *A*, including the diagonal, is overwritten.

m **INTEGER**. The total number of eigenvalues found,
 $0 \leq m \leq n$. If *range* = 'A', $m = n$, and if *range* = 'I',
 $m = i_u - i_l + 1$.

w **REAL** for *cheevr*
DOUBLE PRECISION for *zheevr*.
 Array, **DIMENSION** at least $\max(1, n)$, contains the selected eigenvalues in ascending order, stored in *w*(1) to *w*(*m*).

z **COMPLEX** for *cheevr*
DOUBLE COMPLEX for *zheevr*.
 Array *z*(*ldz*, *); the second dimension of *z* must be at least $\max(1, m)$.
 If *jobz* = 'V', then if *info* = 0, the first *m* columns of *z* contain the orthonormal eigenvectors of the matrix *T* corresponding to the selected eigenvalues, with the *i*-th column of *z* holding the eigenvector associated with *w*(*i*).
 If *jobz* = 'N', then *z* is not referenced.
 Note: you must ensure that at least $\max(1, m)$ columns are supplied in the array *z*; if *range* = 'V', the exact value of *m* is not known in advance and an upper bound must be used.

isuppz **INTEGER**.
 Array, **DIMENSION** at least $2 * \max(1, m)$.

	The support of the eigenvectors in z , i.e., the indices indicating the nonzero elements in z . The i -th eigenvector is nonzero only in elements $isuppz(2i-1)$ through $isuppz(2i)$.
$work(1)$	On exit, if $info = 0$, then $work(1)$ returns the required minimal size of $lwork$.
$rwork(1)$	On exit, if $info = 0$, then $rwork(1)$ returns the required minimal size of $lrwork$.
$iwork(1)$	On exit, if $info = 0$, then $iwork(1)$ returns the required minimal size of $liwork$.
$info$	INTEGER. If $info = 0$, the execution is successful. If $info = -i$, the i th parameter had an illegal value. If $info = i$, an internal error has occurred.

Application Notes

For optimum performance use $lwork \geq (nb+1)*n$, where nb is the maximum of the blocksize for `?hetrd` and `?unmtr` returned by `ilaenv`. If you are in doubt how much workspace to supply, use a generous value of $lwork$ for the first run. On exit, examine $work(1)$ and use this value for subsequent runs.

Normal execution of `?stegr` may create NaNs and infinities and hence may abort due to a floating point exception in environments which do not handle NaNs and infinities in the IEEE standard default manner.

?spev

Computes all eigenvalues and, optionally, eigenvectors of a real symmetric matrix in packed storage.

```
call sspev (jobz, uplo, n, ap, w, z, ldz, work, info)
call dspev (jobz, uplo, n, ap, w, z, ldz, work, info)
```

Discussion

This routine computes all the eigenvalues and, optionally, eigenvectors of a real symmetric matrix A in packed storage.

Input Parameters

jobz CHARACTER*1. Must be 'N' or 'V'.
 If *job* = 'N', then only eigenvalues are computed.
 If *job* = 'V', then eigenvalues and eigenvectors are computed.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', *ap* stores the packed upper triangular part of A .
 If *uplo* = 'L', *ap* stores the packed lower triangular part of A .

n INTEGER. The order of the matrix A ($n \geq 0$).

ap, work REAL for *sspev*
 DOUBLE PRECISION for *dspev*
 Arrays:
ap(*) contains the packed upper or lower triangle of symmetric matrix A , as specified by *uplo*. The dimension of *ap* must be at least $\max(1, n*(n+1)/2)$.
work(*) is a workspace array, DIMENSION at least $\max(1, 3n)$.

ldz **INTEGER**. The leading dimension of the output array *z*.
 Constraints:
 if *jobz* = 'N', then *ldz* ≥ 1;
 if *jobz* = 'V', then *ldz* ≥ max(1, *n*).

Output Parameters

w, z **REAL** for *sspev*
 DOUBLE PRECISION for *dspev*
 Arrays:
w(*), **DIMENSION** at least max(1, *n*).
 If *info* = 0, *w* contains the eigenvalues of the matrix *A*
 in ascending order.
z(*ldz*, *). The second dimension of *z* must be at least
 max(1, *n*).
 If *jobz* = 'V', then if *info* = 0, *z* contains the
 orthonormal eigenvectors of the matrix *A*, with the *i*-th
 column of *z* holding the eigenvector associated with
w(*i*).
 If *jobz* = 'N', then *z* is not referenced.

ap On exit, this array is overwritten by the values generated
 during the reduction to tridiagonal form. The elements
 of the diagonal and the off-diagonal of the tridiagonal
 matrix overwrite the corresponding elements of *A*.

info **INTEGER**.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.
 If *info* = *i*, then the algorithm failed to converge; *i*
 indicates the number of elements of an intermediate
 tridiagonal form which did not converge to zero.

?hpev

Computes all eigenvalues and, optionally, eigenvectors of a Hermitian matrix in packed storage.

```
call chpev (jobz, uplo, n, ap, w, z, ldz, work, rwork, info)
call zhpev (jobz, uplo, n, ap, w, z, ldz, work, rwork, info)
```

Discussion

This routine computes all the eigenvalues and, optionally, eigenvectors of a complex Hermitian matrix A in packed storage.

Input Parameters

jobz CHARACTER*1. Must be 'N' or 'V'.
 If *job* = 'N', then only eigenvalues are computed.
 If *job* = 'V', then eigenvalues and eigenvectors are computed.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', *ap* stores the packed upper triangular part of A .
 If *uplo* = 'L', *ap* stores the packed lower triangular part of A .

n INTEGER. The order of the matrix A ($n \geq 0$).

ap, work COMPLEX for *chpev*
 DOUBLE COMPLEX for *zhpev*.
 Arrays:
ap(*) contains the packed upper or lower triangle of Hermitian matrix A , as specified by *uplo*. The dimension of *ap* must be at least $\max(1, n*(n+1)/2)$.
work(*) is a workspace array, DIMENSION at least $\max(1, 2n-1)$.

ldz **INTEGER**. The leading dimension of the output array *z*.
 Constraints:
 if *jobz* = 'N', then *ldz* ≥ 1;
 if *jobz* = 'V', then *ldz* ≥ max(1, *n*).

rwork **REAL** for *chpev*
DOUBLE PRECISION for *zhpev*.
 Workspace array, **DIMENSION** at least max(1, 3*n*-2).

Output Parameters

w **REAL** for *chpev*
DOUBLE PRECISION for *zhpev*.
 Array, **DIMENSION** at least max(1, *n*).
 If *info* = 0, *w* contains the eigenvalues of the matrix *A*
 in ascending order.

z **COMPLEX** for *chpev*
DOUBLE COMPLEX for *zhpev* .
 Array *z*(*ldz*, *). The second dimension of *z* must be at
 least max(1, *n*).
 If *jobz* = 'V', then if *info* = 0, *z* contains the
 orthonormal eigenvectors of the matrix *A*, with the *i*-th
 column of *z* holding the eigenvector associated with
w(*i*).
 If *jobz* = 'N', then *z* is not referenced.

ap On exit, this array is overwritten by the values generated
 during the reduction to tridiagonal form. The elements
 of the diagonal and the off-diagonal of the tridiagonal
 matrix overwrite the corresponding elements of *A*.

info **INTEGER**.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.
 If *info* = *i*, then the algorithm failed to converge; *i*
 indicates the number of elements of an intermediate
 tridiagonal form which did not converge to zero.

?spevd

Uses divide and conquer algorithm to compute all eigenvalues and (optionally) all eigenvectors of a real symmetric matrix held in packed storage.

```
call sspevd (job,uplo,n,ap,w,z,ldz,work,lwork,iwork,liwork,info)
call dspevd (job,uplo,n,ap,w,z,ldz,work,lwork,iwork,liwork,info)
```

Discussion

This routine computes all the eigenvalues, and optionally all the eigenvectors, of a real symmetric matrix A (held in packed storage). In other words, it can compute the spectral factorization of A as: $A = Z\Lambda Z^T$.

Here Λ is a diagonal matrix whose diagonal elements are the eigenvalues λ_i , and Z is the orthogonal matrix whose columns are the eigenvectors z_i . Thus,

$$Az_i = \lambda_i z_i \text{ for } i = 1, 2, \dots, n.$$

If the eigenvectors are requested, then this routine uses a divide and conquer algorithm to compute eigenvalues and eigenvectors. However, if only eigenvalues are required, then it uses the Pal-Walker-Kahan variant of the QL or QR algorithm.

Input Parameters

job CHARACTER*1. Must be 'N' or 'V'.
 If **job** = 'N', then only eigenvalues are computed.
 If **job** = 'V', then eigenvalues and eigenvectors are computed.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If **uplo** = 'U', **ap** stores the packed upper triangular part of A .
 If **uplo** = 'L', **ap** stores the packed lower triangular part of A .

n INTEGER. The order of the matrix A ($n \geq 0$).

<i>ap, work</i>	<p>REAL for <i>sspevd</i> DOUBLE PRECISION for <i>dspevd</i></p> <p>Arrays: <i>ap(*)</i> contains the packed upper or lower triangle of symmetric matrix A, as specified by <i>uplo</i>. The dimension of <i>ap</i> must be at least $\max(1, n*(n+1)/2)$ <i>work(*)</i> is a workspace array, DIMENSION at least <i>lwork</i>.</p>
<i>ldz</i>	<p>INTEGER. The leading dimension of the output array <i>z</i>. Constraints: if <i>job</i> = 'N', then $ldz \geq 1$; if <i>job</i> = 'V', then $ldz \geq \max(1, n)$.</p>
<i>lwork</i>	<p>INTEGER. The dimension of the array <i>work</i>. Constraints: if $n \leq 1$, then $lwork \geq 1$; if <i>job</i> = 'N' and $n > 1$, then $lwork \geq 2n$; if <i>job</i> = 'V' and $n > 1$, then $lwork \geq 2n^2 + (5+2k) * n + 1$, where <i>k</i> is the smallest integer which satisfies $2^k \geq n$. If <i>lwork</i> = -1, then a workspace query is assumed; the routine only calculates the optimal size of the <i>work</i> array, returns this value as the first entry of the <i>work</i> array, and no error message related to <i>lwork</i> is issued by <i>xerbla</i>.</p>
<i>iwork</i>	<p>INTEGER. Workspace array, DIMENSION at least <i>liwork</i>.</p>
<i>liwork</i>	<p>INTEGER. The dimension of the array <i>iwork</i>. Constraints: if $n \leq 1$, then $liwork \geq 1$; if <i>job</i> = 'N' and $n > 1$, then $liwork \geq 1$; if <i>job</i> = 'V' and $n > 1$, then $liwork \geq 5n+3$. If <i>liwork</i> = -1, then a workspace query is assumed; the routine only calculates the optimal size of the <i>iwork</i> array, returns this value as the first entry of the <i>iwork</i> array, and no error message related to <i>liwork</i> is issued by <i>xerbla</i>.</p>

Output Parameters

<i>w, z</i>	<p>REAL for <code>sspevd</code> DOUBLE PRECISION for <code>dspevd</code> Arrays: <i>w</i>(<i>*</i>), DIMENSION at least $\max(1, n)$. If <i>info</i> = 0, contains the eigenvalues of the matrix A in ascending order. See also <i>info</i>. <i>z</i>(<i>ldz</i>, <i>*</i>) . The second dimension of <i>z</i> must be: at least 1 if <i>job</i> = 'N'; at least $\max(1, n)$ if <i>job</i> = 'V'. If <i>job</i> = 'V', then this array is overwritten by the orthogonal matrix Z which contains the eigenvectors of A. If <i>job</i> = 'N', then <i>z</i> is not referenced.</p>
<i>ap</i>	<p>On exit, this array is overwritten by the values generated during the reduction to tridiagonal form. The elements of the diagonal and the off-diagonal of the tridiagonal matrix overwrite the corresponding elements of A.</p>
<i>work</i> (1)	<p>On exit, if <i>info</i> = 0, then <i>work</i>(1) returns the optimal <i>lwork</i>.</p>
<i>iwork</i> (1)	<p>On exit, if <i>info</i> = 0, then <i>iwork</i>(1) returns the optimal <i>liwork</i>.</p>
<i>info</i>	<p>INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = <i>i</i>, then the algorithm failed to converge; <i>i</i> indicates the number of elements of an intermediate tridiagonal form which did not converge to zero. If <i>info</i> = -<i>i</i>, the <i>i</i>th parameter had an illegal value.</p>

Application Notes

The computed eigenvalues and eigenvectors are exact for a matrix $T + E$ such that $\|E\|_2 = O(\epsilon) \|T\|_2$, where ϵ is the machine precision.

The complex analogue of this routine is [?hpevd](#).

See also [?syevd](#) for matrices held in full storage, and [?sbevd](#) for banded matrices.

?hpevd

Uses divide and conquer algorithm to compute all eigenvalues and (optionally) all eigenvectors of a complex Hermitian matrix held in packed storage.

```
call chpevd (job, uplo, n, ap, w, z, ldz, work, lwork, rwork,
            lrwork, iwork, liwork, info)
call zhpevd (job, uplo, n, ap, w, z, ldz, work, lwork, rwork,
            lrwork, iwork, liwork, info)
```

Discussion

This routine computes all the eigenvalues, and optionally all the eigenvectors, of a complex Hermitian matrix A (held in packed storage). In other words, it can compute the spectral factorization of A as: $A = Z\Lambda Z^H$. Here Λ is a real diagonal matrix whose diagonal elements are the eigenvalues λ_i , and Z is the (complex) unitary matrix whose columns are the eigenvectors z_i . Thus,

$$Az_i = \lambda_i z_i \text{ for } i = 1, 2, \dots, n.$$

If the eigenvectors are requested, then this routine uses a divide and conquer algorithm to compute eigenvalues and eigenvectors. However, if only eigenvalues are required, then it uses the Pal-Walker-Kahan variant of the QL or QR algorithm.

Input Parameters

<i>job</i>	CHARACTER*1. Must be 'N' or 'V'. If <i>job</i> = 'N', then only eigenvalues are computed. If <i>job</i> = 'V', then eigenvalues and eigenvectors are computed.
<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'.

If `uplo = 'U'`, `ap` stores the packed upper triangular part of `A`.
 If `uplo = 'L'`, `ap` stores the packed lower triangular part of `A`.

`n` **INTEGER**. The order of the matrix `A` ($n \geq 0$).

`ap,work` **COMPLEX** for `chpevd`
DOUBLE COMPLEX for `zhpevd`
 Arrays:
`ap(*)` contains the packed upper or lower triangle of Hermitian matrix `A`, as specified by `uplo`. The dimension of `ap` must be at least $\max(1, n*(n+1)/2)$
`work(*)` is a workspace array, **DIMENSION** at least `lwork`.

`ldz` **INTEGER**. The leading dimension of the output array `z`.
 Constraints:
 if `job = 'N'`, then $ldz \geq 1$;
 if `job = 'V'`, then $ldz \geq \max(1, n)$.

`lwork` **INTEGER**. The dimension of the array `work`.
 Constraints:
 if $n \leq 1$, then $lwork \geq 1$;
 if `job = 'N'` and $n > 1$, then $lwork \geq n$;
 if `job = 'V'` and $n > 1$, then $lwork \geq 2n$

`rwork` **REAL** for `chpevd`
DOUBLE PRECISION for `zhpevd`
 Workspace array, **DIMENSION** at least `lrwork`.

`lrwork` **INTEGER**. The dimension of the array `rwork`.
 Constraints:
 if $n \leq 1$, then $lrwork \geq 1$;
 if `job = 'N'` and $n > 1$, then $lrwork \geq n$;
 if `job = 'V'` and $n > 1$, then
 $lrwork \geq 3n^2 + (4+2k)*n+1$, where k is the smallest integer which satisfies $2^k \geq n$.

`iwork` **INTEGER**.
 Workspace array, **DIMENSION** at least `liwork`.

liwork INTEGER. The dimension of the array *iwork*.
 Constraints:
 if $n \leq 1$, then $liwork \geq 1$;
 if $job = 'N'$ and $n > 1$, then $liwork \geq 1$;
 if $job = 'V'$ and $n > 1$, then $liwork \geq 5n+2$.

Output Parameters

w REAL for *chpevd*
 DOUBLE PRECISION for *zhpevd*
 Array, DIMENSION at least $\max(1, n)$.
 If $info = 0$, contains the eigenvalues of the matrix *A* in ascending order. See also *info*.

z COMPLEX for *chpevd*
 DOUBLE COMPLEX for *zhpevd*
 Array, DIMENSION (*ldz*, *). The second dimension of *z* must be:
 at least 1 if $job = 'N'$;
 at least $\max(1, n)$ if $job = 'V'$.
 If $job = 'V'$, then this array is overwritten by the unitary matrix *Z* which contains the eigenvectors of *A*. If $job = 'N'$, then *z* is not referenced.

ap On exit, this array is overwritten by the values generated during the reduction to tridiagonal form. The elements of the diagonal and the off-diagonal of the tridiagonal matrix overwrite the corresponding elements of *A*.

work(1) On exit, if $lwork > 0$, then the real part of *work(1)* returns the required minimal size of *lwork*.

rwork(1) On exit, if $lrwork > 0$, then *rwork(1)* returns the required minimal size of *lrwork*.

iwork(1) On exit, if $liwork > 0$, then *iwork(1)* returns the required minimal size of *liwork*.

info INTEGER.
 If $info = 0$, the execution is successful.
 If $info = i$, then the algorithm failed to converge; *i*

indicates the number of elements of an intermediate tridiagonal form which did not converge to zero.
If $info = -i$, the i th parameter had an illegal value.

Application Notes

The computed eigenvalues and eigenvectors are exact for a matrix $T + E$ such that $\|E\|_2 = O(\epsilon) \|T\|_2$, where ϵ is the machine precision.

The real analogue of this routine is [?spevd](#).

See also [?heevd](#) for matrices held in full storage, and [?hbevd](#) for banded matrices.

?spevx

Computes selected eigenvalues and, optionally, eigenvectors of a real symmetric matrix in packed storage.

```
call sspevx (jobz, range, uplo, n, ap, vl, vu, il, iu, abstol,
            m, w, z, ldz, work, iwork, ifail, info)
call dspevx (jobz, range, uplo, n, ap, vl, vu, il, iu, abstol,
            m, w, z, ldz, work, iwork, ifail, info)
```

Discussion

This routine computes selected eigenvalues and, optionally, eigenvectors of a real symmetric matrix A in packed storage. Eigenvalues and eigenvectors can be selected by specifying either a range of values or a range of indices for the desired eigenvalues.

Input Parameters

<i>jobz</i>	CHARACTER*1. Must be 'N' or 'V'. If <i>job</i> = 'N', then only eigenvalues are computed. If <i>job</i> = 'V', then eigenvalues and eigenvectors are computed.
<i>range</i>	CHARACTER*1. Must be 'A' or 'V' or 'I'. If <i>range</i> = 'A', the routine computes all eigenvalues. If <i>range</i> = 'V', the routine computes eigenvalues λ_i in the half-open interval: $vl < \lambda_i \leq vu$. If <i>range</i> = 'I', the routine computes eigenvalues with indices <i>il</i> to <i>iu</i> .
<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. If <i>uplo</i> = 'U', <i>ap</i> stores the packed upper triangular part of A . If <i>uplo</i> = 'L', <i>ap</i> stores the packed lower triangular part of A .
<i>n</i>	INTEGER. The order of the matrix A ($n \geq 0$).

ap, work REAL for *sspevx*
DOUBLE PRECISION for *dspevx*
Arrays:
ap()* contains the packed upper or lower triangle of the symmetric matrix *A*, as specified by *uplo*. The dimension of *ap* must be at least $\max(1, n*(n+1)/2)$.
work()* is a workspace array, DIMENSION at least $\max(1, 8n)$.

vl, vu REAL for *sspevx*
DOUBLE PRECISION for *dspevx*
If *range* = 'V', the lower and upper bounds of the interval to be searched for eigenvalues.
Constraint: $vl < vu$.
If *range* = 'A' or 'I', *vl* and *vu* are not referenced.

il, iu INTEGER.
If *range* = 'I', the indices in ascending order of the smallest and largest eigenvalues to be returned.
Constraint: $1 \leq il \leq iu \leq n$, if $n > 0$; $il=1$ and $iu=0$ if $n = 0$.
If *range* = 'A' or 'V', *il* and *iu* are not referenced.

abstol REAL for *sspevx*
DOUBLE PRECISION for *dspevx*
The absolute error tolerance to which each eigenvalue is required. See *Application notes* for details on error tolerance.

ldz INTEGER. The leading dimension of the output array *z*.
Constraints:
if *jobz* = 'N', then $ldz \geq 1$;
if *jobz* = 'V', then $ldz \geq \max(1, n)$.

iwork INTEGER.
Workspace array, DIMENSION at least $\max(1, 5n)$.

Output Parameters

<i>ap</i>	On exit, this array is overwritten by the values generated during the reduction to tridiagonal form. The elements of the diagonal and the off-diagonal of the tridiagonal matrix overwrite the corresponding elements of <i>A</i> .
<i>m</i>	INTEGER . The total number of eigenvalues found, $0 \leq m \leq n$. If <i>range</i> = 'A', $m = n$, and if <i>range</i> = 'I', $m = iu - il + 1$.
<i>w, z</i>	REAL for sspevx DOUBLE PRECISION for dspevx Arrays: <i>w</i> (*), DIMENSION at least $\max(1, n)$. If <i>info</i> = 0, contains the selected eigenvalues of the matrix <i>A</i> in ascending order. <i>z</i> (<i>ldz</i> , *) . The second dimension of <i>z</i> must be at least $\max(1, m)$. If <i>jobz</i> = 'V', then if <i>info</i> = 0, the first <i>m</i> columns of <i>z</i> contain the orthonormal eigenvectors of the matrix <i>A</i> corresponding to the selected eigenvalues, with the <i>i</i> -th column of <i>z</i> holding the eigenvector associated with <i>w</i> (<i>i</i>). If an eigenvector fails to converge, then that column of <i>z</i> contains the latest approximation to the eigenvector, and the index of the eigenvector is returned in <i>ifail</i> . If <i>jobz</i> = 'N', then <i>z</i> is not referenced. Note: you must ensure that at least $\max(1, m)$ columns are supplied in the array <i>z</i> ; if <i>range</i> = 'V', the exact value of <i>m</i> is not known in advance and an upper bound must be used.
<i>ifail</i>	INTEGER . Array, DIMENSION at least $\max(1, n)$. If <i>jobz</i> = 'V', then if <i>info</i> = 0, the first <i>m</i> elements of <i>ifail</i> are zero; if <i>info</i> > 0, the <i>ifail</i> contains the indices the eigenvectors that failed to converge. If <i>jobz</i> = 'N', then <i>ifail</i> is not referenced.

info

INTEGER.

If *info* = 0, the execution is successful.If *info* = -*i*, the *i*th parameter had an illegal value.If *info* = *i*, then *i* eigenvectors failed to converge; their indices are stored in the array *ifail*.

Application Notes

An approximate eigenvalue is accepted as converged when it is determined to lie in an interval [a,b] of width less than or equal to

$abstol + \epsilon * \max(|a|, |b|)$, where ϵ is the machine precision. If *abstol* is less than or equal to zero, then $\epsilon * \|T\|_1$ will be used in its place, where *T* is the tridiagonal matrix obtained by reducing *A* to tridiagonal form.

Eigenvalues will be computed most accurately when *abstol* is set to twice the underflow threshold $2 * ?lamch('S')$, not zero. If this routine returns with *info* > 0, indicating that some eigenvectors did not converge, try setting *abstol* to $2 * ?lamch('S')$.

?hpevx

Computes selected eigenvalues and, optionally, eigenvectors of a Hermitian matrix in packed storage.

```
call chpevx (jobz, range, uplo, n, ap, vl, vu, il, iu, abstol,
            m, w, z, ldz, work, rwork, iwork, ifail, info)
call zhpevx (jobz, range, uplo, n, ap, vl, vu, il, iu, abstol,
            m, w, z, ldz, work, rwork, iwork, ifail, info)
```

Discussion

This routine computes selected eigenvalues and, optionally, eigenvectors of a complex Hermitian matrix A in packed storage. Eigenvalues and eigenvectors can be selected by specifying either a range of values or a range of indices for the desired eigenvalues.

Input Parameters

<i>jobz</i>	CHARACTER*1. Must be 'N' or 'V'. If <i>job</i> = 'N', then only eigenvalues are computed. If <i>job</i> = 'V', then eigenvalues and eigenvectors are computed.
<i>range</i>	CHARACTER*1. Must be 'A' or 'V' or 'I'. If <i>range</i> = 'A', the routine computes all eigenvalues. If <i>range</i> = 'V', the routine computes eigenvalues λ_i in the half-open interval: $vl < \lambda_i \leq vu$. If <i>range</i> = 'I', the routine computes eigenvalues with indices <i>il</i> to <i>iu</i> .
<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. If <i>uplo</i> = 'U', <i>ap</i> stores the packed upper triangular part of A . If <i>uplo</i> = 'L', <i>ap</i> stores the packed lower triangular part of A .
<i>n</i>	INTEGER. The order of the matrix A ($n \geq 0$).

<code>ap, work</code>	<p>COMPLEX for <code>chpevx</code> DOUBLE COMPLEX for <code>zhpevx</code></p> <p>Arrays: <code>ap(*)</code> contains the packed upper or lower triangle of the Hermitian matrix A, as specified by <code>uplo</code>. The dimension of <code>ap</code> must be at least $\max(1, n*(n+1)/2)$. <code>work(*)</code> is a workspace array, DIMENSION at least $\max(1, 2n)$.</p>
<code>vl, vu</code>	<p>REAL for <code>chpevx</code> DOUBLE PRECISION for <code>zhpevx</code></p> <p>If <code>range = 'V'</code>, the lower and upper bounds of the interval to be searched for eigenvalues. Constraint: <code>vl < vu</code>. If <code>range = 'A'</code> or <code>'I'</code>, <code>vl</code> and <code>vu</code> are not referenced.</p>
<code>il, iu</code>	<p>INTEGER.</p> <p>If <code>range = 'I'</code>, the indices in ascending order of the smallest and largest eigenvalues to be returned. Constraint: $1 \leq il \leq iu \leq n$, if $n > 0$; <code>il=1</code> and <code>iu=0</code> if $n = 0$. If <code>range = 'A'</code> or <code>'V'</code>, <code>il</code> and <code>iu</code> are not referenced.</p>
<code>abstol</code>	<p>REAL for <code>chpevx</code> DOUBLE PRECISION for <code>zhpevx</code></p> <p>The absolute error tolerance to which each eigenvalue is required. See <i>Application notes</i> for details on error tolerance.</p>
<code>ldz</code>	<p>INTEGER. The leading dimension of the output array <code>z</code>. Constraints: if <code>jobz = 'N'</code>, then <code>ldz</code> ≥ 1; if <code>jobz = 'V'</code>, then <code>ldz</code> $\geq \max(1, n)$.</p>
<code>rwork</code>	<p>REAL for <code>chpevx</code> DOUBLE PRECISION for <code>zhpevx</code></p> <p>Workspace array, DIMENSION at least $\max(1, 7n)$.</p>
<code>iwork</code>	<p>INTEGER. Workspace array, DIMENSION at least $\max(1, 5n)$.</p>

Output Parameters

<i>ap</i>	On exit, this array is overwritten by the values generated during the reduction to tridiagonal form. The elements of the diagonal and the off-diagonal of the tridiagonal matrix overwrite the corresponding elements of <i>A</i> .
<i>m</i>	INTEGER . The total number of eigenvalues found, $0 \leq m \leq n$. If <i>range</i> = 'A', $m = n$, and if <i>range</i> = 'I', $m = iu - il + 1$.
<i>w</i>	REAL for <i>chpevx</i> DOUBLE PRECISION for <i>zhpevx</i> Array, DIMENSION at least $\max(1, n)$. If <i>info</i> = 0, contains the selected eigenvalues of the matrix <i>A</i> in ascending order.
<i>z</i>	COMPLEX for <i>chpevx</i> DOUBLE COMPLEX for <i>zhpevx</i> Array <i>z</i> (<i>ldz</i> , *). The second dimension of <i>z</i> must be at least $\max(1, m)$. If <i>jobz</i> = 'V', then if <i>info</i> = 0, the first <i>m</i> columns of <i>z</i> contain the orthonormal eigenvectors of the matrix <i>A</i> corresponding to the selected eigenvalues, with the <i>i</i> -th column of <i>z</i> holding the eigenvector associated with <i>w</i> (<i>i</i>). If an eigenvector fails to converge, then that column of <i>z</i> contains the latest approximation to the eigenvector, and the index of the eigenvector is returned in <i>ifail</i> . If <i>jobz</i> = 'N', then <i>z</i> is not referenced. Note: you must ensure that at least $\max(1, m)$ columns are supplied in the array <i>z</i> ; if <i>range</i> = 'V', the exact value of <i>m</i> is not known in advance and an upper bound must be used.
<i>ifail</i>	INTEGER . Array, DIMENSION at least $\max(1, n)$. If <i>jobz</i> = 'V', then if <i>info</i> = 0, the first <i>m</i> elements of <i>ifail</i> are zero; if <i>info</i> > 0, the <i>ifail</i> contains the indices the eigenvectors that failed to converge. If <i>jobz</i> = 'N', then <i>ifail</i> is not referenced.

info

INTEGER.

If *info* = 0, the execution is successful.If *info* = -*i*, the *i*th parameter had an illegal value.If *info* = *i*, then *i* eigenvectors failed to converge; their indices are stored in the array *ifail*.

Application Notes

An approximate eigenvalue is accepted as converged when it is determined to lie in an interval [a,b] of width less than or equal to

$abstol + \epsilon * \max(|a|, |b|)$, where ϵ is the machine precision. If *abstol* is less than or equal to zero, then $\epsilon * \|T\|_1$ will be used in its place, where *T* is the tridiagonal matrix obtained by reducing *A* to tridiagonal form.

Eigenvalues will be computed most accurately when *abstol* is set to twice the underflow threshold $2 * ?lamch('S')$, not zero. If this routine returns with *info* > 0, indicating that some eigenvectors did not converge, try setting *abstol* to $2 * ?lamch('S')$.

?sbev

Computes all eigenvalues and, optionally, eigenvectors of a real symmetric band matrix.

```
call ssbev (jobz, uplo, n, kd, ab, ldab, w, z, ldz, work, info)
call dsbev (jobz, uplo, n, kd, ab, ldab, w, z, ldz, work, info)
```

Discussion

This routine computes all eigenvalues and, optionally, eigenvectors of a real symmetric band matrix A .

Input Parameters

<i>jobz</i>	CHARACTER*1. Must be 'N' or 'V'. If <i>jobz</i> = 'N', then only eigenvalues are computed. If <i>jobz</i> = 'V', then eigenvalues and eigenvectors are computed.
<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. If <i>uplo</i> = 'U', <i>ab</i> stores the upper triangular part of A . If <i>uplo</i> = 'L', <i>ab</i> stores the lower triangular part of A .
<i>n</i>	INTEGER. The order of the matrix A ($n \geq 0$).
<i>kd</i>	INTEGER. The number of super- or sub-diagonals in A ($kd \geq 0$).
<i>ab, work</i>	REAL for <i>ssbev</i> DOUBLE PRECISION for <i>dsbev</i> . Arrays: <i>ab</i> (<i>ldab</i> , *) is an array containing either upper or lower triangular part of the symmetric matrix A (as specified by <i>uplo</i>) in band storage format. The second dimension of <i>ab</i> must be at least $\max(1, n)$. <i>work</i> (*) is a workspace array. The dimension of <i>work</i> must be at least $\max(1, 3n-2)$.

ldab **INTEGER**. The leading dimension of *ab*; must be at least $kd + 1$.

ldz **INTEGER**. The leading dimension of the output array *z*.
Constraints:
if *jobz* = 'N', then $ldz \geq 1$;
if *jobz* = 'V', then $ldz \geq \max(1, n)$.

Output Parameters

w, z **REAL** for *ssbev*
DOUBLE PRECISION for *dsbev*
Arrays:
w(*), **DIMENSION** at least $\max(1, n)$.
If *info* = 0, contains the eigenvalues of the matrix *A* in ascending order.
z(*ldz*, *). The second dimension of *z* must be at least $\max(1, n)$.
If *jobz* = 'V', then if *info* = 0, *z* contains the orthonormal eigenvectors of the matrix *A*, with the *i*-th column of *z* holding the eigenvector associated with *w*(*i*).
If *jobz* = 'N', then *z* is not referenced.

ab On exit, this array is overwritten by the values generated during the reduction to tridiagonal form. If *uplo* = 'U', the first superdiagonal and the diagonal of the tridiagonal matrix *T* are returned in rows *kd* and *kd*+1 of *ab*, and if *uplo* = 'L', the diagonal and first subdiagonal of *T* are returned in the first two rows of *ab*.

info **INTEGER**.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i*, then the algorithm failed to converge; *i* indicates the number of elements of an intermediate tridiagonal form which did not converge to zero.

?hbev

Computes all eigenvalues and, optionally, eigenvectors of a Hermitian band matrix.

```
call chbev(jobz, uplo, n, kd, ab, ldab, w, z, ldz, work, rwork, info)
call zhbev(jobz, uplo, n, kd, ab, ldab, w, z, ldz, work, rwork, info)
```

Discussion

This routine computes all eigenvalues and, optionally, eigenvectors of a complex Hermitian band matrix A.

Input Parameters

<i>jobz</i>	CHARACTER*1. Must be 'N' or 'V'. If <i>jobz</i> = 'N', then only eigenvalues are computed. If <i>jobz</i> = 'V', then eigenvalues and eigenvectors are computed.
<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. If <i>uplo</i> = 'U', <i>ab</i> stores the upper triangular part of A. If <i>uplo</i> = 'L', <i>ab</i> stores the lower triangular part of A.
<i>n</i>	INTEGER. The order of the matrix A ($n \geq 0$).
<i>kd</i>	INTEGER. The number of super- or sub-diagonals in A ($kd \geq 0$).
<i>ab, work</i>	COMPLEX for <i>chbev</i> DOUBLE COMPLEX for <i>zhbev</i> . Arrays: <i>ab</i> (<i>ldab</i> , *) is an array containing either upper or lower triangular part of the Hermitian matrix A (as specified by <i>uplo</i>) in band storage format. The second dimension of <i>ab</i> must be at least $\max(1, n)$. <i>work</i> (*) is a workspace array. The dimension of <i>work</i> must be at least $\max(1, n)$.

ldab **INTEGER**. The leading dimension of *ab*; must be at least $kd + 1$.

ldz **INTEGER**. The leading dimension of the output array *z*.
Constraints:
if *jobz* = 'N', then $ldz \geq 1$;
if *jobz* = 'V', then $ldz \geq \max(1, n)$.

rwork **REAL** for *chbev*
DOUBLE PRECISION for *zhbev*
Workspace array, **DIMENSION** at least $\max(1, 3n-2)$.

Output Parameters

w **REAL** for *chbev*
DOUBLE PRECISION for *zhbev*
Array, **DIMENSION** at least $\max(1, n)$. If *info* = 0, contains the eigenvalues in ascending order.

z **COMPLEX** for *chbev*
DOUBLE COMPLEX for *zhbev*.
Array $z(ldz, *)$. The second dimension of *z* must be at least $\max(1, n)$. If *jobz* = 'V', then if *info* = 0, *z* contains the orthonormal eigenvectors of the matrix *A*, with the *i*-th column of *z* holding the eigenvector associated with *w*(*i*). If *jobz* = 'N', then *z* is not referenced.

ab On exit, this array is overwritten by the values generated during the reduction to tridiagonal form. If *uplo* = 'U', the first superdiagonal and the diagonal of the tridiagonal matrix *T* are returned in rows *kd* and *kd*+1 of *ab*, and if *uplo* = 'L', the diagonal and first subdiagonal of *T* are returned in the first two rows of *ab*.

info **INTEGER**.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i*, then the algorithm failed to converge;
i indicates the number of elements of an intermediate tridiagonal form which did not converge to zero.

?sbevd

Computes all eigenvalues and (optionally) all eigenvectors of a real symmetric band matrix using divide and conquer algorithm.

```
call ssbevd (job, uplo, n, kd, ab, ldab, w, z, ldz, work, lwork,
            iwork, liwork, info)
call dsbevd (job, uplo, n, kd, ab, ldab, w, z, ldz, work, lwork,
            iwork, liwork, info)
```

Discussion

This routine computes all the eigenvalues, and optionally all the eigenvectors, of a real symmetric band matrix A . In other words, it can compute the spectral factorization of A as:

$$A = Z\Lambda Z^T$$

Here Λ is a diagonal matrix whose diagonal elements are the eigenvalues λ_i , and Z is the orthogonal matrix whose columns are the eigenvectors z_i . Thus,

$$Az_i = \lambda_i z_i \text{ for } i = 1, 2, \dots, n.$$

If the eigenvectors are requested, then this routine uses a divide and conquer algorithm to compute eigenvalues and eigenvectors. However, if only eigenvalues are required, then it uses the Pal-Walker-Kahan variant of the QL or QR algorithm.

Input Parameters

job CHARACTER*1. Must be 'N' or 'V'.
 If *job* = 'N', then only eigenvalues are computed.
 If *job* = 'V', then eigenvalues and eigenvectors are computed.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', *ab* stores the upper triangular part of A .
 If *uplo* = 'L', *ab* stores the lower triangular part of A .

<i>n</i>	INTEGER . The order of the matrix <i>A</i> ($n \geq 0$).
<i>kd</i>	INTEGER . The number of super- or sub-diagonals in <i>A</i> ($kd \geq 0$).
<i>ab, work</i>	REAL for ssbevd DOUBLE PRECISION for dsbevd . Arrays: <i>ab</i> (<i>ldab</i> , *) is an array containing either upper or lower triangular part of the symmetric matrix <i>A</i> (as specified by <i>uplo</i>) in band storage format. The second dimension of <i>ab</i> must be at least $\max(1, n)$. <i>work</i> (*) is a workspace array. The dimension of <i>work</i> must be at least <i>lwork</i> .
<i>ldab</i>	INTEGER . The leading dimension of <i>ab</i> ; must be at least $kd+1$.
<i>ldz</i>	INTEGER . The leading dimension of the output array <i>z</i> . Constraints: if <i>job</i> = 'N', then $ldz \geq 1$; if <i>job</i> = 'V', then $ldz \geq \max(1, n)$.
<i>lwork</i>	INTEGER . The dimension of the array <i>work</i> . Constraints: if $n \leq 1$, then $lwork \geq 1$; if <i>job</i> = 'N' and $n > 1$, then $lwork \geq 2n$; if <i>job</i> = 'V' and $n > 1$, then $lwork \geq 3n^2 + (4+2k) * n + 1$, where <i>k</i> is the smallest integer which satisfies $2^k \geq n$.
<i>iwork</i>	INTEGER . Workspace array, DIMENSION at least <i>liwork</i> .
<i>liwork</i>	INTEGER . The dimension of the array <i>iwork</i> . Constraints: if $n \leq 1$, then $liwork \geq 1$; if <i>job</i> = 'N' and $n > 1$, then $liwork \geq 1$; if <i>job</i> = 'V' and $n > 1$, then $liwork \geq 5n+2$.

Output Parameters

<i>w, z</i>	<p>REAL for <i>ssbevd</i> DOUBLE PRECISION for <i>dsbevd</i></p> <p>Arrays: <i>w(*)</i>, DIMENSION at least $\max(1, n)$. If <i>info</i> = 0, contains the eigenvalues of the matrix <i>A</i> in ascending order. See also <i>info</i>. <i>z(ldz, *)</i> . The second dimension of <i>z</i> must be: at least 1 if <i>job</i> = 'N'; at least $\max(1, n)$ if <i>job</i> = 'V'. If <i>job</i> = 'V', then this array is overwritten by the orthogonal matrix <i>Z</i> which contains the eigenvectors of <i>A</i>. The <i>i</i>th column of <i>Z</i> contains the eigenvector which corresponds to the eigenvalue <i>w(i)</i>. If <i>job</i> = 'N', then <i>z</i> is not referenced.</p>
<i>ab</i>	On exit, this array is overwritten by the values generated during the reduction to tridiagonal form.
<i>work(1)</i>	On exit, if <i>lwork</i> > 0, then <i>work(1)</i> returns the required minimal size of <i>lwork</i> .
<i>iwork(1)</i>	On exit, if <i>liwork</i> > 0, then <i>iwork(1)</i> returns the required minimal size of <i>liwork</i> .
<i>info</i>	<p>INTEGER.</p> <p>If <i>info</i> = 0, the execution is successful. If <i>info</i> = <i>i</i>, then the algorithm failed to converge; <i>i</i> indicates the number of elements of an intermediate tridiagonal form which did not converge to zero. If <i>info</i> = -<i>i</i>, the <i>i</i>th parameter had an illegal value.</p>

Application Notes

The computed eigenvalues and eigenvectors are exact for a matrix $T + E$ such that $\|E\|_2 = O(\epsilon) \|T\|_2$, where ϵ is the machine precision.

The complex analogue of this routine is [?hbevd](#).

See also [?syevd](#) for matrices held in full storage, and [?spevd](#) for matrices held in packed storage.

?hbevd

Computes all eigenvalues and (optionally) all eigenvectors of a complex Hermitian band matrix using divide and conquer algorithm.

```
call chbevd (job, uplo, n, kd, ab, ldab, w, z, ldz, work, lwork,
            rwork, lrwork, iwork, liwork, info)
call zhbevd (job, uplo, n, kd, ab, ldab, w, z, ldz, work, lwork,
            rwork, lrwork, iwork, liwork, info)
```

Discussion

This routine computes all the eigenvalues, and optionally all the eigenvectors, of a complex Hermitian band matrix A . In other words, it can compute the spectral factorization of A as: $A = Z\Lambda Z^H$.

Here Λ is a real diagonal matrix whose diagonal elements are the eigenvalues λ_i , and Z is the (complex) unitary matrix whose columns are the eigenvectors z_i . Thus,

$$Az_i = \lambda_i z_i \text{ for } i = 1, 2, \dots, n.$$

If the eigenvectors are requested, then this routine uses a divide and conquer algorithm to compute eigenvalues and eigenvectors. However, if only eigenvalues are required, then it uses the Pal-Walker-Kahan variant of the QL or QR algorithm.

Input Parameters

job CHARACTER*1. Must be 'N' or 'V'.
 If *job* = 'N', then only eigenvalues are computed.
 If *job* = 'V', then eigenvalues and eigenvectors are computed.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', *ab* stores the upper triangular part of A .
 If *uplo* = 'L', *ab* stores the lower triangular part of A .

n INTEGER. The order of the matrix A ($n \geq 0$).

<i>kd</i>	INTEGER . The number of super- or sub-diagonals in <i>A</i> ($kd \geq 0$).
<i>ab, work</i>	COMPLEX for <i>chbevd</i> DOUBLE COMPLEX for <i>zhbevd</i> . Arrays: <i>ab (ldab, *)</i> is an array containing either upper or lower triangular part of the Hermitian matrix <i>A</i> (as specified by <i>uplo</i>) in band storage format. The second dimension of <i>ab</i> must be at least $\max(1, n)$. <i>work (*)</i> is a workspace array. The dimension of <i>work</i> must be at least <i>lwork</i> .
<i>ldab</i>	INTEGER . The leading dimension of <i>ab</i> ; must be at least $kd+1$.
<i>ldz</i>	INTEGER . The leading dimension of the output array <i>z</i> . Constraints: if <i>job</i> = 'N', then $ldz \geq 1$; if <i>job</i> = 'V', then $ldz \geq \max(1, n)$.
<i>lwork</i>	INTEGER . The dimension of the array <i>work</i> . Constraints: if $n \leq 1$, then $lwork \geq 1$; if <i>job</i> = 'N' and $n > 1$, then $lwork \geq n$; if <i>job</i> = 'V' and $n > 1$, then $lwork \geq 2n^2$
<i>rwork</i>	REAL for <i>chbevd</i> DOUBLE PRECISION for <i>zhbevd</i> Workspace array, DIMENSION at least <i>lrwork</i> .
<i>lrwork</i>	INTEGER . The dimension of the array <i>rwork</i> . Constraints: if $n \leq 1$, then $lrwork \geq 1$; if <i>job</i> = 'N' and $n > 1$, then $lrwork \geq n$; if <i>job</i> = 'V' and $n > 1$, then $lrwork \geq 3n^2 + (4+2k) * n + 1$, where <i>k</i> is the smallest integer which satisfies $2^k \geq n$.
<i>iwork</i>	INTEGER . Workspace array, DIMENSION at least <i>liwork</i> .

liwork INTEGER. The dimension of the array *iwork*.
 Constraints:
 if *job* = 'N' or $n \leq 1$, then $liwork \geq 1$;
 if *job* = 'V' and $n > 1$, then $liwork \geq 5n+2$.

Output Parameters

w REAL for *chbevd*
 DOUBLE PRECISION for *zhbevd*
 Array, DIMENSION at least $\max(1, n)$.
 If *info* = 0, contains the eigenvalues of the matrix *A* in ascending order. See also *info*.

z COMPLEX for *chbevd*
 DOUBLE COMPLEX for *zhbevd*
 Array, DIMENSION (*ldz*, *). The second dimension of *z* must be:
 at least 1 if *job* = 'N';
 at least $\max(1, n)$ if *job* = 'V'.
 If *job* = 'V', then this array is overwritten by the unitary matrix *Z* which contains the eigenvectors of *A*. The *i*th column of *Z* contains the eigenvector which corresponds to the eigenvalue $w(i)$.
 If *job* = 'N', then *z* is not referenced.

ab On exit, this array is overwritten by the values generated during the reduction to tridiagonal form.

work(1) On exit, if *lwork* > 0, then the real part of *work(1)* returns the required minimal size of *lwork*.

rwork(1) On exit, if *lrwork* > 0, then *rwork(1)* returns the required minimal size of *lrwork*.

iwork(1) On exit, if *liwork* > 0, then *iwork(1)* returns the required minimal size of *liwork*.

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = *i*, then the algorithm failed to converge; *i*

indicates the number of elements of an intermediate tridiagonal form which did not converge to zero.

If $info = -i$, the i th parameter had an illegal value.

Application Notes

The computed eigenvalues and eigenvectors are exact for a matrix $T + E$ such that $\|E\|_2 = O(\epsilon) \|T\|_2$, where ϵ is the machine precision.

The real analogue of this routine is [?sbevd](#).

See also [?heevd](#) for matrices held in full storage, and [?hpevd](#) for matrices held in packed storage.

?sbev

Computes selected eigenvalues and, optionally, eigenvectors of a real symmetric band matrix.

```
call ssbev ( jobz, range, uplo, n, kd, ab, ldab, q, ldq, vl, vu, il,
            iu, abstol, m, w, z, ldz, work, iwork, ifail, info)
call dsbev ( jobz, range, uplo, n, kd, ab, ldab, q, ldq, vl, vu, il,
            iu, abstol, m, w, z, ldz, work, iwork, ifail, info)
```

Discussion

This routine computes selected eigenvalues and, optionally, eigenvectors of a real symmetric band matrix A . Eigenvalues and eigenvectors can be selected by specifying either a range of values or a range of indices for the desired eigenvalues.

Input Parameters

jobz CHARACTER*1. Must be 'N' or 'V'.
 If *job* = 'N', then only eigenvalues are computed.
 If *job* = 'V', then eigenvalues and eigenvectors are computed.

range CHARACTER*1. Must be 'A' or 'V' or 'I'.
 If *range* = 'A', the routine computes all eigenvalues.
 If *range* = 'V', the routine computes eigenvalues λ_i in the half-open interval: $vl < \lambda_i \leq vu$.
 If *range* = 'I', the routine computes eigenvalues with indices *il* to *iu*.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', *ab* stores the upper triangular part of A .
 If *uplo* = 'L', *ab* stores the lower triangular part of A .

n INTEGER. The order of the matrix A ($n \geq 0$).

kd INTEGER. The number of super- or sub-diagonals in A ($kd \geq 0$).

<i>ab, work</i>	<p>REAL for <i>ssbevx</i> DOUBLE PRECISION for <i>dsbevx</i>. Arrays: <i>ab (ldab, *)</i> is an array containing either upper or lower triangular part of the symmetric matrix <i>A</i> (as specified by <i>uplo</i>) in band storage format. The second dimension of <i>ab</i> must be at least $\max(1, n)$. <i>work (*)</i> is a workspace array. The dimension of <i>work</i> must be at least $\max(1, 7n)$.</p>
<i>ldab</i>	<p>INTEGER. The leading dimension of <i>ab</i>; must be at least $kd + 1$.</p>
<i>vl, vu</i>	<p>REAL for <i>ssbevx</i> DOUBLE PRECISION for <i>dsbevx</i>. If <i>range</i> = 'V', the lower and upper bounds of the interval to be searched for eigenvalues. Constraint: $vl < vu$. If <i>range</i> = 'A' or 'I', <i>vl</i> and <i>vu</i> are not referenced.</p>
<i>il, iu</i>	<p>INTEGER. If <i>range</i> = 'I', the indices in ascending order of the smallest and largest eigenvalues to be returned. Constraint: $1 \leq il \leq iu \leq n$, if $n > 0$; $il=1$ and $iu=0$ if $n = 0$. If <i>range</i> = 'A' or 'V', <i>il</i> and <i>iu</i> are not referenced.</p>
<i>abstol</i>	<p>REAL for <i>chpevx</i> DOUBLE PRECISION for <i>zhpevx</i> The absolute error tolerance to which each eigenvalue is required. See <i>Application notes</i> for details on error tolerance.</p>
<i>ldq, ldz</i>	<p>INTEGER. The leading dimensions of the output arrays <i>q</i> and <i>z</i>, respectively. Constraints: $ldq \geq 1, ldz \geq 1$; If <i>jobz</i> = 'V', then $ldq \geq \max(1, n)$ and $ldz \geq \max(1, n)$.</p>
<i>iwork</i>	<p>INTEGER. Workspace array, DIMENSION at least $\max(1, 5n)$.</p>

Output Parameters

- m* **INTEGER**. The total number of eigenvalues found, $0 \leq m \leq n$. If *range* = 'A', $m = n$, and if *range* = 'I', $m = i_u - i_l + 1$.
- w, z* **REAL** for *ssbev*x
DOUBLE PRECISION for *dsbev*x
Arrays:
w(*), **DIMENSION** at least $\max(1, n)$.
The first *m* elements of *w* contain the selected eigenvalues of the matrix *A* in ascending order.
z(*ldz*, *) . The second dimension of *z* must be at least $\max(1, m)$.
If *jobz* = 'V', then if *info* = 0, the first *m* columns of *z* contain the orthonormal eigenvectors of the matrix *A* corresponding to the selected eigenvalues, with the *i*-th column of *z* holding the eigenvector associated with *w*(*i*). If an eigenvector fails to converge, then that column of *z* contains the latest approximation to the eigenvector, and the index of the eigenvector is returned in *ifail*.
If *jobz* = 'N', then *z* is not referenced.
Note: you must ensure that at least $\max(1, m)$ columns are supplied in the array *z*; if *range* = 'V', the exact value of *m* is not known in advance and an upper bound must be used.
- ab* On exit, this array is overwritten by the values generated during the reduction to tridiagonal form. If *uplo* = 'U', the first superdiagonal and the diagonal of the tridiagonal matrix *T* are returned in rows *kd* and *kd*+1 of *ab*, and if *uplo* = 'L', the diagonal and first subdiagonal of *T* are returned in the first two rows of *ab*.
- ifail* **INTEGER**.
Array, **DIMENSION** at least $\max(1, n)$.
If *jobz* = 'V', then if *info* = 0, the first *m* elements of

ifail are zero; if *info* > 0, the *ifail* contains the indices the eigenvectors that failed to converge.

If *jobz* = 'N', then *ifail* is not referenced.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

If *info* = *i*, then *i* eigenvectors failed to converge; their indices are stored in the array *ifail*.

Application Notes

An approximate eigenvalue is accepted as converged when it is determined to lie in an interval [a,b] of width less than or equal to

$abstol + \epsilon * \max(|a|, |b|)$, where ϵ is the machine precision. If *abstol* is less than or equal to zero, then $\epsilon * \|T\|_1$ will be used in its place, where *T* is the tridiagonal matrix obtained by reducing *A* to tridiagonal form.

Eigenvalues will be computed most accurately when *abstol* is set to twice the underflow threshold $2 * ?lamch('S')$, not zero. If this routine returns with *info* > 0, indicating that some eigenvectors did not converge, try setting *abstol* to $2 * ?lamch('S')$.

?hbev

Computes selected eigenvalues and, optionally, eigenvectors of a Hermitian band matrix.

```
call chbev ( jobz, range, uplo, n, kd, ab, ldab, q, ldq, vl, vu, il,
            iu, abstol, m, w, z, ldz, work, rwork, iwork, ifail, info)
call zhbev ( jobz, range, uplo, n, kd, ab, ldab, q, ldq, vl, vu, il,
            iu, abstol, m, w, z, ldz, work, rwork, iwork, ifail, info)
```

Discussion

This routine computes selected eigenvalues and, optionally, eigenvectors of a complex Hermitian band matrix A . Eigenvalues and eigenvectors can be selected by specifying either a range of values or a range of indices for the desired eigenvalues.

Input Parameters

jobz CHARACTER*1. Must be 'N' or 'V'.
 If *job* = 'N', then only eigenvalues are computed.
 If *job* = 'V', then eigenvalues and eigenvectors are computed.

range CHARACTER*1. Must be 'A' or 'V' or 'I'.
 If *range* = 'A', the routine computes all eigenvalues.
 If *range* = 'V', the routine computes eigenvalues λ_i in the half-open interval: $vl < \lambda_i \leq vu$.
 If *range* = 'I', the routine computes eigenvalues with indices *il* to *iu*.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', *ab* stores the upper triangular part of A .
 If *uplo* = 'L', *ab* stores the lower triangular part of A .

n INTEGER. The order of the matrix A ($n \geq 0$).

kd INTEGER. The number of super- or sub-diagonals in A ($kd \geq 0$).

<i>ab, work</i>	<p>COMPLEX for <i>chbev</i>x DOUBLE COMPLEX for <i>zhbev</i>x. Arrays: <i>ab (ldab, *)</i> is an array containing either upper or lower triangular part of the Hermitian matrix <i>A</i> (as specified by <i>uplo</i>) in band storage format. The second dimension of <i>ab</i> must be at least $\max(1, n)$. <i>work (*)</i> is a workspace array. The dimension of <i>work</i> must be at least $\max(1, n)$.</p>
<i>ldab</i>	<p>INTEGER. The leading dimension of <i>ab</i>; must be at least $kd + 1$.</p>
<i>vl, vu</i>	<p>REAL for <i>chbev</i>x DOUBLE PRECISION for <i>zhbev</i>x. If <i>range</i> = 'V', the lower and upper bounds of the interval to be searched for eigenvalues. Constraint: $vl < vu$. If <i>range</i> = 'A' or 'I', <i>vl</i> and <i>vu</i> are not referenced.</p>
<i>il, iu</i>	<p>INTEGER. If <i>range</i> = 'I', the indices in ascending order of the smallest and largest eigenvalues to be returned. Constraint: $1 \leq il \leq iu \leq n$, if $n > 0$; $il=1$ and $iu=0$ if $n = 0$. If <i>range</i> = 'A' or 'V', <i>il</i> and <i>iu</i> are not referenced.</p>
<i>abstol</i>	<p>REAL for <i>chbev</i>x DOUBLE PRECISION for <i>zhbev</i>x. The absolute error tolerance to which each eigenvalue is required. See <i>Application notes</i> for details on error tolerance.</p>
<i>ldq, ldz</i>	<p>INTEGER. The leading dimensions of the output arrays <i>q</i> and <i>z</i>, respectively. Constraints: $ldq \geq 1, ldz \geq 1$; If <i>jobz</i> = 'V', then $ldq \geq \max(1, n)$ and $ldz \geq \max(1, n)$.</p>

rwork REAL for *chbev*x
 DOUBLE PRECISION for *zhbev*x
 Workspace array, DIMENSION at least $\max(1, 7n)$.

iwork INTEGER.
 Workspace array, DIMENSION at least $\max(1, 5n)$.

Output Parameters

m INTEGER. The total number of eigenvalues found, $0 \leq m \leq n$. If *range* = 'A', $m = n$, and if *range* = 'I', $m = i_u - i_l + 1$.

w REAL for *chbev*x
 DOUBLE PRECISION for *zhbev*x
 Array, DIMENSION at least $\max(1, n)$.
 The first *m* elements contain the selected eigenvalues of the matrix *A* in ascending order.

z COMPLEX for *chbev*x
 DOUBLE COMPLEX for *zhbev*x.
 Array *z*(*ldz*,*) . The second dimension of *z* must be at least $\max(1, m)$.
 If *jobz* = 'V', then if *info* = 0, the first *m* columns of *z* contain the orthonormal eigenvectors of the matrix *A* corresponding to the selected eigenvalues, with the *i*-th column of *z* holding the eigenvector associated with *w*(*i*). If an eigenvector fails to converge, then that column of *z* contains the latest approximation to the eigenvector, and the index of the eigenvector is returned in *ifail*.
 If *jobz* = 'N', then *z* is not referenced.
 Note: you must ensure that at least $\max(1, m)$ columns are supplied in the array *z* ; if *range* = 'V', the exact value of *m* is not known in advance and an upper bound must be used.

ab On exit, this array is overwritten by the values generated during the reduction to tridiagonal form. If *uplo* = 'U', the first superdiagonal and the diagonal of the

tridiagonal matrix T are returned in rows kd and $kd+1$ of ab , and if $uplo = 'L'$, the diagonal and first subdiagonal of T are returned in the first two rows of ab .

ifail

INTEGER.

Array, DIMENSION at least $\max(1, n)$.

If $jobz = 'V'$, then if $info = 0$, the first m elements of *ifail* are zero; if $info > 0$, the *ifail* contains the indices of the eigenvectors that failed to converge.

If $jobz = 'N'$, then *ifail* is not referenced.

info

INTEGER.

If $info = 0$, the execution is successful.

If $info = -i$, the i th parameter had an illegal value.

If $info = i$, then i eigenvectors failed to converge; their indices are stored in the array *ifail*.

Application Notes

An approximate eigenvalue is accepted as converged when it is determined to lie in an interval $[a,b]$ of width less than or equal to

$abstol + \varepsilon * \max(|a|,|b|)$, where ε is the machine precision. If *abstol* is less than or equal to zero, then $\varepsilon * \|T\|_1$ will be used in its place, where T is the tridiagonal matrix obtained by reducing A to tridiagonal form.

Eigenvalues will be computed most accurately when *abstol* is set to twice the underflow threshold $2 * ?lamch('S')$, not zero. If this routine returns with $info > 0$, indicating that some eigenvectors did not converge, try setting *abstol* to $2 * ?lamch('S')$.

?stev

Computes all eigenvalues and, optionally, eigenvectors of a real symmetric tridiagonal matrix.

```
call sstev (jobz, n, d, e, z, ldz, work, info)
call dstev (jobz, n, d, e, z, ldz, work, info)
```

Discussion

This routine computes all eigenvalues and, optionally, eigenvectors of a real symmetric tridiagonal matrix A .

Input Parameters

jobz CHARACTER*1. Must be 'N' or 'V'.
 If *jobz* = 'N', then only eigenvalues are computed.
 If *jobz* = 'V', then eigenvalues and eigenvectors are computed.

n INTEGER. The order of the matrix A ($n \geq 0$).

d, *e*, *work* REAL for *sstev*
 DOUBLE PRECISION for *dstev*.

Arrays:
d(*) contains the n diagonal elements of the tridiagonal matrix A .
 The dimension of *d* must be at least $\max(1, n)$.
e(*) contains the $n-1$ subdiagonal elements of the tridiagonal matrix A .
 The dimension of *e* must be at least $\max(1, n)$. The n th element of this array is used as workspace.
work(*) is a workspace array.
 The dimension of *work* must be at least $\max(1, 2n-2)$.
 If *jobz* = 'N', *work* is not referenced.

ldz **INTEGER**. The leading dimension of the output array *z*;
ldz ≥ 1. If *jobz* = 'V' then *ldz* ≥ max(1, *n*).

Output Parameters

d On exit, if *info* = 0, contains the eigenvalues of the matrix *A* in ascending order.

z **REAL** for *sstev*
DOUBLE PRECISION for *dstev*
Array, **DIMENSION** (*ldz*, *).
The second dimension of *z* must be at least max(1, *n*).
If *jobz* = 'V', then if *info* = 0, *z* contains the orthonormal eigenvectors of the matrix *A*, with the *i*-th column of *z* holding the eigenvector associated with the eigenvalue returned in *d*(*i*).
If *jobz* = 'N', then *z* is not referenced.

e On exit, this array is overwritten with intermediate results.

info **INTEGER**.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i*, then the algorithm failed to converge;
i elements of *e* did not converge to zero.

?stevd

Computes all eigenvalues and (optionally) all eigenvectors of a real symmetric tridiagonal matrix using divide and conquer algorithm.

```
call sstevd (job, n, d, e, z, ldz, work, lwork, iwork, liwork, info)
call dstevd (job, n, d, e, z, ldz, work, lwork, iwork, liwork, info)
```

Discussion

This routine computes all the eigenvalues, and optionally all the eigenvectors, of a real symmetric tridiagonal matrix T . In other words, the routine can compute the spectral factorization of T as: $T = Z\Lambda Z^T$.

Here Λ is a diagonal matrix whose diagonal elements are the eigenvalues λ_i , and Z is the orthogonal matrix whose columns are the eigenvectors z_i . Thus,

$$Tz_i = \lambda_i z_i \text{ for } i = 1, 2, \dots, n.$$

If the eigenvectors are requested, then this routine uses a divide and conquer algorithm to compute eigenvalues and eigenvectors. However, if only eigenvalues are required, then it uses the Pal-Walker-Kahan variant of the QL or QR algorithm.

There is no complex analogue of this routine.

Input Parameters

job CHARACTER*1. Must be 'N' or 'V'.
 If **job** = 'N', then only eigenvalues are computed.
 If **job** = 'V', then eigenvalues and eigenvectors are computed.

n INTEGER. The order of the matrix T ($n \geq 0$).

d, e, work REAL for **ssstevd**
 DOUBLE PRECISION for **dstevd**.

Arrays:

	<p>$d(*)$ contains the n diagonal elements of the tridiagonal matrix T. The dimension of d must be at least $\max(1, n)$.</p> <p>$e(*)$ contains the $n-1$ off-diagonal elements of T. The dimension of e must be at least $\max(1, n)$. The nth element of this array is used as workspace.</p> <p>$work(*)$ is a workspace array. The dimension of $work$ must be at least $lwork$.</p>
ldz	<p>INTEGER. The leading dimension of the output array z. Constraints: $ldz \geq 1$ if $job = 'N'$; $ldz \geq \max(1, n)$ if $job = 'V'$.</p>
$lwork$	<p>INTEGER. The dimension of the array $work$. Constraints: if $job = 'N'$ or $n \leq 1$, then $lwork \geq 1$; if $job = 'V'$ and $n > 1$, then $lwork \geq 2n^2 + (3+2k) * n + 1$, where k is the smallest integer which satisfies $2^k \geq n$.</p>
$iwork$	<p>INTEGER. Workspace array, DIMENSION at least $liwork$.</p>
$liwork$	<p>INTEGER. The dimension of the array $iwork$. Constraints: if $job = 'N'$ or $n \leq 1$, then $liwork \geq 1$; if $job = 'V'$ and $n > 1$, then $liwork \geq 5n+2$.</p>

Output Parameters

d	<p>On exit, if $info = 0$, contains the eigenvalues of the matrix T in ascending order. See also $info$.</p>
z	<p>REAL for <code>sstevd</code> DOUBLE PRECISION for <code>dstevd</code> Array, DIMENSION ($ldz, *$). The second dimension of z must be: at least 1 if $job = 'N'$; at least $\max(1, n)$ if $job = 'V'$.</p>

	If $job = 'V'$, then this array is overwritten by the orthogonal matrix Z which contains the eigenvectors of T . If $job = 'N'$, then z is not referenced.
e	On exit, this array is overwritten with intermediate results.
$work(1)$	On exit, if $lwork > 0$, then $work(1)$ returns the required minimal size of $lwork$.
$iwork(1)$	On exit, if $liwork > 0$, then $iwork(1)$ returns the required minimal size of $liwork$.
$info$	<p>INTEGER.</p> <p>If $info = 0$, the execution is successful.</p> <p>If $info = i$, then the algorithm failed to converge; i indicates the number of elements of an intermediate tridiagonal form which did not converge to zero.</p> <p>If $info = -i$, the ith parameter had an illegal value.</p>

Application Notes

The computed eigenvalues and eigenvectors are exact for a matrix $T + E$ such that $\|E\|_2 = O(\epsilon) \|T\|_2$, where ϵ is the machine precision.

If λ_i is an exact eigenvalue, and μ_i is the corresponding computed value, then

$$|\mu_i - \lambda_i| \leq c(n)\epsilon \|T\|_2$$

where $c(n)$ is a modestly increasing function of n .

If z_i is the corresponding exact eigenvector, and w_i is the corresponding computed vector, then the angle $\theta(z_i, w_i)$ between them is bounded as follows:

$$\theta(z_i, w_i) \leq c(n)\epsilon \|T\|_2 / \min_{i \neq j} |\lambda_i - \lambda_j|.$$

Thus the accuracy of a computed eigenvector depends on the gap between its eigenvalue and all the other eigenvalues.

?stevx

Computes selected eigenvalues and eigenvectors of a real symmetric tridiagonal matrix.

```
call sstevx ( jobz, range, n, d, e, vl, vu, il, iu, abstol, m, w, z,
             ldz, work, iwork, ifail, info)
call dstevx ( jobz, range, n, d, e, vl, vu, il, iu, abstol, m, w, z,
             ldz, work, iwork, ifail, info)
```

Discussion

This routine computes selected eigenvalues and, optionally, eigenvectors of a real symmetric tridiagonal matrix A . Eigenvalues and eigenvectors can be selected by specifying either a range of values or a range of indices for the desired eigenvalues.

Input Parameters

<i>jobz</i>	CHARACTER*1. Must be 'N' or 'V'. If <i>job</i> = 'N', then only eigenvalues are computed. If <i>job</i> = 'V', then eigenvalues and eigenvectors are computed.
<i>range</i>	CHARACTER*1. Must be 'A' or 'V' or 'I'. If <i>range</i> = 'A', the routine computes all eigenvalues. If <i>range</i> = 'V', the routine computes eigenvalues λ_i in the half-open interval: $vl < \lambda_i \leq vu$. If <i>range</i> = 'I', the routine computes eigenvalues with indices <i>il</i> to <i>iu</i> .
<i>n</i>	INTEGER. The order of the matrix A ($n \geq 0$).
<i>d, e, work</i>	REAL for sstevx DOUBLE PRECISION for dstevx. Arrays:

$d(*)$ contains the n diagonal elements of the tridiagonal matrix A .
The dimension of d must be at least $\max(1, n)$.

$e(*)$ contains the $n-1$ subdiagonal elements of A .
The dimension of e must be at least $\max(1, n)$. The n th element of this array is used as workspace.

$work(*)$ is a workspace array.
The dimension of $work$ must be at least $\max(1, 5n)$.

vl, vu REAL for `sstevx`
DOUBLE PRECISION for `dstevx`.
If $range = 'V'$, the lower and upper bounds of the interval to be searched for eigenvalues.
Constraint: $vl < vu$.
If $range = 'A'$ or $'I'$, vl and vu are not referenced.

il, iu INTEGER.
If $range = 'I'$, the indices in ascending order of the smallest and largest eigenvalues to be returned.
Constraint: $1 \leq il \leq iu \leq n$, if $n > 0$; $il=1$ and $iu=0$ if $n = 0$.
If $range = 'A'$ or $'V'$, il and iu are not referenced.

$abstol$ REAL for `sstevx`
DOUBLE PRECISION for `dstevx`.
The absolute error tolerance to which each eigenvalue is required. See *Application notes* for details on error tolerance.

ldz INTEGER. The leading dimensions of the output array z ;
 $ldz \geq 1$. If $jobz = 'V'$, then $ldz \geq \max(1, n)$.

$iwork$ INTEGER.
Workspace array, DIMENSION at least $\max(1, 5n)$.

Output Parameters

<i>m</i>	INTEGER. The total number of eigenvalues found, $0 \leq m \leq n$. If <i>range</i> = 'A', $m = n$, and if <i>range</i> = 'I', $m = i_u - i_l + 1$.
<i>w</i> , <i>z</i>	REAL for <i>sstevx</i> DOUBLE PRECISION for <i>dstevx</i> . Arrays: <i>w</i> (*), DIMENSION at least $\max(1, n)$. The first <i>m</i> elements of <i>w</i> contain the selected eigenvalues of the matrix <i>A</i> in ascending order. <i>z</i> (<i>ldz</i> , *). The second dimension of <i>z</i> must be at least $\max(1, m)$. If <i>jobz</i> = 'V', then if <i>info</i> = 0, the first <i>m</i> columns of <i>z</i> contain the orthonormal eigenvectors of the matrix <i>A</i> corresponding to the selected eigenvalues, with the <i>i</i> -th column of <i>z</i> holding the eigenvector associated with <i>w</i> (<i>i</i>). If an eigenvector fails to converge, then that column of <i>z</i> contains the latest approximation to the eigenvector, and the index of the eigenvector is returned in <i>ifail</i> . If <i>jobz</i> = 'N', then <i>z</i> is not referenced. Note: you must ensure that at least $\max(1, m)$ columns are supplied in the array <i>z</i> ; if <i>range</i> = 'V', the exact value of <i>m</i> is not known in advance and an upper bound must be used.
<i>d</i> , <i>e</i>	On exit, these arrays may be multiplied by a constant factor chosen to avoid overflow or underflow in computing the eigenvalues.
<i>ifail</i>	INTEGER. Array, DIMENSION at least $\max(1, n)$. If <i>jobz</i> = 'V', then if <i>info</i> = 0, the first <i>m</i> elements of <i>ifail</i> are zero; if <i>info</i> > 0, the <i>ifail</i> contains the indices of the eigenvectors that failed to converge. If <i>jobz</i> = 'N', then <i>ifail</i> is not referenced.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

If *info* = *i*, then *i* eigenvectors failed to converge; their indices are stored in the array *ifail*.

Application Notes

An approximate eigenvalue is accepted as converged when it is determined to lie in an interval [a,b] of width less than or equal to

$abstol + \epsilon * \max(|a|, |b|)$, where ϵ is the machine precision. If *abstol* is less than or equal to zero, then $\epsilon * \|A\|_1$ will be used in its place.

Eigenvalues will be computed most accurately when *abstol* is set to twice the underflow threshold $2 * ?lamch('S')$, not zero. If this routine returns with *info* > 0, indicating that some eigenvectors did not converge, try setting *abstol* to $2 * ?lamch('S')$.

?stevr

Computes selected eigenvalues and, optionally, eigenvectors of a real symmetric tridiagonal matrix using the Relatively Robust Representations.

```
call sstevr ( jobz, range, n, d, e, vl, vu, il, iu, abstol, m, w, z,
             ldz, isuppz, work, lwork, iwork, liwork, info)
call dstevr ( jobz, range, n, d, e, vl, vu, il, iu, abstol, m, w, z,
             ldz, isuppz, work, lwork, iwork, liwork, info)
```

Discussion

This routine computes selected eigenvalues and, optionally, eigenvectors of a real symmetric tridiagonal matrix T . Eigenvalues and eigenvectors can be selected by specifying either a range of values or a range of indices for the desired eigenvalues.

Whenever possible, `?stevr` calls `sstegr/dstegr` to compute the eigenspectrum using Relatively Robust Representations. `?stegr` computes eigenvalues by the *dqds* algorithm, while orthogonal eigenvectors are computed from various “good” LDL^T representations (also known as Relatively Robust Representations). Gram-Schmidt orthogonalization is avoided as far as possible. More specifically, the various steps of the algorithm are as follows. For the i -th unreduced block of T ,

- (a) Compute $T - \mathbf{Q}_i = L_i D_i L_i^T$, such that $L_i D_i L_i^T$ is a relatively robust representation;
- (b) Compute the eigenvalues, λ_j , of $L_i D_i L_i^T$ to high relative accuracy by the *dqds* algorithm;
- (c) If there is a cluster of close eigenvalues, “choose” \mathbf{Q}_i close to the cluster, and go to step (a);
- (d) Given the approximate eigenvalue λ_j of $L_i D_i L_i^T$, compute the corresponding eigenvector by forming a rank-revealing twisted factorization.

The desired accuracy of the output can be specified by the input parameter *abstol*.

The routine `?stevr` calls `sstegr/dstegr` when the full spectrum is requested on machines which conform to the IEEE-754 floating point standard. `?stevr` calls `sstebz/dstebz` and `sstein/dstein` on non-IEEE machines and when partial spectrum requests are made.

Input Parameters

jobz CHARACTER*1. Must be 'N' or 'V'.
 If *jobz* = 'N', then only eigenvalues are computed.
 If *jobz* = 'V', then eigenvalues and eigenvectors are computed.

range CHARACTER*1. Must be 'A' or 'V' or 'I'.
 If *range* = 'A', the routine computes all eigenvalues.
 If *range* = 'V', the routine computes eigenvalues λ_i in the half-open interval: $vl < \lambda_i \leq vu$.
 If *range* = 'I', the routine computes eigenvalues with indices *il* to *iu*.
 For *range* = 'V' or 'I' and $iu - il < n - 1$, `sstebz/dstebz` and `sstein/dstein` are called.

n INTEGER. The order of the matrix *T* ($n \geq 0$).

d, *e*, *work* REAL for `sstevr`
 DOUBLE PRECISION for `dstevr`.
 Arrays:
d(*) contains the *n* diagonal elements of the tridiagonal matrix *T*.
 The dimension of *d* must be at least $\max(1, n)$.
e(*) contains the *n*-1 subdiagonal elements of *A*.
 The dimension of *e* must be at least $\max(1, n)$. The *n*th element of this array is used as workspace.
work(*lwork*) is a workspace array.

<i>vl, vu</i>	<p>REAL for <i>sstevr</i> DOUBLE PRECISION for <i>dstevr</i>. If <i>range</i> = 'V', the lower and upper bounds of the interval to be searched for eigenvalues. Constraint: <i>vl</i> < <i>vu</i>. If <i>range</i> = 'A' or 'I', <i>vl</i> and <i>vu</i> are not referenced.</p>
<i>il, iu</i>	<p>INTEGER. If <i>range</i> = 'I', the indices in ascending order of the smallest and largest eigenvalues to be returned. Constraint: $1 \leq il \leq iu \leq n$, if $n > 0$; $il=1$ and $iu=0$ if $n = 0$. If <i>range</i> = 'A' or 'V', <i>il</i> and <i>iu</i> are not referenced.</p>
<i>abstol</i>	<p>REAL for <i>ssyevr</i> DOUBLE PRECISION for <i>dsyevr</i>. The absolute error tolerance to which each eigenvalue/eigenvector is required. If <i>jobz</i> = 'V', the eigenvalues and eigenvectors output have residual norms bounded by <i>abstol</i>, and the dot products between different eigenvectors are bounded by <i>abstol</i>. If $abstol < n\epsilon$, then $n\epsilon$ will be used in its place, where ϵ is the machine precision. The eigenvalues are computed to an accuracy of ϵ irrespective of <i>abstol</i>. If high relative accuracy is important, set <i>abstol</i> to <i>?lamch</i>('S').</p>
<i>ldz</i>	<p>INTEGER. The leading dimension of the output array <i>z</i>. Constraints: $ldz \geq 1$ if <i>jobz</i> = 'N'; $ldz \geq \max(1, n)$ if <i>jobz</i> = 'V'.</p>
<i>lwork</i>	<p>INTEGER. The dimension of the array <i>work</i>. Constraint: $lwork \geq \max(1, 20n)$.</p>
<i>iwork</i>	<p>INTEGER. Workspace array, DIMENSION (<i>liwork</i>).</p>
<i>liwork</i>	<p>INTEGER. The dimension of the array <i>iwork</i>, $lwork \geq \max(1, 10n)$.</p>

Output Parameters

- m* **INTEGER**. The total number of eigenvalues found, $0 \leq m \leq n$. If *range* = 'A', $m = n$, and if *range* = 'I', $m = iu - il + 1$.
- w*, *z* **REAL** for *sstevr*
DOUBLE PRECISION for *dstevr*.
 Arrays:
w(*), **DIMENSION** at least $\max(1, n)$.
 The first *m* elements of *w* contain the selected eigenvalues of the matrix *T* in ascending order.
z(*ldz*, *) . The second dimension of *z* must be at least $\max(1, m)$.
 If *jobz* = 'V', then if *info* = 0, the first *m* columns of *z* contain the orthonormal eigenvectors of the matrix *T* corresponding to the selected eigenvalues, with the *i*-th column of *z* holding the eigenvector associated with *w*(*i*).
 If *jobz* = 'N', then *z* is not referenced.
 Note: you must ensure that at least $\max(1, m)$ columns are supplied in the array *z*; if *range* = 'V', the exact value of *m* is not known in advance and an upper bound must be used.
- d*, *e* On exit, these arrays may be multiplied by a constant factor chosen to avoid overflow or underflow in computing the eigenvalues.
- isuppz* **INTEGER**.
 Array, **DIMENSION** at least $2 * \max(1, m)$.
 The support of the eigenvectors in *z*, i.e., the indices indicating the nonzero elements in *z*. The *i*-th eigenvector is nonzero only in elements *isuppz*(2*i*-1) through *isuppz*(2*i*).
 Implemented only for *range* = 'A' or 'I' and $iu - il = n - 1$.
- work*(1) On exit, if *info* = 0, then *work*(1) returns the required minimal size of *lwork*.

iwork(1) On exit, if *info* = 0, then *iwork(1)* returns the required minimal size of *liwork*.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i*, an internal error has occurred.

Application Notes

Normal execution of the routine `?stegr` may create NaNs and infinities and hence may abort due to a floating point exception in environments which do not handle NaNs and infinities in the IEEE standard default manner.

Nonsymmetric Eigenproblems

This section describes LAPACK driver routines used for solving nonsymmetric eigenproblems. See also [computational routines](#) that can be called to solve these problems.

[Table 5-12](#) lists routines described in more detail below.

Table 5-11 Driver Routines for Solving Nonsymmetric Eigenproblems

Routine Name	Operation performed
?gees	Computes the eigenvalues and Schur factorization of a general matrix, and orders the factorization so that selected eigenvalues are at the top left of the Schur form.
?geesx	Computes the eigenvalues and Schur factorization of a general matrix, orders the factorization and computes reciprocal condition numbers.
?geev	Computes the eigenvalues and left and right eigenvectors of a general matrix.
?geevx	Computes the eigenvalues and left and right eigenvectors of a general matrix, with preliminary matrix balancing, and computes reciprocal condition numbers for the eigenvalues and right eigenvectors.

?gees

Computes the eigenvalues and Schur factorization of a general matrix, and orders the factorization so that selected eigenvalues are at the top left of the Schur form.

```
call sgees ( jobvs, sort, select, n, a, lda, sdim, wr, wi, vs, ldvs,
            work, lwork, bwork, info)
call dgees ( jobvs, sort, select, n, a, lda, sdim, wr, wi, vs, ldvs,
            work, lwork, bwork, info)
call cgees ( jobvs, sort, select, n, a, lda, sdim, w, vs, ldvs,
            work, lwork, rwork, bwork, info)
call zgees ( jobvs, sort, select, n, a, lda, sdim, w, vs, ldvs,
            work, lwork, rwork, bwork, info)
```


Discussion

This routine computes for an n -by- n real/complex nonsymmetric matrix A , the eigenvalues, the real Schur form T , and, optionally, the matrix of Schur vectors Z . This gives the Schur factorization $A = ZTZ^H$.

Optionally, it also orders the eigenvalues on the diagonal of the real-Schur/Schur form so that selected eigenvalues are at the top left. The leading columns of Z then form an orthonormal basis for the invariant subspace corresponding to the selected eigenvalues.

A real matrix is in real-Schur form if it is upper quasi-triangular with 1-by-1 and 2-by-2 blocks. 2-by-2 blocks will be standardized in the form

$$\begin{pmatrix} a & b \\ c & a \end{pmatrix}$$

where $b*c < 0$. The eigenvalues of such a block are $a \pm \sqrt{bc}$.

A complex matrix is in Schur form if it is upper triangular.

Input Parameters

- jobvs* CHARACTER*1. Must be 'N' or 'V'.
 If *jobvs* = 'N', then Schur vectors are not computed.
 If *jobvs* = 'V', then Schur vectors are computed.
- sort* CHARACTER*1. Must be 'N' or 'S'.
 Specifies whether or not to order the eigenvalues on the diagonal of the Schur form.
 If *sort* = 'N', then eigenvalues are not ordered.
 If *sort* = 'S', eigenvalues are ordered (see *select*).
- select* LOGICAL FUNCTION of two REAL arguments for real flavors.
 LOGICAL FUNCTION of one COMPLEX argument for complex flavors.
select must be declared EXTERNAL in the calling subroutine.
 If *sort* = 'S', *select* is used to select eigenvalues to sort to the top left of the Schur form.
 If *sort* = 'N', *select* is not referenced.

For real flavors:

An eigenvalue $wr(j) + \sqrt{-1} * wi(j)$ is selected if `select(wr(j), wi(j))` is true; that is, if either one of a complex conjugate pair of eigenvalues is selected, then both complex eigenvalues are selected. Note that a selected complex eigenvalue may no longer satisfy `select(wr(j), wi(j)) = .TRUE.` after ordering, since ordering may change the value of complex eigenvalues (especially if the eigenvalue is ill-conditioned); in this case `info` may be set to `n+2` (see `info` below).

For complex flavors:

An eigenvalue $w(j)$ is selected if `select(w(j))` is true.

`n` `INTEGER`. The order of the matrix A ($n \geq 0$).

`a, work` `REAL` for `sgees`
 `DOUBLE PRECISION` for `dgees`
 `COMPLEX` for `cgees`
 `DOUBLE COMPLEX` for `zgees`.

Arrays:
`a(lda,*)` is an array containing the n -by- n matrix A .
 The second dimension of `a` must be at least $\max(1, n)$.

`work(lwork)` is a workspace array.

`lda` `INTEGER`. The first dimension of the array `a`.
 Must be at least $\max(1, n)$.

`ldvs` `INTEGER`. The leading dimension of the output array `vs`.
 Constraints:
`ldvs` ≥ 1 ;
`ldvs` $\geq \max(1, n)$ if `jobvs = 'V'`.

`lwork` `INTEGER`. The dimension of the array `work`.
 Constraint:
`lwork` $\geq \max(1, 3n)$ for real flavors;
`lwork` $\geq \max(1, 2n)$ for complex flavors.

`rwork` `REAL` for `cgees`
 `DOUBLE PRECISION` for `zgees`
 Workspace array, `DIMENSION` at least $\max(1, n)$. Used
 in complex flavors only.

bwork LOGICAL.
 Workspace array, DIMENSION at least $\max(1, n)$. Not referenced if *sort* = 'N'.

Output Parameters

a On exit, this array is overwritten by the real-Schur/Schur form T .

sdim INTEGER.
 If *sort* = 'N', *sdim* = 0.
 If *sort* = 'S', *sdim* is equal to the number of eigenvalues (after sorting) for which *select* is true. Note that for real flavors complex conjugate pairs for which *select* is true for either eigenvalue count as 2.

wr, wi REAL for *sgees*
 DOUBLE PRECISION for *dgees*
 Arrays, DIMENSION at least $\max(1, n)$ each. Contain the real and imaginary parts, respectively, of the computed eigenvalues, in the same order that they appear on the diagonal of the output real-Schur form T . Complex conjugate pairs of eigenvalues appear consecutively with the eigenvalue having positive imaginary part first.

w COMPLEX for *cgees*
 DOUBLE COMPLEX for *zgees*.
 Array, DIMENSION at least $\max(1, n)$. Contains the computed eigenvalues. The eigenvalues are stored in the same order as they appear on the diagonal of the output Schur form T .

vs REAL for *sgees*
 DOUBLE PRECISION for *dgees*
 COMPLEX for *cgees*
 DOUBLE COMPLEX for *zgees*.
 Array *vs*(*ldvs*, *); the second dimension of *vs* must be at least $\max(1, n)$.

If *jobvs* = 'V', *vs* contains the orthogonal/unitary matrix *Z* of Schur vectors.

If *jobvs* = 'N', *vs* is not referenced.

work(1) On exit, if *info* = 0, then *work(1)* returns the required minimal size of *lwork*.

info INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

If *info* = *i*, and
i ≤ *n* :

the QR algorithm failed to compute all the eigenvalues; elements 1:*ilo*-1 and *i*+1:*n* of *wr* and *wi* (for real flavors) or *w* (for complex flavors) contain those eigenvalues which have converged; if *jobvs* = 'V', *vs* contains the matrix which reduces *A* to its partially converged Schur form;

i = *n*+1 :

the eigenvalues could not be reordered because some eigenvalues were too close to separate (the problem is very ill-conditioned);

i = *n*+2 :

after reordering, roundoff changed values of some complex eigenvalues so that leading eigenvalues in the Schur form no longer satisfy *select* = .TRUE.. This could also be caused by underflow due to scaling.

Application Notes

If you are in doubt how much workspace to supply for the array *work*, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

?geesx

Computes the eigenvalues and Schur factorization of a general matrix, orders the factorization and computes reciprocal condition numbers.

```
call sgeesx(jobvs, sort, select, sense, n, a, lda, sdim, wr, wi, vs,
           ldvs, rconde, rcondv, work, lwork, iwork, liwork, bwork, info)
call dgeesx(jobvs, sort, select, sense, n, a, lda, sdim, wr, wi, vs,
           ldvs, rconde, rcondv, work, lwork, iwork, liwork, bwork, info)
call cgeesx(jobvs, sort, select, sense, n, a, lda, sdim, w, vs,
           ldvs, rconde, rcondv, work, lwork, rwork, bwork, info)
call zgeesx(jobvs, sort, select, sense, n, a, lda, sdim, w, vs,
           ldvs, rconde, rcondv, work, lwork, rwork, bwork, info)
```

Discussion

This routine computes for an n -by- n real/complex nonsymmetric matrix A , the eigenvalues, the real-Schur/Schur form T , and, optionally, the matrix of Schur vectors Z . This gives the Schur factorization $A = ZTZ^H$.

Optionally, it also orders the eigenvalues on the diagonal of the real-Schur/Schur form so that selected eigenvalues are at the top left; computes a reciprocal condition number for the average of the selected eigenvalues (*rconde*); and computes a reciprocal condition number for the right invariant subspace corresponding to the selected eigenvalues (*rcondv*). The leading columns of Z form an orthonormal basis for this invariant subspace.

For further explanation of the reciprocal condition numbers *rconde* and *rcondv*, see [LUG], Section 4.10 (where these quantities are called *s* and *sep* respectively).

A real matrix is in real-Schur form if it is upper quasi-triangular with 1-by-1 and 2-by-2 blocks. 2-by-2 blocks will be standardized in the form

$$\begin{pmatrix} a & b \\ c & a \end{pmatrix}$$

where $b*c < 0$. The eigenvalues of such a block are $a \pm \sqrt{bc}$.
 A complex matrix is in Schur form if it is upper triangular.

Input Parameters

jobvs CHARACTER*1. Must be 'N' or 'V'.
 If *jobvs* = 'N', then Schur vectors are not computed.
 If *jobvs* = 'V', then Schur vectors are computed.

sort CHARACTER*1. Must be 'N' or 'S'.
 Specifies whether or not to order the eigenvalues on the diagonal of the Schur form.
 If *sort* = 'N', then eigenvalues are not ordered.
 If *sort* = 'S', eigenvalues are ordered (see *select*).

select LOGICAL FUNCTION of two REAL arguments for real flavors.
 LOGICAL FUNCTION of one COMPLEX argument for complex flavors.
select must be declared EXTERNAL in the calling subroutine.
 If *sort* = 'S', *select* is used to select eigenvalues to sort to the top left of the Schur form.
 If *sort* = 'N', *select* is not referenced.
 For real flavors:
 An eigenvalue $wr(j) + \sqrt{-1} * wi(j)$ is selected if *select*(*wr*(j), *wi*(j)) is true; that is, if either one of a complex conjugate pair of eigenvalues is selected, then both complex eigenvalues are selected. Note that a selected complex eigenvalue may no longer satisfy *select*(*wr*(j), *wi*(j)) = .TRUE. after ordering, since ordering may change the value of complex eigenvalues (especially if the eigenvalue is ill-conditioned); in this case *info* may be set to *n*+2 (see *info* below).
 For complex flavors:
 An eigenvalue *w*(j) is selected if *select*(*w*(j)) is true.

sense CHARACTER*1. Must be 'N', 'E', 'V', or 'B'.
Determines which reciprocal condition number are computed.
If *sense* = 'N', none are computed;
If *sense* = 'E', computed for average of selected eigenvalues only;
If *sense* = 'V', computed for selected right invariant subspace only;
If *sense* = 'B', computed for both.
If *sense* is 'E', 'V', or 'B', then *sort* must equal 'S'.

n INTEGER. The order of the matrix A ($n \geq 0$).

a, *work* REAL for sgeesx
DOUBLE PRECISION for dgeesx
COMPLEX for cgeesx
DOUBLE COMPLEX for zgeesx.
Arrays:
a(*lda*,*) is an array containing the *n*-by-*n* matrix A.
The second dimension of *a* must be at least max(1, *n*).
work(*lwork*) is a workspace array.

lda INTEGER. The first dimension of the array *a*.
Must be at least max(1, *n*).

ldvs INTEGER. The leading dimension of the output array *vs*.
Constraints:
ldvs ≥ 1 ;
ldvs $\geq \max(1, n)$ if *jobvs* = 'V'.

lwork INTEGER. The dimension of the array *work*.
Constraint:
lwork $\geq \max(1, 3n)$ for real flavors;
lwork $\geq \max(1, 2n)$ for complex flavors.
Also, if *sense* = 'E', 'V', or 'B', then
lwork $\geq n + 2 * sdim * (n - sdim)$ for real flavors;
lwork $\geq 2 * sdim * (n - sdim)$ for complex flavors;

where *sdim* is the number of selected eigenvalues computed by this routine. Note that $2 * \textit{sdim} * (n - \textit{sdim}) \leq n * n / 2$.

For good performance, *lwork* must generally be larger.

<i>iwork</i>	INTEGER. Workspace array, DIMENSION (<i>liwork</i>). Used in real flavors only. Not referenced if <i>sense</i> = 'N' or 'E'.
<i>liwork</i>	INTEGER. The dimension of the array <i>iwork</i> . Used in real flavors only. Constraint: $\textit{liwork} \geq 1$; if <i>sense</i> = 'V' or 'B', $\textit{liwork} \geq \textit{sdim} * (n - \textit{sdim})$.
<i>rwork</i>	REAL for <i>cgeesx</i> DOUBLE PRECISION for <i>zgeesx</i> Workspace array, DIMENSION at least $\max(1, n)$. Used in complex flavors only.
<i>bwork</i>	LOGICAL. Workspace array, DIMENSION at least $\max(1, n)$. Not referenced if <i>sort</i> = 'N'.

Output Parameters

<i>a</i>	On exit, this array is overwritten by the real-Schur/Schur form <i>T</i> .
<i>sdim</i>	INTEGER. If <i>sort</i> = 'N', <i>sdim</i> = 0. If <i>sort</i> = 'S', <i>sdim</i> is equal to the number of eigenvalues (after sorting) for which <i>select</i> is true. Note that for real flavors complex conjugate pairs for which <i>select</i> is true for either eigenvalue count as 2.
<i>wr, wi</i>	REAL for <i>sgeesx</i> DOUBLE PRECISION for <i>dgeesx</i> Arrays, DIMENSION at least $\max(1, n)$ each. Contain the real and imaginary parts, respectively, of the computed eigenvalues, in the same order that they appear on the diagonal of the output real-Schur form <i>T</i> .

Complex conjugate pairs of eigenvalues appear consecutively with the eigenvalue having positive imaginary part first.

<i>w</i>	<p>COMPLEX for <i>cgeesx</i> DOUBLE COMPLEX for <i>zgeesx</i>. Array, DIMENSION at least $\max(1, n)$. Contains the computed eigenvalues. The eigenvalues are stored in the same order as they appear on the diagonal of the output Schur form <i>T</i>.</p>
<i>vs</i>	<p>REAL for <i>sgeesx</i> DOUBLE PRECISION for <i>dgeesx</i> COMPLEX for <i>cgeesx</i> DOUBLE COMPLEX for <i>zgeesx</i>. Array <i>vs(ldvs, *)</i>; the second dimension of <i>vs</i> must be at least $\max(1, n)$. If <i>jobvs</i> = 'V', <i>vs</i> contains the orthogonal/unitary matrix <i>Z</i> of Schur vectors. If <i>jobvs</i> = 'N', <i>vs</i> is not referenced.</p>
<i>rconde, rcondv</i>	<p>REAL for single precision flavors DOUBLE PRECISION for double precision flavors. If <i>sense</i> = 'E' or 'B', <i>rconde</i> contains the reciprocal condition number for the average of the selected eigenvalues. If <i>sense</i> = 'N' or 'V', <i>rconde</i> is not referenced. If <i>sense</i> = 'V' or 'B', <i>rcondv</i> contains the reciprocal condition number for the selected right invariant subspace. If <i>sense</i> = 'N' or 'E', <i>rcondv</i> is not referenced.</p>
<i>work(1)</i>	<p>On exit, if <i>info</i> = 0, then <i>work(1)</i> returns the required minimal size of <i>lwork</i>.</p>
<i>info</i>	<p>INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = -<i>i</i>, the <i>i</i>th parameter had an illegal value.</p>

If $info = i$, and

$i \leq n$:

the *QR* algorithm failed to compute all the eigenvalues; elements 1:*ilo*-1 and *i*+1:*n* of *wr* and *wi* (for real flavors) or *w* (for complex flavors) contain those eigenvalues which have converged; if *jobvs* = 'V', *vs* contains the transformation which reduces *A* to its partially converged Schur form;

$i = n+1$:

the eigenvalues could not be reordered because some eigenvalues were too close to separate (the problem is very ill-conditioned);

$i = n+2$:

after reordering, roundoff changed values of some complex eigenvalues so that leading eigenvalues in the Schur form no longer satisfy *select* = *.TRUE.*. This could also be caused by underflow due to scaling.

Application Notes

If you are in doubt how much workspace to supply for the array *work*, use a generous value of *lwork* for the first run. On exit, examine *work*(1) and use this value for subsequent runs.

?ggev

Computes the eigenvalues and left and right eigenvectors of a general matrix.

```
call sgeev ( jobvl, jobvr, n, a, lda, wr, wi, vl, ldvl, vr, ldvr,
             work, lwork, info)
call dgeev ( jobvl, jobvr, n, a, lda, wr, wi, vl, ldvl, vr, ldvr,
             work, lwork, info)
call cgeev ( jobvl, jobvr, n, a, lda, w, vl, ldvl, vr, ldvr, work,
             lwork, rwork, info)
call zgeev ( jobvl, jobvr, n, a, lda, w, vl, ldvl, vr, ldvr, work,
             lwork, rwork, info)
```

Discussion

This routine computes for an n -by- n real/complex nonsymmetric matrix A , the eigenvalues and, optionally, the left and/or right eigenvectors. The right eigenvector $v(j)$ of A satisfies

$$A * v(j) = \lambda(j) * v(j)$$

where $\lambda(j)$ is its eigenvalue.

The left eigenvector $u(j)$ of A satisfies

$$u(j)^H * A = \lambda(j) * u(j)^H$$

where $u(j)^H$ denotes the conjugate transpose of $u(j)$.

The computed eigenvectors are normalized to have Euclidean norm equal to 1 and largest component real.

Input Parameters

jobvl CHARACTER*1. Must be 'N' or 'V'.
 If *jobvl* = 'N', then left eigenvectors of A are not computed.
 If *jobvl* = 'V', then left eigenvectors of A are computed.

jobvr CHARACTER*1. Must be 'N' or 'V'.
 If *jobvr* = 'N', then right eigenvectors of *A* are not computed.
 If *jobvr* = 'V', then right eigenvectors of *A* are computed.

n INTEGER. The order of the matrix *A* ($n \geq 0$).

a, *work* REAL for *sggev*
 DOUBLE PRECISION for *dgeev*
 COMPLEX for *cgeev*
 DOUBLE COMPLEX for *zgeev*.
 Arrays:
a(*lda*,*) is an array containing the *n*-by-*n* matrix *A*.
 The second dimension of *a* must be at least $\max(1, n)$.
work(*lwork*) is a workspace array.

lda INTEGER. The first dimension of the array *a*.
 Must be at least $\max(1, n)$.

ldvl, *ldvr* INTEGER. The leading dimensions of the output arrays *vl* and *vr*, respectively. Constraints:
 $ldvl \geq 1$; $ldvr \geq 1$.
 If *jobvl* = 'V', $ldvl \geq \max(1, n)$;
 If *jobvr* = 'V', $ldvr \geq \max(1, n)$.

lwork INTEGER. The dimension of the array *work*.
 Constraint:
 $lwork \geq \max(1, 3n)$, and if *jobvl* = 'V' or *jobvr* = 'V', $lwork \geq \max(1, 4n)$ (for real flavors);
 $lwork \geq \max(1, 2n)$ (for complex flavors).
 For good performance, *lwork* must generally be larger.

rwork REAL for *cgeev*
 DOUBLE PRECISION for *zgeev*
 Workspace array, DIMENSION at least $\max(1, 2n)$. Used in complex flavors only.

Output Parameters

<i>a</i>	On exit, this array is overwritten by intermediate results.
<i>wr, wi</i>	<p>REAL for <i>sgeev</i> DOUBLE PRECISION for <i>dgeev</i></p> <p>Arrays, DIMENSION at least max (1, <i>n</i>) each. Contain the real and imaginary parts, respectively, of the computed eigenvalues. Complex conjugate pairs of eigenvalues appear consecutively with the eigenvalue having positive imaginary part first.</p>
<i>w</i>	<p>COMPLEX for <i>cgeev</i> DOUBLE COMPLEX for <i>zgeev</i>.</p> <p>Array, DIMENSION at least max(1,<i>n</i>). Contains the computed eigenvalues.</p>
<i>v1, vr</i>	<p>REAL for <i>sgeev</i> DOUBLE PRECISION for <i>dgeev</i> COMPLEX for <i>cgeev</i> DOUBLE COMPLEX for <i>zgeev</i>.</p> <p>Arrays: <i>v1(ldv1, *)</i>; the second dimension of <i>v1</i> must be at least max(1, <i>n</i>). If <i>jobv1</i> = 'V', the left eigenvectors <i>u</i>(<i>j</i>) are stored one after another in the columns of <i>v1</i>, in the same order as their eigenvalues. If <i>jobv1</i> = 'N', <i>v1</i> is not referenced. <i>For real flavors:</i> If the <i>j</i>-th eigenvalue is real, then <i>u</i>(<i>j</i>) = <i>v1</i>(:,<i>j</i>), the <i>j</i>-th column of <i>v1</i>. If the <i>j</i>-th and (<i>j</i>+1)-st eigenvalues form a complex conjugate pair, then <i>u</i>(<i>j</i>) = <i>v1</i>(:,<i>j</i>) + <i>i</i>*<i>v1</i>(:,<i>j</i>+1) and <i>u</i>(<i>j</i>+1) = <i>v1</i>(:,<i>j</i>) - <i>i</i>*<i>v1</i>(:,<i>j</i>+1), where <i>i</i> = $\sqrt{-1}$. <i>For complex flavors:</i> <i>u</i>(<i>j</i>) = <i>v1</i>(:,<i>j</i>), the <i>j</i>-th column of <i>v1</i>. <i>vr(ldvr, *)</i>; the second dimension of <i>vr</i> must be at least max(1, <i>n</i>). If <i>jobvr</i> = 'V', the right eigenvectors <i>v</i>(<i>j</i>) are stored one after another in the columns of <i>vr</i>, in the same order as their eigenvalues. If <i>jobvr</i> = 'N', <i>vr</i> is not referenced.</p>

For real flavors:

If the j -th eigenvalue is real, then $v(j) = \mathbf{vr}(:,j)$, the j -th column of \mathbf{vr} . If the j -th and $(j+1)$ -st eigenvalues form a complex conjugate pair, then $v(j) = \mathbf{vr}(:,j) + i*\mathbf{vr}(:,j+1)$ and $v(j+1) = \mathbf{vr}(:,j) - i*\mathbf{vr}(:,j+1)$, where $i = \sqrt{-1}$.

For complex flavors:

$v(j) = \mathbf{vr}(:,j)$, the j -th column of \mathbf{vr} .

`work(1)` On exit, if `info = 0`, then `work(1)` returns the required minimal size of `lwork`.

`info` INTEGER.

If `info = 0`, the execution is successful.

If `info = -i`, the i th parameter had an illegal value.

If `info = i`, the QR algorithm failed to compute all the eigenvalues, and no eigenvectors have been computed; elements $i+1:n$ of `wr` and `wi` (for real flavors) or `w` (for complex flavors) contain those eigenvalues which have converged.

Application Notes

If you are in doubt how much workspace to supply for the array `work`, use a generous value of `lwork` for the first run. On exit, examine `work(1)` and use this value for subsequent runs.

?ggev

Computes the eigenvalues and left and right eigenvectors of a general matrix, with preliminary matrix balancing, and computes reciprocal condition numbers for the eigenvalues and right eigenvectors.

```
call sgev ( balanc, jobvl, jobvr, sense, n, a, lda, wr, wi, vl,
            ldvl, vr, ldvr, ilo, ihi, scale, abnrm, rconde,
            rcondv, work, lwork, iwork, info)
call dgev ( balanc, jobvl, jobvr, sense, n, a, lda, wr, wi, vl,
            ldvl, vr, ldvr, ilo, ihi, scale, abnrm, rconde,
            rcondv, work, lwork, iwork, info)
call cgev ( balanc, jobvl, jobvr, sense, n, a, lda, w, vl, ldvl,
            vr, ldvr, ilo, ihi, scale, abnrm, rconde, rcondv,
            work, lwork, rwork, info)
call zgev ( balanc, jobvl, jobvr, sense, n, a, lda, w, vl, ldvl,
            vr, ldvr, ilo, ihi, scale, abnrm, rconde, rcondv,
            work, lwork, rwork, info)
```

Discussion

This routine computes for an n -by- n real/complex nonsymmetric matrix A , the eigenvalues and, optionally, the left and/or right eigenvectors.

Optionally also, it computes a balancing transformation to improve the conditioning of the eigenvalues and eigenvectors (ilo , ihi , $scale$, and $abnrm$), reciprocal condition numbers for the eigenvalues ($rconde$), and reciprocal condition numbers for the right eigenvectors ($rcondv$).

The right eigenvector $v(j)$ of A satisfies

$$A * v(j) = \lambda(j) * v(j)$$

where $\lambda(j)$ is its eigenvalue.

The left eigenvector $u(j)$ of A satisfies

$$u(j)^H * A = \lambda(j) * u(j)^H$$

where $u(j)^H$ denotes the conjugate transpose of $u(j)$.

The computed eigenvectors are normalized to have Euclidean norm equal to 1 and largest component real.

Balancing a matrix means permuting the rows and columns to make it more nearly upper triangular, and applying a diagonal similarity transformation $DA D^{-1}$, where D is a diagonal matrix, to make its rows and columns closer in norm and the condition numbers of its eigenvalues and eigenvectors smaller. The computed reciprocal condition numbers correspond to the balanced matrix.

Permuting rows and columns will not change the condition numbers in exact arithmetic) but diagonal scaling will. For further explanation of balancing, see [[LUG](#)], Section 4.10.

Input Parameters

- balanc* CHARACTER*1. Must be 'N', 'P', 'S', or 'B'.
Indicates how the input matrix should be diagonally scaled and/or permuted to improve the conditioning of its eigenvalues.
If *balanc* = 'N', do not diagonally scale or permute;
If *balanc* = 'P', perform permutations to make the matrix more nearly upper triangular. Do not diagonally scale;
If *balanc* = 'S', Diagonally scale the matrix, i.e. replace A by $DA D^{-1}$, where D is a diagonal matrix chosen to make the rows and columns of A more equal in norm. Do not permute;
If *balanc* = 'B', both diagonally scale and permute A .
Computed reciprocal condition numbers will be for the matrix after balancing and/or permuting. Permuting does not change condition numbers (in exact arithmetic), but balancing does.
- jobvl* CHARACTER*1. Must be 'N' or 'V'.
If *jobvl* = 'N', left eigenvectors of A are not computed;
If *jobvl* = 'V', left eigenvectors of A are computed.
If *sense* = 'E' or 'B', then *jobvl* must be 'V'.

<i>jobvr</i>	<p>CHARACTER*1. Must be 'N' or 'V'.</p> <p>If <i>jobvr</i> = 'N', right eigenvectors of <i>A</i> are not computed;</p> <p>If <i>jobvr</i> = 'V', right eigenvectors of <i>A</i> are computed.</p> <p>If <i>sense</i> = 'E' or 'B', then <i>jobvr</i> must be 'V'.</p>
<i>sense</i>	<p>CHARACTER*1. Must be 'N', 'E', 'V', or 'B'.</p> <p>Determines which reciprocal condition number are computed.</p> <p>If <i>sense</i> = 'N', none are computed;</p> <p>If <i>sense</i> = 'E', computed for eigenvalues only;</p> <p>If <i>sense</i> = 'V', computed for right eigenvectors only;</p> <p>If <i>sense</i> = 'B', computed for eigenvalues and right eigenvectors.</p> <p>If <i>sense</i> is 'E' or 'B', both left and right eigenvectors must also be computed (<i>jobvl</i> = 'V' and <i>jobvr</i> = 'V').</p>
<i>n</i>	INTEGER. The order of the matrix <i>A</i> ($n \geq 0$).
<i>a, work</i>	<p>REAL for <i>sggeevx</i></p> <p>DOUBLE PRECISION for <i>dgeevx</i></p> <p>COMPLEX for <i>cgeevx</i></p> <p>DOUBLE COMPLEX for <i>zgeevx</i>.</p> <p>Arrays:</p> <p><i>a</i>(<i>lda</i>,*) is an array containing the <i>n</i>-by-<i>n</i> matrix <i>A</i>. The second dimension of <i>a</i> must be at least $\max(1, n)$.</p> <p><i>work</i>(<i>lwork</i>) is a workspace array.</p>
<i>lda</i>	<p>INTEGER. The first dimension of the array <i>a</i>.</p> <p>Must be at least $\max(1, n)$.</p>
<i>ldvl, ldvr</i>	<p>INTEGER. The leading dimensions of the output arrays <i>vl</i> and <i>vr</i>, respectively. Constraints:</p> <p>$ldvl \geq 1$; $ldvr \geq 1$.</p> <p>If <i>jobvl</i> = 'V', $ldvl \geq \max(1, n)$;</p> <p>If <i>jobvr</i> = 'V', $ldvr \geq \max(1, n)$.</p>
<i>lwork</i>	<p>INTEGER. The dimension of the array <i>work</i>.</p> <p>For real flavors:</p> <p>If <i>sense</i> = 'N' or 'E', $lwork \geq \max(1, 2n)$, and</p>

if $jobvl = 'V'$ or $jobvr = 'V'$, $lwork \geq 3n$;

If $sense = 'V'$ or $'B'$, $lwork \geq n(n+6)$.

For good performance, $lwork$ must generally be larger.

For complex flavors:

If $sense = 'N'$ or $'E'$, $lwork \geq \max(1, 2n)$;

If $sense = 'V'$ or $'B'$, $lwork \geq n^2 + 2n$.

For good performance, $lwork$ must generally be larger.

$rwork$ REAL for `cgeevx`
 DOUBLE PRECISION for `zgeevx`
 Workspace array, DIMENSION at least $\max(1, 2n)$. Used in complex flavors only.

$iwork$ INTEGER.
 Workspace array, DIMENSION at least $\max(1, 2n-2)$. Used in real flavors only. Not referenced if $sense = 'N'$ or $'E'$.

Output Parameters

a On exit, this array is overwritten. If $jobvl = 'V'$ or $jobvr = 'V'$, it contains the real-Schur/Schur form of the balanced version of the input matrix A .

wr, wi REAL for `sgeevx`
 DOUBLE PRECISION for `dgeevx`
 Arrays, DIMENSION at least $\max(1, n)$ each.
 Contain the real and imaginary parts, respectively, of the computed eigenvalues. Complex conjugate pairs of eigenvalues appear consecutively with the eigenvalue having positive imaginary part first.

w COMPLEX for `cgeevx`
 DOUBLE COMPLEX for `zgeevx`.
 Array, DIMENSION at least $\max(1, n)$.
 Contains the computed eigenvalues.

vl, vr REAL for `sgeevx`
 DOUBLE PRECISION for `dgeevx`
 COMPLEX for `cgeevx`
 DOUBLE COMPLEX for `zgeevx`.

Arrays:

$\mathbf{vl}(\mathit{ldvl}, *)$; the second dimension of \mathbf{vl} must be at least $\max(1, n)$.

If $\mathit{jobvl} = 'V'$, the left eigenvectors $u(j)$ are stored one after another in the columns of \mathbf{vl} , in the same order as their eigenvalues. If $\mathit{jobvl} = 'N'$, \mathbf{vl} is not referenced.

For real flavors:

If the j -th eigenvalue is real, then $u(j) = \mathbf{vl}(:,j)$, the j -th column of \mathbf{vl} . If the j -th and $(j+1)$ -st eigenvalues form a complex conjugate pair, then $u(j) = \mathbf{vl}(:,j) + i*\mathbf{vl}(:,j+1)$ and $u(j+1) = \mathbf{vl}(:,j) - i*\mathbf{vl}(:,j+1)$, where $i = \sqrt{-1}$.

For complex flavors:

$u(j) = \mathbf{vl}(:,j)$, the j -th column of \mathbf{vl} .

$\mathbf{vr}(\mathit{ldvr}, *)$; the second dimension of \mathbf{vr} must be at least $\max(1, n)$.

If $\mathit{jobvr} = 'V'$, the right eigenvectors $v(j)$ are stored one after another in the columns of \mathbf{vr} , in the same order as their eigenvalues. If $\mathit{jobvr} = 'N'$, \mathbf{vr} is not referenced.

For real flavors:

If the j -th eigenvalue is real, then $v(j) = \mathbf{vr}(:,j)$, the j -th column of \mathbf{vr} . If the j -th and $(j+1)$ -st eigenvalues form a complex conjugate pair, then $v(j) = \mathbf{vr}(:,j) + i*\mathbf{vr}(:,j+1)$ and $v(j+1) = \mathbf{vr}(:,j) - i*\mathbf{vr}(:,j+1)$, where $i = \sqrt{-1}$.

For complex flavors:

$v(j) = \mathbf{vr}(:,j)$, the j -th column of \mathbf{vr} .

$\mathit{ilo}, \mathit{ihi}$

INTEGER.

ilo and ihi are integer values determined when A was balanced.

The balanced $A(i,j) = 0$ if $i > j$ and $j = 1, \dots, \mathit{ilo}-1$ or $i = \mathit{ihi}+1, \dots, n$.

If $\mathit{balanc} = 'N'$ or $'S'$, $\mathit{ilo} = 1$ and $\mathit{ihi} = n$.

scale

REAL for single-precision flavors

DOUBLE PRECISION for double-precision flavors.

Array, DIMENSION at least $\max(1, n)$.

Details of the permutations and scaling factors applied

when balancing A . If $P(j)$ is the index of the row and column interchanged with row and column j , and $D(j)$ is the scaling factor applied to row and column j , then

$$\begin{aligned} \text{scale}(j) &= P(j), & \text{for } j = 1, \dots, \text{ilo}-1 \\ &= D(j), & \text{for } j = \text{ilo}, \dots, \text{ihi} \\ &= P(j) & \text{for } j = \text{ihi}+1, \dots, n. \end{aligned}$$

The order in which the interchanges are made is n to $\text{ihi}+1$, then 1 to $\text{ilo}-1$.

abnrm

REAL for single-precision flavors

DOUBLE PRECISION for double-precision flavors.

The one-norm of the balanced matrix (the maximum of the sum of absolute values of elements of any column).

rconde, rcondv

REAL for single precision flavors

DOUBLE PRECISION for double precision flavors.

Arrays, DIMENSION at least $\max(1, n)$ each.

rconde(j) is the reciprocal condition number of the j -th eigenvalue.

rcondv(j) is the reciprocal condition number of the j -th right eigenvector.

work(1)

On exit, if *info* = 0, then *work(1)* returns the required minimal size of *lwork*.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = $-i$, the i th parameter had an illegal value.

If *info* = i , the QR algorithm failed to compute all the eigenvalues, and no eigenvectors or condition numbers have been computed; elements $1:\text{ilo}-1$ and $i+1:n$ of *wr* and *wi* (for real flavors) or *w* (for complex flavors) contain eigenvalues which have converged.

Application Notes

If you are in doubt how much workspace to supply for the array *work*, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

Singular Value Decomposition

This section describes LAPACK driver routines used for solving singular value problems. See also [computational routines](#) that can be called to solve these problems.

[Table 5-12](#) lists routines described in more detail below.

Table 5-12 Driver Routines for Singular Value Decomposition

Routine Name	Operation performed
?gesvd	Computes the singular value decomposition of a general rectangular matrix.
?gesdd	Computes the singular value decomposition of a general rectangular matrix using a divide and conquer method.
?ggsvd	Computes the generalized singular value decomposition of a pair of general rectangular matrices.

[?gesvd](#)

Computes the singular value decomposition of a general rectangular matrix.

```
call sgesvd ( jobu, jobvt, m, n, a, lda, s, u, ldu, vt, ldvt,
             work, lwork, info)
call dgesvd ( jobu, jobvt, m, n, a, lda, s, u, ldu, vt, ldvt,
             work, lwork, info)
call cgesvd ( jobu, jobvt, m, n, a, lda, s, u, ldu, vt, ldvt,
             work, lwork, rwork, info)
call zgesvd ( jobu, jobvt, m, n, a, lda, s, u, ldu, vt, ldvt,
             work, lwork, rwork, info)
```

Discussion

This routine computes the singular value decomposition (SVD) of a real/complex m -by- n matrix A , optionally computing the left and/or right singular vectors. The SVD is written

$$A = U \Sigma V^H$$

where Σ is an m -by- n matrix which is zero except for its $\min(m,n)$ diagonal elements, U is an m -by- m orthogonal/unitary matrix, and V is an n -by- n orthogonal/unitary matrix. The diagonal elements of Σ are the singular values of A ; they are real and non-negative, and are returned in descending order. The first $\min(m,n)$ columns of U and V are the left and right singular vectors of A .

Note that the routine returns V^H , not V .

Input Parameters

- jobu* CHARACTER*1. Must be 'A', 'S', 'O', or 'N'. Specifies options for computing all or part of the matrix U .
- If *jobu* = 'A', all m columns of U are returned in the array *u*;
 if *jobu* = 'S', the first $\min(m,n)$ columns of U (the left singular vectors) are returned in the array *u*;
 if *jobu* = 'O', the first $\min(m,n)$ columns of U (the left singular vectors) are overwritten on the array *a*;
 if *jobu* = 'N', no columns of U (no left singular vectors) are computed.
- jobvt* CHARACTER*1. Must be 'A', 'S', 'O', or 'N'. Specifies options for computing all or part of the matrix V^H .
- If *jobvt* = 'A', all n rows of V^H are returned in the array *vt*;
 if *jobvt* = 'S', the first $\min(m,n)$ rows of V^H (the right singular vectors) are returned in the array *vt*;
 if *jobvt* = 'O', the first $\min(m,n)$ rows of V^H (the right singular vectors) are overwritten on the array *a*;
 if *jobvt* = 'N', no rows of V^H (no right singular vectors) are computed.
- jobvt* and *jobu* cannot both be 'O'.
- m* INTEGER. The number of rows of the matrix A ($m \geq 0$).
- n* INTEGER. The number of columns in A ($n \geq 0$).

<i>a, work</i>	<p>REAL for <i>sgesvd</i> DOUBLE PRECISION for <i>dgesvd</i> COMPLEX for <i>cgesvd</i> DOUBLE COMPLEX for <i>zgesvd</i>.</p> <p>Arrays: <i>a(lda,*)</i> is an array containing the <i>m</i>-by-<i>n</i> matrix <i>A</i>. The second dimension of <i>a</i> must be at least $\max(1, n)$. <i>work(lwork)</i> is a workspace array.</p>
<i>lda</i>	<p>INTEGER. The first dimension of the array <i>a</i>. Must be at least $\max(1, m)$.</p>
<i>ldu, ldvt</i>	<p>INTEGER. The leading dimensions of the output arrays <i>u</i> and <i>vt</i>, respectively. Constraints: <i>ldu</i> ≥ 1 ; <i>ldvt</i> ≥ 1. If <i>jobu</i> = 'S' or 'A', <i>ldu</i> $\geq m$; If <i>jobvt</i> = 'A', <i>ldvt</i> $\geq n$; If <i>jobvt</i> = 'S', <i>ldvt</i> $\geq \min(m, n)$.</p>
<i>lwork</i>	<p>INTEGER. The dimension of the array <i>work</i>; <i>lwork</i> ≥ 1. Constraints: <i>lwork</i> $\geq \max(3 * \min(m, n) + \max(m, n), 5 * \min(m, n))$ (for real flavors); <i>lwork</i> $\geq 2 * \min(m, n) + \max(m, n)$ (for complex flavors). For good performance, <i>lwork</i> must generally be larger.</p>
<i>rwork</i>	<p>REAL for <i>cgesvd</i> DOUBLE PRECISION for <i>zgesvd</i> Workspace array, DIMENSION at least $\max(1, 5 * \min(m, n))$. Used in complex flavors only.</p>

Output Parameters

<i>a</i>	<p>On exit, If <i>jobu</i> = 'O', <i>a</i> is overwritten with the first $\min(m, n)$ columns of <i>U</i> (the left singular vectors, stored columnwise); If <i>jobvt</i> = 'O', <i>a</i> is overwritten with the first $\min(m, n)$</p>
----------	--

rows of V^H (the right singular vectors, stored rowwise);
 If $jobu \neq 'O'$ and $jobvt \neq 'O'$, the contents of a are destroyed.

s **REAL** for single precision flavors
 DOUBLE PRECISION for double precision flavors.
 Array, **DIMENSION** at least $\max(1, \min(m,n))$.
 Contains the singular values of A sorted so that
 $s(i) \geq s(i+1)$.

u, vt **REAL** for `sgesvd`
 DOUBLE PRECISION for `dgesvd`
 COMPLEX for `cgesvd`
 DOUBLE COMPLEX for `zgesvd`.
 Arrays:
 $u(ldu, *)$; the second dimension of u must be at least
 $\max(1, m)$ if $jobu = 'A'$, and at least $\max(1, \min(m,n))$ if
 $jobu = 'S'$.
 If $jobu = 'A'$, u contains the m -by- m orthogonal/unitary
 matrix U .
 If $jobu = 'S'$, u contains the first $\min(m,n)$ columns of
 U (the left singular vectors, stored columnwise).
 If $jobu = 'N'$ or $'O'$, u is not referenced.
 $vt(ldvt, *)$; the second dimension of vt must be at
 least $\max(1, n)$.
 If $jobvt = 'A'$, vt contains the n -by- n
 orthogonal/unitary matrix V^H .
 If $jobvt = 'S'$, vt contains the first $\min(m,n)$ rows of
 V^H (the right singular vectors, stored rowwise).
 If $jobvt = 'N'$ or $'O'$, vt is not referenced.

$work$ On exit, if $info = 0$, then $work(1)$ returns the required
 minimal size of $lwork$.
 For real flavors:
 If $info > 0$, $work(2:\min(m,n))$ contains the
 unconverged superdiagonal elements of an upper
 bidiagonal matrix B whose diagonal is in s (not

necessarily sorted). B satisfies $A = u * B * vt$, so it has the same singular values as A , and singular vectors related by u and vt .

rwork

On exit (for complex flavors), if $info > 0$, $rwork(1:\min(m,n)-1)$ contains the unconverged superdiagonal elements of an upper bidiagonal matrix B whose diagonal is in s (not necessarily sorted). B satisfies $A = u * B * vt$, so it has the same singular values as A , and singular vectors related by u and vt .

info

INTEGER.

If $info = 0$, the execution is successful.

If $info = -i$, the i th parameter had an illegal value.

If $info = i$, then if `?bdsqr` did not converge, i specifies how many superdiagonals of the intermediate bidiagonal form B did not converge to zero.

Application Notes

If you are in doubt how much workspace to supply for the array *rwork*, use a generous value of *lwork* for the first run. On exit, examine $rwork(1)$ and use this value for subsequent runs.

?gesdd

Computes the singular value decomposition of a general rectangular matrix using a divide and conquer method.

```
call sgesdd ( jobz, m, n, a, lda, s, u, ldu, vt, ldvt,  
             work, lwork, iwork, info)  
call dgesdd ( jobz, m, n, a, lda, s, u, ldu, vt, ldvt,  
             work, lwork, iwork, info)  
call cgesdd ( jobz, m, n, a, lda, s, u, ldu, vt, ldvt,  
             work, lwork, rwork, iwork, info)  
call zgesdd ( jobz, m, n, a, lda, s, u, ldu, vt, ldvt,  
             work, lwork, rwork, iwork, info)
```

Discussion

This routine computes the singular value decomposition (SVD) of a real/complex m -by- n matrix A , optionally computing the left and/or right singular vectors. If singular vectors are desired, it uses a divide and conquer algorithm.

The SVD is written

$$A = U \Sigma V^H$$

where Σ is an m -by- n matrix which is zero except for its $\min(m,n)$ diagonal elements, U is an m -by- m orthogonal/unitary matrix, and V is an n -by- n orthogonal/unitary matrix. The diagonal elements of Σ are the singular values of A ; they are real and non-negative, and are returned in descending order. The first $\min(m,n)$ columns of U and V are the left and right singular vectors of A .

Note that the routine returns V^H , not V .

Input Parameters

jobz CHARACTER*1. Must be 'A', 'S', 'O', or 'N'.
Specifies options for computing all or part of the matrix U .

If $jobz = 'A'$, all m columns of U and all n rows of V^T are returned in the arrays u and vt ;

if $jobz = 'S'$, the first $\min(m,n)$ columns of U and the first $\min(m,n)$ rows of V^T are returned in the arrays u and vt ;

if $jobz = 'O'$, then

if $m \geq n$, the first n columns of U are overwritten on the array a and all rows of V^T are returned in the array vt ;

if $m < n$, all columns of U are returned in the array u and the first m rows of V^T are overwritten in the array vt ;

if $jobz = 'N'$, no columns of U or rows of V^T are computed.

m INTEGER. The number of rows of the matrix A ($m \geq 0$).

n INTEGER. The number of columns in A ($n \geq 0$).

$a, work$ REAL for `sgesdd`
DOUBLE PRECISION for `dgesdd`
COMPLEX for `cgesdd`
DOUBLE COMPLEX for `zgesdd`.

Arrays:

$a(lda, *)$ is an array containing the m -by- n matrix A . The second dimension of a must be at least $\max(1, n)$.

$work(lwork)$ is a workspace array.

lda INTEGER. The first dimension of the array a . Must be at least $\max(1, m)$.

$ldu, ldvt$ INTEGER. The leading dimensions of the output arrays u and vt , respectively. Constraints:
 $ldu \geq 1$; $ldvt \geq 1$.

If $jobz = 'S'$ or $'A'$, or $jobz = 'O'$ and $m < n$, then $ldu \geq m$;

If $jobz = 'A'$ or $jobz = 'O'$ and $m \geq n$, then $ldvt \geq n$;

If $jobz = 'S'$, $ldvt \geq \min(m, n)$.

lwork **INTEGER**. The dimension of the array *work*; *lwork* ≥ 1 . See *Application Notes* for the suggested value of *lwork*.

rwork **REAL** for *cgesdd*
DOUBLE PRECISION for *zgesdd*
 Workspace array, **DIMENSION** at least $\max(1, 5 * \min(m, n))$ if *jobz* = 'N'. Otherwise, the dimension of *rwork* must be at least $5 * (\min(m, n))^2 + 7 * \min(m, n)$. This array is used in complex flavors only.

iwork **INTEGER**. Workspace array, **DIMENSION** at least $\max(1, 8 * \min(m, n))$.

Output Parameters

a On exit:
 If *jobz* = 'O', then if $m \geq n$, *a* is overwritten with the first *n* columns of *U* (the left singular vectors, stored columnwise). If $m < n$, *a* is overwritten with the first *m* rows of V^T (the right singular vectors, stored rowwise); If *jobz* \neq 'O', the contents of *a* are destroyed.

s **REAL** for single precision flavors
DOUBLE PRECISION for double precision flavors.
 Array, **DIMENSION** at least $\max(1, \min(m, n))$.
 Contains the singular values of *A* sorted so that $s(i) \geq s(i+1)$.

u, vt **REAL** for *sgesdd*
DOUBLE PRECISION for *dgesdd*
COMPLEX for *cgesdd*
DOUBLE COMPLEX for *zgesdd*.
 Arrays:
u(*ldu*, *); the second dimension of *u* must be at least $\max(1, m)$ if *jobz* = 'A' or *jobz* = 'O' and $m < n$.
 If *jobz* = 'S', the second dimension of *u* must be at least $\max(1, \min(m, n))$.
 If *jobz* = 'A' or *jobz* = 'O' and $m < n$, *u* contains the *m*-by-*m* orthogonal/unitary matrix *U*.
 If *jobz* = 'S', *u* contains the first $\min(m, n)$ columns of

U (the left singular vectors, stored columnwise).

If $jobz = 'O'$ and $m \geq n$, or $jobz = 'N'$, u is not referenced.

$vt(ldvt, *)$; the second dimension of vt must be at least $\max(1, n)$.

If $jobz = 'A'$ or $jobz = 'O'$ and $m \geq n$, vt contains the n -by- n orthogonal/unitary matrix V^T .

If $jobz = 'S'$, vt contains the first $\min(m, n)$ rows of V^T (the right singular vectors, stored rowwise).

If $jobz = 'O'$ and $m < n$, or $jobz = 'N'$, vt is not referenced.

$work(1)$ On exit, if $info = 0$, then $work(1)$ returns the required minimal size of $lwork$.

$info$ INTEGER.

If $info = 0$, the execution is successful.

If $info = -i$, the i th parameter had an illegal value.

If $info = i$, then $?bdsdc$ did not converge, updating process failed.

Application Notes

For real flavors:

If $jobz = 'N'$, $lwork \geq 3 * \min(m, n) + \max(\max(m, n), 6 * \min(m, n))$;

If $jobz = 'O'$, $lwork \geq 3 * (\min(m, n))^2 + \max(\max(m, n), 5 * (\min(m, n))^2 + 4 * \min(m, n))$;

If $jobz = 'S'$ or $'A'$, $lwork \geq 3 * (\min(m, n))^2 + \max(\max(m, n), 4 * (\min(m, n))^2 + 4 * \min(m, n))$.

For complex flavors:

If $jobz = 'N'$, $lwork \geq 2 * \min(m, n) + \max(m, n)$;

If $jobz = 'O'$, $lwork \geq 2 * (\min(m, n))^2 + \max(m, n) + 2 * \min(m, n)$;

If $jobz = 'S'$ or $'A'$, $lwork \geq (\min(m, n))^2 + \max(m, n) + 2 * \min(m, n)$;

For good performance, $lwork$ should generally be larger.

If you are in doubt how much workspace to supply for the array $work$, use a generous value of $lwork$ for the first run. On exit, examine $work(1)$ and use this value for subsequent runs.

?ggsvd

Computes the generalized singular value decomposition of a pair of general rectangular matrices.

```
call sggsvd ( jobu, jobv, jobq, m, n, p, k, l, a, lda, b, ldb, alpha,
              beta, u, ldu, v, ldv, q, ldq, work, iwork, info)
call dggsvd ( jobu, jobv, jobq, m, n, p, k, l, a, lda, b, ldb, alpha,
              beta, u, ldu, v, ldv, q, ldq, work, iwork, info)
call cggsvd ( jobu, jobv, jobq, m, n, p, k, l, a, lda, b, ldb, alpha,
              beta, u, ldu, v, ldv, q, ldq, work, rwork, iwork, info)
call zggsvd ( jobu, jobv, jobq, m, n, p, k, l, a, lda, b, ldb, alpha,
              beta, u, ldu, v, ldv, q, ldq, work, rwork, iwork, info)
```

Discussion

This routine computes the generalized singular value decomposition (GSVD) of an m -by- n real/complex matrix A and p -by- n real/complex matrix B :

$$U^H A Q = D_1^* (0 \ R), \quad V^H B Q = D_2^* (0 \ R),$$

where U , V and Q are orthogonal/unitary matrices.

Let $k+1$ = the effective numerical rank of the matrix $(A^H, B^H)^H$, then R is a $(k+1)$ -by- $(k+1)$ nonsingular upper triangular matrix, D_1 and D_2 are m -by- $(k+1)$ and p -by- $(k+1)$ "diagonal" matrices and of the following structures, respectively:

If $m-k-1 \geq 0$,

$$D_1 = \begin{matrix} & k & l \\ & \begin{pmatrix} I & 0 \\ 0 & C \end{pmatrix} \\ m-k-l & \begin{pmatrix} 0 & 0 \end{pmatrix} \end{matrix}$$

$$D_2 = \begin{matrix} & k & l \\ & \begin{pmatrix} 0 & S \\ 0 & 0 \end{pmatrix} \\ p-l & \begin{pmatrix} 0 & 0 \end{pmatrix} \end{matrix}$$

$$\begin{pmatrix} 0 & R \end{pmatrix} = \begin{matrix} & n-k-l & k & l \\ k & \begin{pmatrix} 0 & R_{11} & R_{12} \\ 0 & 0 & R_{22} \end{pmatrix} \\ l & \end{matrix}$$

where

$$C = \text{diag}(\alpha(k+1), \dots, \alpha(k+l))$$

$$S = \text{diag}(\beta(k+1), \dots, \beta(k+l))$$

$$C^2 + S^2 = I$$

R is stored in $a(1:k+l, n-k-l+1:n)$ on exit.

If $m-k-l < 0$,

$$D_1 = \begin{matrix} & k & m-k & k+l-m \\ m-k & \begin{pmatrix} I & 0 & 0 \\ 0 & C & 0 \end{pmatrix} \\ & \end{matrix}$$

$$D_2 = \begin{matrix} & k & m-k & k+l-m \\ m-k & \begin{pmatrix} 0 & S & 0 \\ 0 & 0 & I \\ 0 & 0 & 0 \end{pmatrix} \\ k+l-m & \\ p-l & \end{matrix}$$

$$\begin{pmatrix} 0 & R \end{pmatrix} = \begin{matrix} & n-k-l & k & m-k & k+l-m \\ k & \begin{pmatrix} 0 & R_{11} & R_{12} & R_{13} \\ 0 & 0 & R_{22} & R_{23} \\ 0 & 0 & 0 & R_{33} \end{pmatrix} \\ m-k & \\ k+l-m & \end{matrix}$$

where

$$C = \text{diag}(\alpha(k+1), \dots, \alpha(m)),$$

$$S = \text{diag}(\beta(k+1), \dots, \beta(m)),$$

$$C^2 + S^2 = I$$

On exit, $\begin{pmatrix} R_{11} & R_{12} & R_{13} \\ 0 & R_{22} & R_{23} \end{pmatrix}$ is stored in $a(1:m, n-k-1+1:n)$ and R_{33} is stored

in $b(m-k+1:l, n+m-k-1+1:n)$.

The routine computes C , S , R , and optionally the orthogonal/unitary transformation matrices U , V and Q .

In particular, if B is an n -by- n nonsingular matrix, then the GSVD of A and B implicitly gives the SVD of AB^{-1} :

$$AB^{-1} = U(D_1 D_2^{-1})V^H.$$

If $(A^H, B^H)^H$ has orthonormal columns, then the GSVD of A and B is also equal to the CS decomposition of A and B . Furthermore, the GSVD can be used to derive the solution of the eigenvalue problem:

$$A^H A x = \lambda B^H B x.$$

Input Parameters

<i>jobu</i>	CHARACTER*1. Must be 'U' or 'N'. If <i>jobu</i> = 'U', orthogonal/unitary matrix U is computed. If <i>jobu</i> = 'N', U is not computed.
<i>jobv</i>	CHARACTER*1. Must be 'V' or 'N'. If <i>jobv</i> = 'V', orthogonal/unitary matrix V is computed. If <i>jobv</i> = 'N', V is not computed.
<i>jobq</i>	CHARACTER*1. Must be 'Q' or 'N'. If <i>jobq</i> = 'Q', orthogonal/unitary matrix Q is computed. If <i>jobq</i> = 'N', Q is not computed.
<i>m</i>	INTEGER. The number of rows of the matrix A ($m \geq 0$).
<i>n</i>	INTEGER. The number of columns of the matrices A and B ($n \geq 0$).
<i>p</i>	INTEGER. The number of rows of the matrix B ($p \geq 0$).
<i>a</i> , <i>b</i> , <i>work</i>	REAL for <i>sggsvd</i> DOUBLE PRECISION for <i>dggsvd</i> COMPLEX for <i>cggsvd</i> DOUBLE COMPLEX for <i>zggsvd</i> .

Arrays:

$a(lda, *)$ contains the m -by- n matrix A .

The second dimension of a must be at least $\max(1, n)$.

$b(ldb, *)$ contains the p -by- n matrix B .

The second dimension of b must be at least $\max(1, n)$.

$work(*)$ is a workspace array. The dimension of $work$ must be at least $\max(3n, m, p) + n$.

lda	INTEGER. The first dimension of a ; at least $\max(1, m)$.
ldb	INTEGER. The first dimension of b ; at least $\max(1, p)$.
ldu	INTEGER. The first dimension of the array u . $ldu \geq \max(1, m)$ if $jobu = 'U'$; $ldu \geq 1$ otherwise.
ldv	INTEGER. The first dimension of the array v . $ldv \geq \max(1, p)$ if $jobv = 'V'$; $ldv \geq 1$ otherwise.
ldq	INTEGER. The first dimension of the array q . $ldq \geq \max(1, n)$ if $jobq = 'Q'$; $ldq \geq 1$ otherwise.
$iwork$	INTEGER. Workspace array, DIMENSION at least $\max(1, n)$.
$rwork$	REAL for <code>cggsvd</code> DOUBLE PRECISION for <code>zggsvd</code> . Workspace array, DIMENSION at least $\max(1, 2n)$. Used in complex flavors only.

Output Parameters

k, l	INTEGER. On exit, k and l specify the dimension of the subblocks. The sum $k+l$ is equal to the effective numerical rank of $(A^H, B^H)^H$.
a	On exit, a contains the triangular matrix R or part of R .
b	On exit, b contains part of the triangular matrix R if $m-k-l < 0$.
$alpha, beta$	REAL for single-precision flavors DOUBLE PRECISION for double-precision flavors. Arrays, DIMENSION at least $\max(1, n)$ each. Contain the generalized singular value pairs of A and B .

$\alpha(1:k) = 1,$
 $\beta(1:k) = 0,$
 and if $m-k-1 \geq 0,$
 $\alpha(k+1:k+1) = C,$
 $\beta(k+1:k+1) = S,$
 or if $m-k-1 < 0,$
 $\alpha(k+1:m) = C, \alpha(m+1:k+1) = 0$
 $\beta(k+1:m) = S, \beta(m+1:k+1) = 1$
 and
 $\alpha(k+1+1:n) = 0$
 $\beta(k+1+1:n) = 0.$

u, v, q

REAL for *s*ggsvd
 DOUBLE PRECISION for *d*ggsvd
 COMPLEX for *c*ggsvd
 DOUBLE COMPLEX for *z*ggsvd.

Arrays:

$u(ldu, *)$; the second dimension of u must be at least $\max(1, m)$.

If $jobu = 'U'$, u contains the m -by- m orthogonal/unitary matrix U .

If $jobu = 'N'$, u is not referenced.

$v(ldv, *)$; the second dimension of v must be at least $\max(1, p)$.

If $jobv = 'V'$, v contains the p -by- p orthogonal/unitary matrix V .

If $jobv = 'N'$, v is not referenced.

$q(ldq, *)$; the second dimension of q must be at least $\max(1, n)$.

If $jobq = 'Q'$, q contains the n -by- n orthogonal/unitary matrix Q .

If $jobq = 'N'$, q is not referenced.

$iwork$

On exit, $iwork$ stores the sorting information.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

If *info* = 1, the Jacobi-type procedure failed to converge. For further details, see subroutine [?tgsja](#).

Generalized Symmetric Definite Eigenproblems

This section describes LAPACK driver routines used for solving generalized symmetric definite eigenproblems. See also [computational routines](#) that can be called to solve these problems.

[Table 5-13](#) lists routines described in more detail below.

Table 5-13 Driver Routines for Solving Generalized Symmetric Definite Eigenproblems

Routine Name	Operation performed
?sygv/?hegv	Computes all eigenvalues and, optionally, eigenvectors of a real / complex generalized symmetric /Hermitian definite eigenproblem.
?sygvd/?hegvd	Computes all eigenvalues and, optionally, eigenvectors of a real / complex generalized symmetric /Hermitian definite eigenproblem. If eigenvectors are desired, it uses a divide and conquer method.
?sygvx/?hegvx	Computes selected eigenvalues and, optionally, eigenvectors of a real / complex generalized symmetric /Hermitian definite eigenproblem.
?spgv/?hpgv	Computes all eigenvalues and, optionally, eigenvectors of a real / complex generalized symmetric /Hermitian definite eigenproblem with matrices in packed storage.
?spgvd/?hpgvd	Computes all eigenvalues and, optionally, eigenvectors of a real / complex generalized symmetric /Hermitian definite eigenproblem with matrices in packed storage. If eigenvectors are desired, it uses a divide and conquer method.
?spgvx/?hpgvx	Computes selected eigenvalues and, optionally, eigenvectors of a real / complex generalized symmetric /Hermitian definite eigenproblem with matrices in packed storage.
?sbgv/?hbgv	Computes all eigenvalues and, optionally, eigenvectors of a real / complex generalized symmetric /Hermitian definite eigenproblem with banded matrices.
?sbgvd/?hbgvd	Computes all eigenvalues and, optionally, eigenvectors of a real / complex generalized symmetric /Hermitian definite eigenproblem with banded matrices. If eigenvectors are desired, it uses a divide and conquer method.
?sbgvx/?hbgvx	Computes selected eigenvalues and, optionally, eigenvectors of a real / complex generalized symmetric /Hermitian definite eigenproblem with banded matrices.

?sygv

Computes all eigenvalues and, optionally, eigenvectors of a real generalized symmetric definite eigenproblem.

```
call ssgv ( itype, jobz, uplo, n, a, lda, b, ldb, w, work,
           lwork, info )
call dsgv ( itype, jobz, uplo, n, a, lda, b, ldb, w, work,
           lwork, info )
```

Discussion

This routine computes all the eigenvalues, and optionally, the eigenvectors of a real generalized symmetric-definite eigenproblem, of the form

$$Ax = \lambda Bx, \quad ABx = \lambda x, \quad \text{or} \quad BAx = \lambda x.$$

Here A and B are assumed to be symmetric and B is also positive definite.

Input Parameters

<i>itype</i>	INTEGER. Must be 1 or 2 or 3. Specifies the problem type to be solved: if <i>itype</i> = 1, the problem type is $Ax = \lambda Bx$; if <i>itype</i> = 2, the problem type is $ABx = \lambda x$; if <i>itype</i> = 3, the problem type is $BAx = \lambda x$.
<i>jobz</i>	CHARACTER*1. Must be 'N' or 'V'. If <i>jobz</i> = 'N', then compute eigenvalues only. If <i>jobz</i> = 'V', then compute eigenvalues and eigenvectors.
<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. If <i>uplo</i> = 'U', arrays <i>a</i> and <i>b</i> store the upper triangles of A and B ; If <i>uplo</i> = 'L', arrays <i>a</i> and <i>b</i> store the lower triangles of A and B .
<i>n</i>	INTEGER. The order of the matrices A and B ($n \geq 0$).

a, *b*, *work* REAL for *ssygv*
DOUBLE PRECISION for *dsygv*.
Arrays:
a(lda,)* contains the upper or lower triangle of the symmetric matrix *A*, as specified by *uplo*.
The second dimension of *a* must be at least $\max(1, n)$.
b(ldb,)* contains the upper or lower triangle of the symmetric positive definite matrix *B*, as specified by *uplo*.
The second dimension of *b* must be at least $\max(1, n)$.
work(lwork) is a workspace array.

lda INTEGER. The first dimension of *a*; at least $\max(1, n)$.

ldb INTEGER. The first dimension of *b*; at least $\max(1, n)$.

lwork INTEGER. The dimension of the array *work*;
lwork $\geq \max(1, 3n-1)$.
See *Application Notes* for the suggested value of *lwork*.

Output Parameters

a On exit, if *jobz* = 'V', then if *info* = 0, *a* contains the matrix *Z* of eigenvectors. The eigenvectors are normalized as follows:
if *itype* = 1 or 2, $Z^T B Z = I$;
if *itype* = 3, $Z^T B^{-1} Z = I$;
If *jobz* = 'N', then on exit the upper triangle (if *uplo* = 'U') or the lower triangle (if *uplo* = 'L') of *A*, including the diagonal, is destroyed.

b On exit, if *info* $\leq n$, the part of *b* containing the matrix is overwritten by the triangular factor *U* or *L* from the Cholesky factorization $B = U^T U$ or $B = L L^T$.

w REAL for *ssygv*
DOUBLE PRECISION for *dsygv*.
Array, DIMENSION at least $\max(1, n)$.
If *info* = 0, contains the eigenvalues in ascending order.

work(1) On exit, if *info* = 0, then *work(1)* returns the required minimal size of *lwork*.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th argument had an illegal value.
If *info* > 0, *spotrf/dpotrf* and *ssyev/dsyev* returned an error code:

If *info* = $i \leq n$, *ssyev/dsyev* failed to converge, and *i* off-diagonal elements of an intermediate tridiagonal did not converge to zero;

If *info* = $n + i$, for $1 \leq i \leq n$, then the leading minor of order *i* of *B* is not positive-definite. The factorization of *B* could not be completed and no eigenvalues or eigenvectors were computed.

Application Notes

For optimum performance use $lwork \geq (nb+2)*n$, where *nb* is the blocksize for *ssytrd/dsytrd* returned by *ilaenv*.

If you are in doubt how much workspace to supply for the array *work*, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

?hegv

Computes all eigenvalues and, optionally, eigenvectors of a complex generalized Hermitian definite eigenproblem.

```
call chegv ( itype, jobz, uplo, n, a, lda, b, ldb, w, work,
            lwork, rwork, info )
call zhegv ( itype, jobz, uplo, n, a, lda, b, ldb, w, work,
            lwork, rwork, info )
```

Discussion

This routine computes all the eigenvalues, and optionally, the eigenvectors of a complex generalized Hermitian-definite eigenproblem, of the form

$$Ax = \lambda Bx, \quad ABx = \lambda x, \quad \text{or} \quad BAx = \lambda x.$$

Here A and B are assumed to be Hermitian and B is also positive definite.

Input Parameters

itype **INTEGER**. Must be 1 or 2 or 3.
Specifies the problem type to be solved:
if *itype* = 1, the problem type is $Ax = \lambda Bx$;
if *itype* = 2, the problem type is $ABx = \lambda x$;
if *itype* = 3, the problem type is $BAx = \lambda x$.

jobz **CHARACTER*1**. Must be 'N' or 'V'.
If *jobz* = 'N', then compute eigenvalues only.
If *jobz* = 'V', then compute eigenvalues and eigenvectors.

uplo **CHARACTER*1**. Must be 'U' or 'L'.
If *uplo* = 'U', arrays *a* and *b* store the upper triangles of A and B ;
If *uplo* = 'L', arrays *a* and *b* store the lower triangles of A and B .

n **INTEGER**. The order of the matrices A and B ($n \geq 0$).

a, *b*, *work* COMPLEX for *cheqv*
 DOUBLE COMPLEX for *zhegv*.
 Arrays:
a(lda,)* contains the upper or lower triangle of the Hermitian matrix *A*, as specified by *uplo*.
 The second dimension of *a* must be at least $\max(1, n)$.
b(l db,)* contains the upper or lower triangle of the Hermitian positive definite matrix *B*, as specified by *uplo*.
 The second dimension of *b* must be at least $\max(1, n)$.
work(lwork) is a workspace array.

lda INTEGER. The first dimension of *a*; at least $\max(1, n)$.

ldb INTEGER. The first dimension of *b*; at least $\max(1, n)$.

lwork INTEGER. The dimension of the array *work*;
 $lwork \geq \max(1, 2n-1)$.
 See *Application Notes* for the suggested value of *lwork*.

rwork REAL for *cheqv*
 DOUBLE PRECISION for *zhegv*.
 Workspace array, DIMENSION at least $\max(1, 3n-2)$.

Output Parameters

a On exit, if *jobz* = 'V', then if *info* = 0, *a* contains the matrix *Z* of eigenvectors. The eigenvectors are normalized as follows:
 if *itype* = 1 or 2, $Z^H B Z = I$;
 if *itype* = 3, $Z^H B^{-1} Z = I$;
 If *jobz* = 'N', then on exit the upper triangle (if *uplo* = 'U') or the lower triangle (if *uplo* = 'L') of *A*, including the diagonal, is destroyed.

b On exit, if *info* $\leq n$, the part of *b* containing the matrix is overwritten by the triangular factor *U* or *L* from the Cholesky factorization $B = U^H U$ or $B = L L^H$.

w REAL for *chegv*
DOUBLE PRECISION for *zhegv*.
Array, DIMENSION at least $\max(1, n)$.
If *info* = 0, contains the eigenvalues in ascending order.

work(1) On exit, if *info* = 0, then *work(1)* returns the required minimal size of *lwork*.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th argument had an illegal value.
If *info* > 0, *cpotrf/zpotrf* and *cheev/zheev* returned an error code:
If *info* = $i \leq n$, *cheev/zheev* failed to converge, and *i* off-diagonal elements of an intermediate tridiagonal did not converge to zero;
If *info* = $n + i$, for $1 \leq i \leq n$, then the leading minor of order *i* of *B* is not positive-definite. The factorization of *B* could not be completed and no eigenvalues or eigenvectors were computed.

Application Notes

For optimum performance use $lwork \geq (nb+1)*n$, where *nb* is the blocksize for *chetrd/zhetrd* returned by *ilaenv*.

If you are in doubt how much workspace to supply for the array *work*, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

?sygvd

Computes all eigenvalues and, optionally, eigenvectors of a real generalized symmetric definite eigenproblem. If eigenvectors are desired, it uses a divide and conquer method.

```
call ssygvd ( itype, jobz, uplo, n, a, lda, b, ldb, w, work,
             lwork, iwork, liwork, info )
call dsygvd ( itype, jobz, uplo, n, a, lda, b, ldb, w, work,
             lwork, iwork, liwork, info )
```

Discussion

This routine computes all the eigenvalues, and optionally, the eigenvectors of a real generalized symmetric-definite eigenproblem, of the form

$$Ax = \lambda Bx, \quad ABx = \lambda x, \quad \text{or} \quad B Ax = \lambda x.$$

Here A and B are assumed to be symmetric and B is also positive definite.

If eigenvectors are desired, it uses a divide and conquer algorithm.

Input Parameters

itype **INTEGER**. Must be 1 or 2 or 3.
 Specifies the problem type to be solved:
 if *itype* = 1, the problem type is $Ax = \lambda Bx$;
 if *itype* = 2, the problem type is $ABx = \lambda x$;
 if *itype* = 3, the problem type is $B Ax = \lambda x$.

jobz **CHARACTER*1**. Must be 'N' or 'V'.
 If *jobz* = 'N', then compute eigenvalues only.
 If *jobz* = 'V', then compute eigenvalues and eigenvectors.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', arrays *a* and *b* store the upper triangles of *A* and *B*;
 If *uplo* = 'L', arrays *a* and *b* store the lower triangles of *A* and *B*.

n INTEGER. The order of the matrices *A* and *B* ($n \geq 0$).

a, *b*, *work* REAL for *ssygv*d
 DOUBLE PRECISION for *dsygv*d.
 Arrays:
a(*lda*,*) contains the upper or lower triangle of the symmetric matrix *A*, as specified by *uplo*.
 The second dimension of *a* must be at least $\max(1, n)$.
b(*ldb*,*) contains the upper or lower triangle of the symmetric positive definite matrix *B*, as specified by *uplo*.
 The second dimension of *b* must be at least $\max(1, n)$.
work(*lwork*) is a workspace array.

lda INTEGER. The first dimension of *a*; at least $\max(1, n)$.

ldb INTEGER. The first dimension of *b*; at least $\max(1, n)$.

lwork INTEGER. The dimension of the array *work*.
 Constraints:
 If $n \leq 1$, *lwork* ≥ 1 ;
 If *jobz* = 'N' and $n > 1$, *lwork* $\geq 2n+1$;
 If *jobz* = 'V' and $n > 1$, *lwork* $\geq 2n^2+6n+1$.

iwork INTEGER.
 Workspace array, DIMENSION (*liwork*).

liwork INTEGER. The dimension of the array *iwork*.
 Constraints:
 If $n \leq 1$, *liwork* ≥ 1 ;
 If *jobz* = 'N' and $n > 1$, *liwork* ≥ 1 ;
 If *jobz* = 'V' and $n > 1$, *liwork* $\geq 5n+3$.

Output Parameters

- a* On exit, if *jobz* = 'V', then if *info* = 0, *a* contains the matrix *Z* of eigenvectors. The eigenvectors are normalized as follows:
 if *itype* = 1 or 2, $Z^T B Z = I$;
 if *itype* = 3, $Z^T B^{-1} Z = I$;
- If *jobz* = 'N', then on exit the upper triangle (if *uplo* = 'U') or the lower triangle (if *uplo* = 'L') of *A*, including the diagonal, is destroyed.
- b* On exit, if *info* ≤ *n*, the part of *b* containing the matrix is overwritten by the triangular factor *U* or *L* from the Cholesky factorization $B = U^T U$ or $B = L L^T$.
- w* REAL for *ssygv*d
 DOUBLE PRECISION for *dsygv*d.
 Array, DIMENSION at least max(1, *n*).
 If *info* = 0, contains the eigenvalues in ascending order.
- work(1)* On exit, if *info* = 0, then *work(1)* returns the required minimal size of *lwork*.
- iwork(1)* On exit, if *info* = 0, then *iwork(1)* returns the required minimal size of *liwork*.
- info* INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th argument had an illegal value.
 If *info* > 0, *spotrf/dpotrf* and *ssyev/dsyev* returned an error code:
 If *info* = *i* ≤ *n*, *ssyev/dsyev* failed to converge, and *i* off-diagonal elements of an intermediate tridiagonal did not converge to zero;
 If *info* = *n* + *i*, for 1 ≤ *i* ≤ *n*, then the leading minor of order *i* of *B* is not positive-definite. The factorization of *B* could not be completed and no eigenvalues or eigenvectors were computed.

?hegvd

Computes all eigenvalues and, optionally, eigenvectors of a complex generalized Hermitian definite eigenproblem. If eigenvectors are desired, it uses a divide and conquer method.

```
call chegvd ( itype, jobz, uplo, n, a, lda, b, ldb, w, work,  
             lwork, rwork, lrwork, iwork, liwork, info )  
call zhegvd ( itype, jobz, uplo, n, a, lda, b, ldb, w, work,  
             lwork, rwork, lrwork, iwork, liwork, info )
```

Discussion

This routine computes all the eigenvalues, and optionally, the eigenvectors of a complex generalized Hermitian-definite eigenproblem, of the form

$$Ax = \lambda Bx, \quad ABx = \lambda x, \quad \text{or} \quad BAx = \lambda x.$$

Here A and B are assumed to be Hermitian and B is also positive definite. If eigenvectors are desired, it uses a divide and conquer algorithm.

Input Parameters

itype **INTEGER**. Must be 1 or 2 or 3.
Specifies the problem type to be solved:
if *itype* = 1, the problem type is $Ax = \lambda Bx$;
if *itype* = 2, the problem type is $ABx = \lambda x$;
if *itype* = 3, the problem type is $BAx = \lambda x$.

jobz **CHARACTER*1**. Must be 'N' or 'V'.
If *jobz* = 'N', then compute eigenvalues only.
If *jobz* = 'V', then compute eigenvalues and eigenvectors.

<i>uplo</i>	<p>CHARACTER*1. Must be 'U' or 'L'.</p> <p>If <i>uplo</i> = 'U', arrays <i>a</i> and <i>b</i> store the upper triangles of <i>A</i> and <i>B</i>;</p> <p>If <i>uplo</i> = 'L', arrays <i>a</i> and <i>b</i> store the lower triangles of <i>A</i> and <i>B</i>.</p>
<i>n</i>	<p>INTEGER. The order of the matrices <i>A</i> and <i>B</i> ($n \geq 0$).</p>
<i>a, b, work</i>	<p>COMPLEX for <i>chegvd</i></p> <p>DOUBLE COMPLEX for <i>zhegvd</i>.</p> <p>Arrays:</p> <p><i>a(lda,*)</i> contains the upper or lower triangle of the Hermitian matrix <i>A</i>, as specified by <i>uplo</i>. The second dimension of <i>a</i> must be at least $\max(1, n)$.</p> <p><i>b(ldb,*)</i> contains the upper or lower triangle of the Hermitian positive definite matrix <i>B</i>, as specified by <i>uplo</i>. The second dimension of <i>b</i> must be at least $\max(1, n)$.</p> <p><i>work(lwork)</i> is a workspace array.</p>
<i>lda</i>	<p>INTEGER. The first dimension of <i>a</i>; at least $\max(1, n)$.</p>
<i>ldb</i>	<p>INTEGER. The first dimension of <i>b</i>; at least $\max(1, n)$.</p>
<i>lwork</i>	<p>INTEGER. The dimension of the array <i>work</i>.</p> <p>Constraints:</p> <p>If $n \leq 1$, $lwork \geq 1$;</p> <p>If <i>jobz</i> = 'N' and $n > 1$, $lwork \geq n + 1$;</p> <p>If <i>jobz</i> = 'V' and $n > 1$, $lwork \geq n^2 + 2n$.</p>
<i>rwork</i>	<p>REAL for <i>chegvd</i></p> <p>DOUBLE PRECISION for <i>zhegvd</i>.</p> <p>Workspace array, DIMENSION (<i>lrwork</i>).</p>
<i>lrwork</i>	<p>INTEGER. The dimension of the array <i>rwork</i>.</p> <p>Constraints:</p> <p>If $n \leq 1$, $lrwork \geq 1$;</p> <p>If <i>jobz</i> = 'N' and $n > 1$, $lrwork \geq n$;</p> <p>If <i>jobz</i> = 'V' and $n > 1$, $lrwork \geq 2n^2 + 5n + 1$.</p>
<i>iwork</i>	<p>INTEGER.</p> <p>Workspace array, DIMENSION (<i>liwork</i>).</p>

liwork **INTEGER.** The dimension of the array *iwork*.
 Constraints:
 If $n \leq 1$, $liwork \geq 1$;
 If $jobz = 'N'$ and $n > 1$, $liwork \geq 1$;
 If $jobz = 'V'$ and $n > 1$, $liwork \geq 5n + 3$.

Output Parameters

a On exit, if $jobz = 'V'$, then if $info = 0$, *a* contains the matrix *Z* of eigenvectors. The eigenvectors are normalized as follows:
 if $itype = 1$ or 2 , $Z^H B Z = I$;
 if $itype = 3$, $Z^H B^{-1} Z = I$;
 If $jobz = 'N'$, then on exit the upper triangle (if $uplo = 'U'$) or the lower triangle (if $uplo = 'L'$) of *A*, including the diagonal, is destroyed.

b On exit, if $info \leq n$, the part of *b* containing the matrix is overwritten by the triangular factor *U* or *L* from the Cholesky factorization $B = U^H U$ or $B = L L^H$.

w **REAL** for *chegvd*
DOUBLE PRECISION for *zhgevd*.
 Array, **DIMENSION** at least $\max(1, n)$.
 If $info = 0$, contains the eigenvalues in ascending order.

work(1) On exit, if $info = 0$, then *work(1)* returns the required minimal size of *lwork*.

rwork(1) On exit, if $info = 0$, then *rwork(1)* returns the required minimal size of *lrwork*.

iwork(1) On exit, if $info = 0$, then *iwork(1)* returns the required minimal size of *liwork*.

info **INTEGER.**
 If $info = 0$, the execution is successful.
 If $info = -i$, the *i*th argument had an illegal value.
 If $info > 0$, *cpotrf/zpotrf* and *cheev/zheev* returned an error code:

If $info = i \leq n$, `cheev/zheev` failed to converge, and i off-diagonal elements of an intermediate tridiagonal did not converge to zero;

If $info = n + i$, for $1 \leq i \leq n$, then the leading minor of order i of B is not positive-definite. The factorization of B could not be completed and no eigenvalues or eigenvectors were computed.

?sygvx

Computes selected eigenvalues and, optionally, eigenvectors of a real generalized symmetric definite eigenproblem.

```
call ssygvx(itype, jobz, range, uplo, n, a, lda, b, ldb, vl, vu, il,
            iu, abstol, m, w, z, ldz, work, lwork, iwork, ifail, info)
call dsygvx(itype, jobz, range, uplo, n, a, lda, b, ldb, vl, vu, il,
            iu, abstol, m, w, z, ldz, work, lwork, iwork, ifail, info)
```

Discussion

This routine computes selected eigenvalues, and optionally, the eigenvectors of a real generalized symmetric-definite eigenproblem, of the form

$$Ax = \lambda Bx, \quad ABx = \lambda x, \quad \text{or} \quad BAx = \lambda x.$$

Here A and B are assumed to be symmetric and B is also positive definite. Eigenvalues and eigenvectors can be selected by specifying either a range of values or a range of indices for the desired eigenvalues.

Input Parameters

itype **INTEGER**. Must be 1 or 2 or 3.
 Specifies the problem type to be solved:
 if *itype* = 1, the problem type is $Ax = \lambda Bx$;
 if *itype* = 2, the problem type is $ABx = \lambda x$;
 if *itype* = 3, the problem type is $BAx = \lambda x$.

jobz **CHARACTER*1**. Must be 'N' or 'V'.
 If *jobz* = 'N', then compute eigenvalues only.
 If *jobz* = 'V', then compute eigenvalues and eigenvectors.

range **CHARACTER*1**. Must be 'A' or 'V' or 'I'.

If *range* = 'A', the routine computes all eigenvalues.
 If *range* = 'V', the routine computes eigenvalues λ_i in the half-open interval: $vl < \lambda_i \leq vu$.
 If *range* = 'I', the routine computes eigenvalues with indices *il* to *iu*.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', arrays *a* and *b* store the upper triangles of *A* and *B*;
 If *uplo* = 'L', arrays *a* and *b* store the lower triangles of *A* and *B*.

n INTEGER. The order of the matrices *A* and *B* ($n \geq 0$).

a, *b*, *work* REAL for *ssygvx*
 DOUBLE PRECISION for *dsygvx*.
 Arrays:
a(*lda*,*) contains the upper or lower triangle of the symmetric matrix *A*, as specified by *uplo*.
 The second dimension of *a* must be at least $\max(1, n)$.
b(*ldb*,*) contains the upper or lower triangle of the symmetric positive definite matrix *B*, as specified by *uplo*.
 The second dimension of *b* must be at least $\max(1, n)$.
work(*lwork*) is a workspace array.

lda INTEGER. The first dimension of *a*; at least $\max(1, n)$.

ldb INTEGER. The first dimension of *b*; at least $\max(1, n)$.

vl, *vu* REAL for *ssygvx*
 DOUBLE PRECISION for *dsygvx*.
 If *range* = 'V', the lower and upper bounds of the interval to be searched for eigenvalues.
 Constraint: $vl < vu$.
 If *range* = 'A' or 'I', *vl* and *vu* are not referenced.

<i>il, iu</i>	<p>INTEGER.</p> <p>If <i>range</i> = 'I', the indices in ascending order of the smallest and largest eigenvalues to be returned.</p> <p>Constraint: $1 \leq il \leq iu \leq n$, if $n > 0$; $il=1$ and $iu=0$ if $n = 0$.</p> <p>If <i>range</i> = 'A' or 'V', <i>il</i> and <i>iu</i> are not referenced.</p>
<i>abstol</i>	<p>REAL for <i>ssygvx</i></p> <p>DOUBLE PRECISION for <i>dsygvx</i>.</p> <p>The absolute error tolerance for the eigenvalues.</p> <p>See <i>Application Notes</i> for more information.</p>
<i>ldz</i>	<p>INTEGER. The leading dimension of the output array <i>z</i>.</p> <p>Constraints:</p> <p>$ldz \geq 1$; if <i>jobz</i> = 'V', $ldz \geq \max(1, n)$.</p>
<i>lwork</i>	<p>INTEGER. The dimension of the array <i>work</i>;</p> <p>$lwork \geq \max(1, 8n)$.</p> <p>See <i>Application Notes</i> for the suggested value of <i>lwork</i>.</p>
<i>iwork</i>	<p>INTEGER.</p> <p>Workspace array, DIMENSION at least $\max(1, 5n)$.</p>

Output Parameters

<i>a</i>	<p>On exit, the upper triangle (if <i>uplo</i> = 'U') or the lower triangle (if <i>uplo</i> = 'L') of <i>A</i>, including the diagonal, is overwritten.</p>
<i>b</i>	<p>On exit, if $info \leq n$, the part of <i>b</i> containing the matrix is overwritten by the triangular factor <i>U</i> or <i>L</i> from the Cholesky factorization $B = U^T U$ or $B = L L^T$.</p>
<i>m</i>	<p>INTEGER. The total number of eigenvalues found, $0 \leq m \leq n$. If <i>range</i> = 'A', $m = n$, and if <i>range</i> = 'I', $m = iu - il + 1$.</p>
<i>w, z</i>	<p>REAL for <i>ssygvx</i></p> <p>DOUBLE PRECISION for <i>dsygvx</i>.</p> <p>Arrays:</p>

$w(*)$, `DIMENSION` at least $\max(1, n)$.

The first m elements of w contain the selected eigenvalues in ascending order.

$z(ldz,*)$. The second dimension of z must be at least $\max(1, m)$.

If `jobz` = 'V', then if `info` = 0, the first m columns of z contain the orthonormal eigenvectors of the matrix A corresponding to the selected eigenvalues, with the i -th column of z holding the eigenvector associated with $w(i)$. The eigenvectors are normalized as follows:

if `itype` = 1 or 2, $Z^T B Z = I$;
 if `itype` = 3, $Z^T B^{-1} Z = I$;

If `jobz` = 'N', then z is not referenced.

If an eigenvector fails to converge, then that column of z contains the latest approximation to the eigenvector, and the index of the eigenvector is returned in `ifail`.

Note: you must ensure that at least $\max(1, m)$ columns are supplied in the array z ; if `range` = 'V', the exact value of m is not known in advance and an upper bound must be used.

`work(1)` On exit, if `info` = 0, then `work(1)` returns the required minimal size of `lwork`.

`ifail` INTEGER.

Array, `DIMENSION` at least $\max(1, n)$.

If `jobz` = 'V', then if `info` = 0, the first m elements of `ifail` are zero; if `info` > 0, the `ifail` contains the indices of the eigenvectors that failed to converge.

If `jobz` = 'N', then `ifail` is not referenced.

`info` INTEGER.

If `info` = 0, the execution is successful.

If `info` = $-i$, the i th argument had an illegal value.

If `info` > 0, `spotrf/dpotrf` and `ssyevx/dsyevx` returned an error code:

If $info = i \leq n$, `ssyevx/dsyevx` failed to converge, and i eigenvectors failed to converge. Their indices are stored in the array `ifail`;
If $info = n + i$, for $1 \leq i \leq n$, then the leading minor of order i of B is not positive-definite. The factorization of B could not be completed and no eigenvalues or eigenvectors were computed.

Application Notes

An approximate eigenvalue is accepted as converged when it is determined to lie in an interval $[a,b]$ of width less than or equal to

$abstol + \epsilon * \max(|a|,|b|)$, where ϵ is the machine precision. If $abstol$ is less than or equal to zero, then $\epsilon * \|T\|_1$ will be used in its place, where T is the tridiagonal matrix obtained by reducing A to tridiagonal form.

Eigenvalues will be computed most accurately when $abstol$ is set to twice the underflow threshold $2 * ?lamch('S')$, not zero. If this routine returns with $info > 0$, indicating that some eigenvectors did not converge, try setting $abstol$ to $2 * ?lamch('S')$.

For optimum performance use $lwork \geq (nb+3)*n$, where nb is the blocksize for `ssytrd/dsytrd` returned by `ilaenv`.

If you are in doubt how much workspace to supply for the array `work`, use a generous value of $lwork$ for the first run. On exit, examine `work(1)` and use this value for subsequent runs.

?hegvx

Computes selected eigenvalues and, optionally, eigenvectors of a complex generalized Hermitian definite eigenproblem.

```
call chegvx ( itype, jobz, range, uplo, n, a, lda, b, ldb, vl, vu,
             il, iu, abstol, m, w, z, ldz, work, lwork, rwork,
             iwork, ifail, info)
call zhegvx ( itype, jobz, range, uplo, n, a, lda, b, ldb, vl, vu,
             il, iu, abstol, m, w, z, ldz, work, lwork, rwork,
             iwork, ifail, info)
```

Discussion

This routine computes selected eigenvalues, and optionally, the eigenvectors of a complex generalized Hermitian-definite eigenproblem, of the form

$$Ax = \lambda Bx, \quad ABx = \lambda x, \quad \text{or} \quad BAx = \lambda x.$$

Here A and B are assumed to be Hermitian and B is also positive definite. Eigenvalues and eigenvectors can be selected by specifying either a range of values or a range of indices for the desired eigenvalues.

Input Parameters

<i>itype</i>	INTEGER. Must be 1 or 2 or 3. Specifies the problem type to be solved: if <i>itype</i> = 1, the problem type is $Ax = \lambda Bx$; if <i>itype</i> = 2, the problem type is $ABx = \lambda x$; if <i>itype</i> = 3, the problem type is $BAx = \lambda x$.
<i>jobz</i>	CHARACTER*1. Must be 'N' or 'V'. If <i>jobz</i> = 'N', then compute eigenvalues only. If <i>jobz</i> = 'V', then compute eigenvalues and eigenvectors.
<i>range</i>	CHARACTER*1. Must be 'A' or 'V' or 'I'.

If *range* = 'A', the routine computes all eigenvalues.
 If *range* = 'V', the routine computes eigenvalues λ_i in the half-open interval: $vl < \lambda_i \leq vu$.
 If *range* = 'I', the routine computes eigenvalues with indices *il* to *iu*.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', arrays *a* and *b* store the upper triangles of *A* and *B*;
 If *uplo* = 'L', arrays *a* and *b* store the lower triangles of *A* and *B*.

n INTEGER. The order of the matrices *A* and *B* ($n \geq 0$).

a, *b*, *work* COMPLEX for *chegvx*
 DOUBLE COMPLEX for *zhegvx*.
 Arrays:
a(*lda*,*) contains the upper or lower triangle of the Hermitian matrix *A*, as specified by *uplo*.
 The second dimension of *a* must be at least $\max(1, n)$.
b(*ldb*,*) contains the upper or lower triangle of the Hermitian positive definite matrix *B*, as specified by *uplo*.
 The second dimension of *b* must be at least $\max(1, n)$.
work(*lwork*) is a workspace array.

lda INTEGER. The first dimension of *a*; at least $\max(1, n)$.

ldb INTEGER. The first dimension of *b*; at least $\max(1, n)$.

vl, *vu* REAL for *chegvx*
 DOUBLE PRECISION for *zhegvx*.
 If *range* = 'V', the lower and upper bounds of the interval to be searched for eigenvalues.
 Constraint: $vl < vu$.
 If *range* = 'A' or 'I', *vl* and *vu* are not referenced.

<i>il, iu</i>	<p>INTEGER.</p> <p>If <i>range</i> = 'I', the indices in ascending order of the smallest and largest eigenvalues to be returned.</p> <p>Constraint: $1 \leq il \leq iu \leq n$, if $n > 0$; $il=1$ and $iu=0$ if $n = 0$.</p> <p>If <i>range</i> = 'A' or 'V', <i>il</i> and <i>iu</i> are not referenced.</p>
<i>abstol</i>	<p>REAL for <i>chegvx</i></p> <p>DOUBLE PRECISION for <i>zhegvx</i>.</p> <p>The absolute error tolerance for the eigenvalues.</p> <p>See <i>Application Notes</i> for more information.</p>
<i>ldz</i>	<p>INTEGER. The leading dimension of the output array <i>z</i>.</p> <p>Constraints:</p> <p>$ldz \geq 1$; if <i>jobz</i> = 'V', $ldz \geq \max(1, n)$.</p>
<i>lwork</i>	<p>INTEGER. The dimension of the array <i>work</i>;</p> <p>$lwork \geq \max(1, 2n-1)$.</p> <p>See <i>Application Notes</i> for the suggested value of <i>lwork</i>.</p>
<i>rwork</i>	<p>REAL for <i>chegvx</i></p> <p>DOUBLE PRECISION for <i>zhegvx</i>.</p> <p>Workspace array, DIMENSION at least $\max(1, 7n)$.</p>
<i>iwork</i>	<p>INTEGER.</p> <p>Workspace array, DIMENSION at least $\max(1, 5n)$.</p>

Output Parameters

<i>a</i>	<p>On exit, the upper triangle (if <i>uplo</i> = 'U') or the lower triangle (if <i>uplo</i> = 'L') of <i>A</i>, including the diagonal, is overwritten.</p>
<i>b</i>	<p>On exit, if <i>info</i> $\leq n$, the part of <i>b</i> containing the matrix is overwritten by the triangular factor <i>U</i> or <i>L</i> from the Cholesky factorization $B = U^H U$ or $B = L L^H$.</p>
<i>m</i>	<p>INTEGER. The total number of eigenvalues found, $0 \leq m \leq n$. If <i>range</i> = 'A', $m = n$, and if <i>range</i> = 'I', $m = iu - il + 1$.</p>

w REAL for *chegvx*
DOUBLE PRECISION for *zhegvx*.
Array, DIMENSION at least $\max(1, n)$.
The first m elements of *w* contain the selected eigenvalues in ascending order.

z COMPLEX for *chegvx*
DOUBLE COMPLEX for *zhegvx*.
Array *z*(*ldz*,*) . The second dimension of *z* must be at least $\max(1, m)$.
If *jobz* = 'V', then if *info* = 0, the first m columns of *z* contain the orthonormal eigenvectors of the matrix *A* corresponding to the selected eigenvalues, with the *i*-th column of *z* holding the eigenvector associated with *w*(*i*). The eigenvectors are normalized as follows:
if *itype* = 1 or 2, $Z^H B Z = I$;
if *itype* = 3, $Z^H B^{-1} Z = I$;
If *jobz* = 'N', then *z* is not referenced.
If an eigenvector fails to converge, then that column of *z* contains the latest approximation to the eigenvector, and the index of the eigenvector is returned in *ifail*.
Note: you must ensure that at least $\max(1, m)$ columns are supplied in the array *z*; if *range* = 'V', the exact value of m is not known in advance and an upper bound must be used.

work(1) On exit, if *info* = 0, then *work*(1) returns the required minimal size of *lwork*.

ifail INTEGER.
Array, DIMENSION at least $\max(1, n)$.
If *jobz* = 'V', then if *info* = 0, the first m elements of *ifail* are zero; if *info* > 0, the *ifail* contains the indices of the eigenvectors that failed to converge.
If *jobz* = 'N', then *ifail* is not referenced.

info

INTEGER.

If *info* = 0, the execution is successful.If *info* = -*i*, the *i*th argument had an illegal value.If *info* > 0, *cpotrf/zpotrf* and *cheevx/zheevx* returned an error code:

If *info* = *i* ≤ *n*, *cheevx/zheevx* failed to converge, and *i* eigenvectors failed to converge. Their indices are stored in the array *ifail*;

If *info* = *n* + *i*, for 1 ≤ *i* ≤ *n*, then the leading minor of order *i* of *B* is not positive-definite. The factorization of *B* could not be completed and no eigenvalues or eigenvectors were computed.

Application Notes

An approximate eigenvalue is accepted as converged when it is determined to lie in an interval [a,b] of width less than or equal to

abstol + ε * max(|a|,|b|), where ε is the machine precision. If *abstol* is less than or equal to zero, then ε * ||*T*||₁ will be used in its place, where *T* is the tridiagonal matrix obtained by reducing *A* to tridiagonal form.

Eigenvalues will be computed most accurately when *abstol* is set to twice the underflow threshold 2**?lamch*('S'), not zero. If this routine returns with *info* > 0, indicating that some eigenvectors did not converge, try setting *abstol* to 2**?lamch*('S').

For optimum performance use *lwork* ≥ (*nb*+1)**n*, where *nb* is the blocksize for *chetrd/zhetrd* returned by *ilaenv*.

If you are in doubt how much workspace to supply for the array *work*, use a generous value of *lwork* for the first run. On exit, examine *work*(1) and use this value for subsequent runs.

?spgv

Computes all eigenvalues and, optionally, eigenvectors of a real generalized symmetric definite eigenproblem with matrices in packed storage.

```
call sspgv ( itype, jobz, uplo, n, ap, bp, w, z, ldz, work, info )
call dspgv ( itype, jobz, uplo, n, ap, bp, w, z, ldz, work, info )
```

Discussion

This routine computes all the eigenvalues, and optionally, the eigenvectors of a real generalized symmetric-definite eigenproblem, of the form

$$Ax = \lambda Bx, \quad ABx = \lambda x, \quad \text{or} \quad BAx = \lambda x.$$

Here A and B are assumed to be symmetric, stored in packed format, and B is also positive definite.

Input Parameters

- itype* **INTEGER**. Must be 1 or 2 or 3.
Specifies the problem type to be solved:
if *itype* = 1, the problem type is $Ax = \lambda Bx$;
if *itype* = 2, the problem type is $ABx = \lambda x$;
if *itype* = 3, the problem type is $BAx = \lambda x$.
- jobz* **CHARACTER*1**. Must be 'N' or 'V'.
If *jobz* = 'N', then compute eigenvalues only.
If *jobz* = 'V', then compute eigenvalues and eigenvectors.
- uplo* **CHARACTER*1**. Must be 'U' or 'L'.
If *uplo* = 'U', arrays *ap* and *bp* store the upper triangles of A and B ;
If *uplo* = 'L', arrays *ap* and *bp* store the lower triangles of A and B .
- n* **INTEGER**. The order of the matrices A and B ($n \geq 0$).

ap, *bp*, *work* REAL for *sspgv*
DOUBLE PRECISION for *dspgv*.
Arrays:
ap(*) contains the packed upper or lower triangle of the symmetric matrix *A*, as specified by *uplo*. The dimension of *ap* must be at least $\max(1, n*(n+1)/2)$.
bp(*) contains the packed upper or lower triangle of the symmetric matrix *B*, as specified by *uplo*. The dimension of *bp* must be at least $\max(1, n*(n+1)/2)$.
work(*) is a workspace array, DIMENSION at least $\max(1, 3n)$.

ldz INTEGER. The leading dimension of the output array *z*;
 $ldz \geq 1$. If *jobz* = 'V', $ldz \geq \max(1, n)$.

Output Parameters

ap On exit, the contents of *ap* are overwritten.

bp On exit, contains the triangular factor *U* or *L* from the Cholesky factorization $B = U^T U$ or $B = L L^T$, in the same storage format as *B*.

w, *z* REAL for *sspgv*
DOUBLE PRECISION for *dspgv*.
Arrays:
w(*), DIMENSION at least $\max(1, n)$.
If *info* = 0, contains the eigenvalues in ascending order.
z(*ldz*,*) . The second dimension of *z* must be at least $\max(1, n)$.
If *jobz* = 'V', then if *info* = 0, *z* contains the matrix *Z* of eigenvectors. The eigenvectors are normalized as follows:
if *itype* = 1 or 2, $Z^T B Z = I$;
if *itype* = 3, $Z^T B^{-1} Z = I$;
If *jobz* = 'N', then *z* is not referenced.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th argument had an illegal value.

If *info* > 0, *sptrfs/dptrfs* and *sspev/dspev* returned an error code:

If *info* = $i \leq n$, *sspev/dspev* failed to converge, and *i* off-diagonal elements of an intermediate tridiagonal did not converge to zero;

If *info* = $n + i$, for $1 \leq i \leq n$, then the leading minor of order *i* of *B* is not positive-definite. The factorization of *B* could not be completed and no eigenvalues or eigenvectors were computed.

?hpgv

Computes all eigenvalues and, optionally, eigenvectors of a complex generalized Hermitian definite eigenproblem with matrices in packed storage.

```
call chpgv ( itype, jobz, uplo, n, ap, bp, w, z, ldz, work, rwork,
            info )
call zhpgv ( itype, jobz, uplo, n, ap, bp, w, z, ldz, work, rwork,
            info )
```

Discussion

This routine computes all the eigenvalues, and optionally, the eigenvectors of a complex generalized Hermitian-definite eigenproblem, of the form

$$Ax = \lambda Bx, \quad ABx = \lambda x, \quad \text{or} \quad BAx = \lambda x.$$

Here A and B are assumed to be Hermitian, stored in packed format, and B is also positive definite.

Input Parameters

itype **INTEGER**. Must be 1 or 2 or 3.
Specifies the problem type to be solved:
if *itype* = 1, the problem type is $Ax = \lambda Bx$;
if *itype* = 2, the problem type is $ABx = \lambda x$;
if *itype* = 3, the problem type is $BAx = \lambda x$.

jobz **CHARACTER*1**. Must be 'N' or 'V'.
If *jobz* = 'N', then compute eigenvalues only.
If *jobz* = 'V', then compute eigenvalues and eigenvectors.

uplo **CHARACTER*1**. Must be 'U' or 'L'.
If *uplo* = 'U', arrays *ap* and *bp* store the upper triangles of A and B ;
If *uplo* = 'L', arrays *ap* and *bp* store the lower triangles of A and B .

n **INTEGER**. The order of the matrices *A* and *B* ($n \geq 0$).

ap, *bp*, *work* **COMPLEX** for *chpgv*
DOUBLE COMPLEX for *zhpgv*.
 Arrays:
ap(*) contains the packed upper or lower triangle of the Hermitian matrix *A*, as specified by *uplo*. The dimension of *ap* must be at least $\max(1, n*(n+1)/2)$.
bp(*) contains the packed upper or lower triangle of the Hermitian matrix *B*, as specified by *uplo*. The dimension of *bp* must be at least $\max(1, n*(n+1)/2)$.
work(*) is a workspace array, **DIMENSION** at least $\max(1, 2n-1)$.

ldz **INTEGER**. The leading dimension of the output array *z*;
 $ldz \geq 1$. If *jobz* = 'V', $ldz \geq \max(1, n)$.

rwork **REAL** for *chpgv*
DOUBLE PRECISION for *zhpgv*.
 Workspace array, **DIMENSION** at least $\max(1, 3n-2)$.

Output Parameters

ap On exit, the contents of *ap* are overwritten.

bp On exit, contains the triangular factor *U* or *L* from the Cholesky factorization $B = U^H U$ or $B = L L^H$, in the same storage format as *B*.

w **REAL** for *chpgv*
DOUBLE PRECISION for *zhpgv*.
 Array, **DIMENSION** at least $\max(1, n)$.
 If *info* = 0, contains the eigenvalues in ascending order.

z **COMPLEX** for *chpgv*
DOUBLE COMPLEX for *zhpgv*.
 Array *z*(*ldz*, *). The second dimension of *z* must be at least $\max(1, n)$.
 If *jobz* = 'V', then if *info* = 0, *z* contains the matrix *Z* of eigenvectors. The eigenvectors are normalized as

follows:

if *itype* = 1 or 2, $Z^H B Z = I$;

if *itype* = 3, $Z^H B^{-1} Z = I$;

If *jobz* = 'N', then *z* is not referenced.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th argument had an illegal value.

If *info* > 0, *cpptrf/zpptrf* and *chpev/zhpev* returned an error code:

If *info* = *i* ≤ *n*, *chpev/zhpev* failed to converge, and *i* off-diagonal elements of an intermediate tridiagonal did not converge to zero;

If *info* = *n* + *i*, for 1 ≤ *i* ≤ *n*, then the leading minor of order *i* of *B* is not positive-definite. The factorization of *B* could not be completed and no eigenvalues or eigenvectors were computed.

?spgvd

Computes all eigenvalues and, optionally, eigenvectors of a real generalized symmetric definite eigenproblem with matrices in packed storage. If eigenvectors are desired, it uses a divide and conquer method.

```
call sspgvd ( itype, jobz, uplo, n, ap, bp, w, z, ldz, work, lwork,  
             iwork, liwork, info )  
call dspgvd ( itype, jobz, uplo, n, ap, bp, w, z, ldz, work, lwork,  
             iwork, liwork, info )
```

Discussion

This routine computes all the eigenvalues, and optionally, the eigenvectors of a real generalized symmetric-definite eigenproblem, of the form

$$Ax = \lambda Bx, \quad ABx = \lambda x, \quad \text{or} \quad B Ax = \lambda x.$$

Here A and B are assumed to be symmetric, stored in packed format, and B is also positive definite. If eigenvectors are desired, it uses a divide and conquer algorithm.

Input Parameters

itype **INTEGER**. Must be 1 or 2 or 3.
Specifies the problem type to be solved:
if *itype* = 1, the problem type is $Ax = \lambda Bx$;
if *itype* = 2, the problem type is $ABx = \lambda x$;
if *itype* = 3, the problem type is $B Ax = \lambda x$.

jobz **CHARACTER*1**. Must be 'N' or 'V'.
If *jobz* = 'N', then compute eigenvalues only.
If *jobz* = 'V', then compute eigenvalues and eigenvectors.

<i>uplo</i>	<p>CHARACTER*1. Must be 'U' or 'L'.</p> <p>If <i>uplo</i> = 'U', arrays <i>ap</i> and <i>bp</i> store the upper triangles of <i>A</i> and <i>B</i>;</p> <p>If <i>uplo</i> = 'L', arrays <i>ap</i> and <i>bp</i> store the lower triangles of <i>A</i> and <i>B</i>.</p>
<i>n</i>	<p>INTEGER. The order of the matrices <i>A</i> and <i>B</i> ($n \geq 0$).</p>
<i>ap, bp, work</i>	<p>REAL for <i>sspgvd</i></p> <p>DOUBLE PRECISION for <i>dspgvd</i>.</p> <p>Arrays:</p> <p><i>ap</i>(*) contains the packed upper or lower triangle of the symmetric matrix <i>A</i>, as specified by <i>uplo</i>. The dimension of <i>ap</i> must be at least $\max(1, n*(n+1)/2)$.</p> <p><i>bp</i>(*) contains the packed upper or lower triangle of the symmetric matrix <i>B</i>, as specified by <i>uplo</i>. The dimension of <i>bp</i> must be at least $\max(1, n*(n+1)/2)$.</p> <p><i>work</i>(<i>lwork</i>) is a workspace array.</p>
<i>ldz</i>	<p>INTEGER. The leading dimension of the output array <i>z</i>;</p> <p>$ldz \geq 1$. If <i>jobz</i> = 'V', $ldz \geq \max(1, n)$.</p>
<i>lwork</i>	<p>INTEGER. The dimension of the array <i>work</i>.</p> <p>Constraints:</p> <p>If $n \leq 1$, $lwork \geq 1$;</p> <p>If <i>jobz</i> = 'N' and $n > 1$, $lwork \geq 2n$;</p> <p>If <i>jobz</i> = 'V' and $n > 1$, $lwork \geq 2n^2 + 6n + 1$.</p>
<i>iwork</i>	<p>INTEGER.</p> <p>Workspace array, DIMENSION (<i>liwork</i>).</p>
<i>liwork</i>	<p>INTEGER. The dimension of the array <i>iwork</i>.</p> <p>Constraints:</p> <p>If $n \leq 1$, $liwork \geq 1$;</p> <p>If <i>jobz</i> = 'N' and $n > 1$, $liwork \geq 1$;</p> <p>If <i>jobz</i> = 'V' and $n > 1$, $liwork \geq 5n + 3$.</p>

Output Parameters

ap On exit, the contents of *ap* are overwritten.

<i>bp</i>	On exit, contains the triangular factor U or L from the Cholesky factorization $B = U^T U$ or $B = L L^T$, in the same storage format as B .
<i>w, z</i>	<p>REAL for <code>sppgv</code> DOUBLE PRECISION for <code>dsppgv</code>.</p> <p>Arrays: $w(*)$, DIMENSION at least $\max(1, n)$. If $info = 0$, contains the eigenvalues in ascending order. $z(ldz,*)$. The second dimension of z must be at least $\max(1, n)$. If $jobz = 'V'$, then if $info = 0$, z contains the matrix Z of eigenvectors. The eigenvectors are normalized as follows: if $itype = 1$ or 2, $Z^T B Z = I$; if $itype = 3$, $Z^T B^{-1} Z = I$;</p> <p>If $jobz = 'N'$, then z is not referenced.</p>
<i>work(1)</i>	On exit, if $info = 0$, then <i>work(1)</i> returns the required minimal size of <i>lwork</i> .
<i>iwork(1)</i>	On exit, if $info = 0$, then <i>iwork(1)</i> returns the required minimal size of <i>liwork</i> .
<i>info</i>	<p>INTEGER.</p> <p>If $info = 0$, the execution is successful. If $info = -i$, the ith argument had an illegal value. If $info > 0$, <code>spptrf/dpptrf</code> and <code>sspevd/dspevd</code> returned an error code:</p> <p>If $info = i \leq n$, <code>sspevd/dspevd</code> failed to converge, and i off-diagonal elements of an intermediate tridiagonal did not converge to zero; If $info = n + i$, for $1 \leq i \leq n$, then the leading minor of order i of B is not positive-definite. The factorization of B could not be completed and no eigenvalues or eigenvectors were computed.</p>

?hpgvd

Computes all eigenvalues and, optionally, eigenvectors of a complex generalized Hermitian definite eigenproblem with matrices in packed storage. If eigenvectors are desired, it uses a divide and conquer method.

```
call chpgvd ( itype, jobz, uplo, n, ap, bp, w, z, ldz, work, lwork,
             rwork, lrwork, iwork, liwork, info )
call zhpgvd ( itype, jobz, uplo, n, ap, bp, w, z, ldz, work, lwork,
             rwork, lrwork, iwork, liwork, info )
```

Discussion

This routine computes all the eigenvalues, and optionally, the eigenvectors of a complex generalized Hermitian-definite eigenproblem, of the form

$$Ax = \lambda Bx, \quad ABx = \lambda x, \quad \text{or} \quad B Ax = \lambda x.$$

Here A and B are assumed to be Hermitian, stored in packed format, and B is also positive definite. If eigenvectors are desired, it uses a divide and conquer algorithm.

Input Parameters

itype **INTEGER**. Must be 1 or 2 or 3.
 Specifies the problem type to be solved:
 if *itype* = 1, the problem type is $Ax = \lambda Bx$;
 if *itype* = 2, the problem type is $ABx = \lambda x$;
 if *itype* = 3, the problem type is $B Ax = \lambda x$.

jobz **CHARACTER*1**. Must be 'N' or 'V'.
 If *jobz* = 'N', then compute eigenvalues only.
 If *jobz* = 'V', then compute eigenvalues and eigenvectors.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', arrays *ap* and *bp* store the upper triangles of *A* and *B*;
 If *uplo* = 'L', arrays *ap* and *bp* store the lower triangles of *A* and *B*.

n INTEGER. The order of the matrices *A* and *B* ($n \geq 0$).

ap, *bp*, *work* COMPLEX for *chpgvd*
 DOUBLE COMPLEX for *zhpgvd*.
 Arrays:
ap(*) contains the packed upper or lower triangle of the Hermitian matrix *A*, as specified by *uplo*. The dimension of *ap* must be at least $\max(1, n*(n+1)/2)$.
bp(*) contains the packed upper or lower triangle of the Hermitian matrix *B*, as specified by *uplo*. The dimension of *bp* must be at least $\max(1, n*(n+1)/2)$.
work(*lwork*) is a workspace array.

ldz INTEGER. The leading dimension of the output array *z*;
 $ldz \geq 1$. If *jobz* = 'V', $ldz \geq \max(1, n)$.

lwork INTEGER. The dimension of the array *work*.
 Constraints:
 If $n \leq 1$, $lwork \geq 1$;
 If *jobz* = 'N' and $n > 1$, $lwork \geq n$;
 If *jobz* = 'V' and $n > 1$, $lwork \geq 2n$.

rwork REAL for *chpgvd*
 DOUBLE PRECISION for *zhpgvd*.
 Workspace array, DIMENSION (*lrwork*).

lrwork INTEGER. The dimension of the array *rwork*.
 Constraints:
 If $n \leq 1$, $lrwork \geq 1$;
 If *jobz* = 'N' and $n > 1$, $lrwork \geq n$;
 If *jobz* = 'V' and $n > 1$, $lrwork \geq 2n^2 + 5n + 1$.

iwork INTEGER.
 Workspace array, DIMENSION (*liwork*).

liwork **INTEGER**. The dimension of the array *iwork*.
 Constraints:
 If $n \leq 1$, $liwork \geq 1$;
 If $jobz = 'N'$ and $n > 1$, $liwork \geq 1$;
 If $jobz = 'V'$ and $n > 1$, $liwork \geq 5n + 3$.

Output Parameters

ap On exit, the contents of *ap* are overwritten.

bp On exit, contains the triangular factor U or L from the Cholesky factorization $B = U^H U$ or $B = L L^H$, in the same storage format as B .

w **REAL** for *chpgvd*
DOUBLE PRECISION for *zhpgvd*.
 Array, **DIMENSION** at least $\max(1, n)$.
 If $info = 0$, contains the eigenvalues in ascending order.

z **COMPLEX** for *chpgvd*
DOUBLE COMPLEX for *zhpgvd*.
 Array $z(ldz, *)$. The second dimension of z must be at least $\max(1, n)$.
 If $jobz = 'V'$, then if $info = 0$, z contains the matrix Z of eigenvectors. The eigenvectors are normalized as follows:
 if $itype = 1$ or 2 , $Z^H B Z = I$;
 if $itype = 3$, $Z^H B^{-1} Z = I$;
 If $jobz = 'N'$, then z is not referenced.

work(1) On exit, if $info = 0$, then *work(1)* returns the required minimal size of *lwork*.

rwork(1) On exit, if $info = 0$, then *rwork(1)* returns the required minimal size of *lrwork*.

iwork(1) On exit, if $info = 0$, then *iwork(1)* returns the required minimal size of *liwork*.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th argument had an illegal value.

If *info* > 0, *cpptrf/zpptrf* and *chpevd/zhpevd* returned an error code:

If *info* = *i* ≤ *n*, *chpevd/zhpevd* failed to converge, and *i* off-diagonal elements of an intermediate tridiagonal did not converge to zero;

If *info* = *n* + *i*, for 1 ≤ *i* ≤ *n*, then the leading minor of order *i* of *B* is not positive-definite. The factorization of *B* could not be completed and no eigenvalues or eigenvectors were computed.

?spgvx

Computes selected eigenvalues and, optionally, eigenvectors of a real generalized symmetric definite eigenproblem with matrices in packed storage.

```
call sspgvx ( itype, jobz, range, uplo, n, ap, bp, vl, vu, il, iu,
             abstol, m, w, z, ldz, work, iwork, ifail, info )
call dspgvx ( itype, jobz, range, uplo, n, ap, bp, vl, vu, il, iu,
             abstol, m, w, z, ldz, work, iwork, ifail, info )
```

Discussion

This routine computes selected eigenvalues, and optionally, the eigenvectors of a real generalized symmetric-definite eigenproblem, of the form

$$Ax = \lambda Bx, \quad ABx = \lambda x, \quad \text{or} \quad BAx = \lambda x.$$

Here A and B are assumed to be symmetric, stored in packed format, and B is also positive definite.

Eigenvalues and eigenvectors can be selected by specifying either a range of values or a range of indices for the desired eigenvalues.

Input Parameters

itype **INTEGER**. Must be 1 or 2 or 3.
Specifies the problem type to be solved:
if *itype* = 1, the problem type is $Ax = \lambda Bx$;
if *itype* = 2, the problem type is $ABx = \lambda x$;
if *itype* = 3, the problem type is $BAx = \lambda x$.

jobz **CHARACTER*1**. Must be 'N' or 'V'.
If *jobz* = 'N', then compute eigenvalues only.
If *jobz* = 'V', then compute eigenvalues and eigenvectors.

range **CHARACTER*1**. Must be 'A' or 'V' or 'I'.

If *range* = 'A', the routine computes all eigenvalues.
 If *range* = 'V', the routine computes eigenvalues λ_i in the half-open interval: $v_l < \lambda_i \leq v_u$.
 If *range* = 'I', the routine computes eigenvalues with indices *il* to *iu*.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', arrays *ap* and *bp* store the upper triangles of A and B;
 If *uplo* = 'L', arrays *ap* and *bp* store the lower triangles of A and B.

n INTEGER. The order of the matrices A and B ($n \geq 0$).

ap, *bp*, *work* REAL for *sspgvx*
 DOUBLE PRECISION for *dspgvx*.
 Arrays:
ap(*) contains the packed upper or lower triangle of the symmetric matrix A, as specified by *uplo*. The dimension of *ap* must be at least $\max(1, n*(n+1)/2)$.
bp(*) contains the packed upper or lower triangle of the symmetric matrix B, as specified by *uplo*. The dimension of *bp* must be at least $\max(1, n*(n+1)/2)$.
work(*) is a workspace array, DIMENSION at least $\max(1, 8n)$.

v_l, *v_u* REAL for *sspgvx*
 DOUBLE PRECISION for *dspgvx*.
 If *range* = 'V', the lower and upper bounds of the interval to be searched for eigenvalues.
 Constraint: $v_l < v_u$.
 If *range* = 'A' or 'I', *v_l* and *v_u* are not referenced.

il, *iu* INTEGER.
 If *range* = 'I', the indices in ascending order of the smallest and largest eigenvalues to be returned.
 Constraint: $1 \leq il \leq iu \leq n$, if $n > 0$; $il=1$ and $iu=0$ if $n = 0$.
 If *range* = 'A' or 'V', *il* and *iu* are not referenced.

<code>abstol</code>	<p>REAL for <code>sspgvx</code> DOUBLE PRECISION for <code>dspgvx</code>. The absolute error tolerance for the eigenvalues. See <i>Application Notes</i> for more information.</p>
<code>ldz</code>	<p>INTEGER. The leading dimension of the output array <code>z</code>. Constraints: $ldz \geq 1$; if <code>jobz = 'V'</code>, $ldz \geq \max(1, n)$.</p>
<code>iwork</code>	<p>INTEGER. Workspace array, DIMENSION at least $\max(1, 5n)$.</p>

Output Parameters

<code>ap</code>	On exit, the contents of <code>ap</code> are overwritten.
<code>bp</code>	On exit, contains the triangular factor U or L from the Cholesky factorization $B = U^T U$ or $B = L L^T$, in the same storage format as B .
<code>m</code>	<p>INTEGER. The total number of eigenvalues found, $0 \leq m \leq n$. If <code>range = 'A'</code>, $m = n$, and if <code>range = 'I'</code>, $m = iu - il + 1$.</p>
<code>w, z</code>	<p>REAL for <code>sspgvx</code> DOUBLE PRECISION for <code>dspgvx</code>. Arrays: <code>w(*)</code>, DIMENSION at least $\max(1, n)$. If <code>info = 0</code>, contains the eigenvalues in ascending order. <code>z(ldz,*)</code>. The second dimension of <code>z</code> must be at least $\max(1, n)$. If <code>jobz = 'V'</code>, then if <code>info = 0</code>, the first m columns of <code>z</code> contain the orthonormal eigenvectors of the matrix A corresponding to the selected eigenvalues, with the i-th column of <code>z</code> holding the eigenvector associated with <code>w(i)</code>. The eigenvectors are normalized as follows: if <code>itype = 1</code> or <code>2</code>, $Z^T B Z = I$; if <code>itype = 3</code>, $Z^T B^{-1} Z = I$; If <code>jobz = 'N'</code>, then <code>z</code> is not referenced. If an eigenvector fails to converge, then that column of <code>z</code> contains the latest approximation to the eigenvector, and</p>

the index of the eigenvector is returned in *ifail*.

Note: you must ensure that at least $\max(1,m)$ columns are supplied in the array *z*; if *range* = 'V', the exact value of *m* is not known in advance and an upper bound must be used.

ifail

INTEGER.

Array, DIMENSION at least $\max(1, n)$.

If *jobz* = 'V', then if *info* = 0, the first *m* elements of *ifail* are zero; if *info* > 0, the *ifail* contains the indices of the eigenvectors that failed to converge.

If *jobz* = 'N', then *ifail* is not referenced.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th argument had an illegal value.

If *info* > 0, *spptrf/dpptrf* and *sspevx/dspevx* returned an error code:

If *info* = $i \leq n$, *sspevx/dspevx* failed to converge, and *i* eigenvectors failed to converge. Their indices are stored in the array *ifail*;

If *info* = $n + i$, for $1 \leq i \leq n$, then the leading minor of order *i* of *B* is not positive-definite. The factorization of *B* could not be completed and no eigenvalues or eigenvectors were computed.

Application Notes

An approximate eigenvalue is accepted as converged when it is determined to lie in an interval [a,b] of width less than or equal to

$abstol + \epsilon * \max(|a|,|b|)$, where ϵ is the machine precision. If *abstol* is less than or equal to zero, then $\epsilon * \|T\|_1$ will be used in its place, where *T* is the tridiagonal matrix obtained by reducing *A* to tridiagonal form.

Eigenvalues will be computed most accurately when *abstol* is set to twice the underflow threshold $2 * ?lamch('S')$, not zero. If this routine returns with *info* > 0, indicating that some eigenvectors did not converge, try setting *abstol* to $2 * ?lamch('S')$.

?hpgvx

Computes selected eigenvalues and, optionally, eigenvectors of a generalized Hermitian definite eigenproblem with matrices in packed storage.

```
call chpgvx ( itype, jobz, range, uplo, n, ap, bp, vl, vu, il, iu,
             abstol, m, w, z, ldz, work, rwork, iwork, ifail, info )
call zhpgvx ( itype, jobz, range, uplo, n, ap, bp, vl, vu, il, iu,
             abstol, m, w, z, ldz, work, rwork, iwork, ifail, info )
```

Discussion

This routine computes selected eigenvalues, and optionally, the eigenvectors of a complex generalized Hermitian-definite eigenproblem, of the form

$$Ax = \lambda Bx, \quad ABx = \lambda x, \quad \text{or} \quad B Ax = \lambda x.$$

Here A and B are assumed to be Hermitian, stored in packed format, and B is also positive definite.

Eigenvalues and eigenvectors can be selected by specifying either a range of values or a range of indices for the desired eigenvalues.

Input Parameters

<i>itype</i>	INTEGER. Must be 1 or 2 or 3. Specifies the problem type to be solved: if <i>itype</i> = 1, the problem type is $Ax = \lambda Bx$; if <i>itype</i> = 2, the problem type is $ABx = \lambda x$; if <i>itype</i> = 3, the problem type is $B Ax = \lambda x$.
<i>jobz</i>	CHARACTER*1. Must be 'N' or 'V'. If <i>jobz</i> = 'N', then compute eigenvalues only. If <i>jobz</i> = 'V', then compute eigenvalues and eigenvectors.
<i>range</i>	CHARACTER*1. Must be 'A' or 'V' or 'I'. If <i>range</i> = 'A', the routine computes all eigenvalues. If <i>range</i> = 'V', the routine computes eigenvalues λ_i in the half-open interval: $vl < \lambda_i \leq vu$.

If *range* = 'I', the routine computes eigenvalues with indices *il* to *iu*.

uplo CHARACTER*1. Must be 'U' or 'L'.
 If *uplo* = 'U', arrays *ap* and *bp* store the upper triangles of *A* and *B*;
 If *uplo* = 'L', arrays *ap* and *bp* store the lower triangles of *A* and *B*.

n INTEGER. The order of the matrices *A* and *B* ($n \geq 0$).

ap, *bp*, *work* COMPLEX for *chpgvx*
 DOUBLE COMPLEX for *zhpgvx*.
 Arrays:
ap(*) contains the packed upper or lower triangle of the Hermitian matrix *A*, as specified by *uplo*. The dimension of *ap* must be at least $\max(1, n*(n+1)/2)$.
bp(*) contains the packed upper or lower triangle of the Hermitian matrix *B*, as specified by *uplo*. The dimension of *bp* must be at least $\max(1, n*(n+1)/2)$.
work(*) is a workspace array, DIMENSION at least $\max(1, 2n)$.

vl, *vu* REAL for *chpgvx*
 DOUBLE PRECISION for *zhpgvx*.
 If *range* = 'V', the lower and upper bounds of the interval to be searched for eigenvalues.
 Constraint: $vl < vu$.
 If *range* = 'A' or 'I', *vl* and *vu* are not referenced.

il, *iu* INTEGER.
 If *range* = 'I', the indices in ascending order of the smallest and largest eigenvalues to be returned.
 Constraint: $1 \leq il \leq iu \leq n$, if $n > 0$; $il=1$ and $iu=0$ if $n = 0$.
 If *range* = 'A' or 'V', *il* and *iu* are not referenced.

abstol REAL for *chpgvx*
 DOUBLE PRECISION for *zhpgvx*.
 The absolute error tolerance for the eigenvalues.

See *Application Notes* for more information.

<i>ldz</i>	INTEGER . The leading dimension of the output array <i>z</i> ; $ldz \geq 1$. If <i>jobz</i> = 'V', $ldz \geq \max(1, n)$.
<i>rwork</i>	REAL for <i>chpgvx</i> DOUBLE PRECISION for <i>zhpgvx</i> . Workspace array, DIMENSION at least $\max(1, 7n)$.
<i>iwork</i>	INTEGER . Workspace array, DIMENSION at least $\max(1, 5n)$.

Output Parameters

<i>ap</i>	On exit, the contents of <i>ap</i> are overwritten.
<i>bp</i>	On exit, contains the triangular factor <i>U</i> or <i>L</i> from the Cholesky factorization $B = U^H U$ or $B = L L^H$, in the same storage format as <i>B</i> .
<i>m</i>	INTEGER . The total number of eigenvalues found, $0 \leq m \leq n$. If <i>range</i> = 'A', $m = n$, and if <i>range</i> = 'I', $m = iu - il + 1$.
<i>w</i>	REAL for <i>chpgvx</i> DOUBLE PRECISION for <i>zhpgvx</i> . Array, DIMENSION at least $\max(1, n)$. If <i>info</i> = 0, contains the eigenvalues in ascending order.
<i>z</i>	COMPLEX for <i>chpgvx</i> DOUBLE COMPLEX for <i>zhpgvx</i> . Array <i>z</i> (<i>ldz</i> , *). The second dimension of <i>z</i> must be at least $\max(1, n)$. If <i>jobz</i> = 'V', then if <i>info</i> = 0, the first <i>m</i> columns of <i>z</i> contain the orthonormal eigenvectors of the matrix <i>A</i> corresponding to the selected eigenvalues, with the <i>i</i> -th column of <i>z</i> holding the eigenvector associated with <i>w</i> (<i>i</i>). The eigenvectors are normalized as follows: if <i>itype</i> = 1 or 2, $Z^H B Z = I$; if <i>itype</i> = 3, $Z^H B^{-1} Z = I$; If <i>jobz</i> = 'N', then <i>z</i> is not referenced.

If an eigenvector fails to converge, then that column of z contains the latest approximation to the eigenvector, and the index of the eigenvector is returned in *ifail*.

Note: you must ensure that at least $\max(1, m)$ columns are supplied in the array z ; if *range* = 'V', the exact value of m is not known in advance and an upper bound must be used.

ifail

INTEGER.

Array, DIMENSION at least $\max(1, n)$.

If *jobz* = 'V', then if *info* = 0, the first m elements of *ifail* are zero; if *info* > 0, the *ifail* contains the indices of the eigenvectors that failed to converge.

If *jobz* = 'N', then *ifail* is not referenced.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = - i , the i th argument had an illegal value.

If *info* > 0, *cpptrf/zpptrf* and *chpevx/zhpevx* returned an error code:

If *info* = $i \leq n$, *chpevx/zhpevx* failed to converge, and i eigenvectors failed to converge. Their indices are stored in the array *ifail*;

If *info* = $n + i$, for $1 \leq i \leq n$, then the leading minor of order i of B is not positive-definite. The factorization of B could not be completed and no eigenvalues or eigenvectors were computed.

Application Notes

An approximate eigenvalue is accepted as converged when it is determined to lie in an interval $[a, b]$ of width less than or equal to

$abstol + \epsilon * \max(|a|, |b|)$, where ϵ is the machine precision. If *abstol* is less than or equal to zero, then $\epsilon * \|T\|_1$ will be used in its place, where T is the tridiagonal matrix obtained by reducing A to tridiagonal form.

Eigenvalues will be computed most accurately when *abstol* is set to twice the underflow threshold $2 * ?lamch('S')$, not zero. If this routine returns with *info* > 0, indicating that some eigenvectors did not converge, try setting *abstol* to $2 * ?lamch('S')$.

?sbgv

Computes all eigenvalues and, optionally, eigenvectors of a real generalized symmetric definite eigenproblem with banded matrices.

```
call ssbgv ( jobz, uplo, n, ka, kb, ab, ldab, bb, ldbb, w, z, ldz,
            work, info )
call dsbgv ( jobz, uplo, n, ka, kb, ab, ldab, bb, ldbb, w, z, ldz,
            work, info )
```

Discussion

This routine computes all the eigenvalues, and optionally, the eigenvectors of a real generalized symmetric-definite banded eigenproblem, of the form $Ax = \lambda Bx$. Here A and B are assumed to be symmetric and banded, and B is also positive definite.

Input Parameters

<i>jobz</i>	CHARACTER*1. Must be 'N' or 'V'. If <i>jobz</i> = 'N', then compute eigenvalues only. If <i>jobz</i> = 'V', then compute eigenvalues and eigenvectors.
<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. If <i>uplo</i> = 'U', arrays <i>ab</i> and <i>bb</i> store the upper triangles of A and B ; If <i>uplo</i> = 'L', arrays <i>ab</i> and <i>bb</i> store the lower triangles of A and B .
<i>n</i>	INTEGER. The order of the matrices A and B ($n \geq 0$).
<i>ka</i>	INTEGER. The number of super- or sub-diagonals in A ($ka \geq 0$).
<i>kb</i>	INTEGER. The number of super- or sub-diagonals in B ($kb \geq 0$).

ab, bb, work REAL for *ssbgv*
DOUBLE PRECISION for *dsbgv*
Arrays:
*ab (ldab, *)* is an array containing either upper or lower triangular part of the symmetric matrix *A* (as specified by *uplo*) in band storage format. The second dimension of the array *ab* must be at least $\max(1, n)$.
*bb (ldbb, *)* is an array containing either upper or lower triangular part of the symmetric matrix *B* (as specified by *uplo*) in band storage format. The second dimension of the array *bb* must be at least $\max(1, n)$.
*work(*)* is a workspace array, DIMENSION at least $\max(1, 3n)$

ldab INTEGER. The first dimension of the array *ab*; must be at least *ka*+1.

ldbb INTEGER. The first dimension of the array *bb*; must be at least *kb*+1.

ldz INTEGER. The leading dimension of the output array *z*; *ldz* ≥ 1. If *jobz* = 'V', *ldz* ≥ $\max(1, n)$.

Output Parameters

ab On exit, the contents of *ab* are overwritten.

bb On exit, contains the factor *S* from the split Cholesky factorization $B = S^T S$, as returned by *spbstf/dpbstf*.

w, z REAL for *ssbgv*
DOUBLE PRECISION for *dsbgv*
Arrays:
*w(*)*, DIMENSION at least $\max(1, n)$.
If *info* = 0, contains the eigenvalues in ascending order.
*z(ldz, *)*. The second dimension of *z* must be at least $\max(1, n)$.
If *jobz* = 'V', then if *info* = 0, *z* contains the matrix *Z* of eigenvectors, with the *i*-th column of *z* holding the

eigenvector associated with $w(i)$. The eigenvectors are normalized so that $Z^T B Z = I$.

If $jobz = 'N'$, then z is not referenced.

info

INTEGER.

If $info = 0$, the execution is successful.

If $info = -i$, the i th argument had an illegal value.

If $info > 0$, and

if $i \leq n$, the algorithm failed to converge, and i off-diagonal elements of an intermediate tridiagonal did not converge to zero;

if $info = n + i$, for $1 \leq i \leq n$, then `spbstf/dpbstf` returned $info = i$ and B is not positive-definite. The factorization of B could not be completed and no eigenvalues or eigenvectors were computed.

?hbgv

Computes all eigenvalues and, optionally, eigenvectors of a complex generalized Hermitian definite eigenproblem with banded matrices.

```
call chbgv ( jobz, uplo, n, ka, kb, ab, ldab, bb, ldbb, w, z, ldz,
            work, rwork, info )
call zhbgv ( jobz, uplo, n, ka, kb, ab, ldab, bb, ldbb, w, z, ldz,
            work, rwork, info )
```

Discussion

This routine computes all the eigenvalues, and optionally, the eigenvectors of a complex generalized Hermitian-definite banded eigenproblem, of the form $Ax = \lambda Bx$. Here A and B are assumed to be Hermitian and banded, and B is also positive definite.

Input Parameters

jobz CHARACTER*1. Must be 'N' or 'V'.
If *jobz* = 'N', then compute eigenvalues only.
If *jobz* = 'V', then compute eigenvalues and eigenvectors.

uplo CHARACTER*1. Must be 'U' or 'L'.
If *uplo* = 'U', arrays *ab* and *bb* store the upper triangles of A and B ;
If *uplo* = 'L', arrays *ab* and *bb* store the lower triangles of A and B .

n INTEGER. The order of the matrices A and B ($n \geq 0$).

ka INTEGER. The number of super- or sub-diagonals in A ($ka \geq 0$).

kb INTEGER. The number of super- or sub-diagonals in B ($kb \geq 0$).

<i>ab,bb,work</i>	<p>COMPLEX for <i>chbgv</i> DOUBLE COMPLEX for <i>zhbgv</i></p> <p>Arrays: <i>ab (ldab,*)</i> is an array containing either upper or lower triangular part of the Hermitian matrix <i>A</i> (as specified by <i>uplo</i>) in band storage format. The second dimension of the array <i>ab</i> must be at least $\max(1, n)$. <i>bb (ldb,*)</i> is an array containing either upper or lower triangular part of the Hermitian matrix <i>B</i> (as specified by <i>uplo</i>) in band storage format. The second dimension of the array <i>bb</i> must be at least $\max(1, n)$. <i>work(*)</i> is a workspace array, DIMENSION at least $\max(1, n)$.</p>
<i>ldab</i>	<p>INTEGER. The first dimension of the array <i>ab</i>; must be at least <i>ka</i>+1.</p>
<i>ldb</i>	<p>INTEGER. The first dimension of the array <i>bb</i>; must be at least <i>kb</i>+1.</p>
<i>ldz</i>	<p>INTEGER. The leading dimension of the output array <i>z</i>; $ldz \geq 1$. If <i>jobz</i> = 'V', $ldz \geq \max(1, n)$.</p>
<i>rwork</i>	<p>REAL for <i>chbgv</i> DOUBLE PRECISION for <i>zhbgv</i>. Workspace array, DIMENSION at least $\max(1, 3n)$.</p>

Output Parameters

<i>ab</i>	On exit, the contents of <i>ab</i> are overwritten.
<i>bb</i>	On exit, contains the factor <i>S</i> from the split Cholesky factorization $B = S^H S$, as returned by <i>cpbstf/zpbstf</i> .
<i>w</i>	<p>REAL for <i>chbgv</i> DOUBLE PRECISION for <i>zhbgv</i>. Array, DIMENSION at least $\max(1, n)$. If <i>info</i> = 0, contains the eigenvalues in ascending order.</p>

z **COMPLEX** for **chbgv**
 DOUBLE COMPLEX for **zhbgv**
Array **z(ldz,*)** . The second dimension of **z** must be at least $\max(1, n)$.
If **jobz** = 'V', then if **info** = 0, **z** contains the matrix **Z** of eigenvectors, with the *i*-th column of **z** holding the eigenvector associated with $w(i)$. The eigenvectors are normalized so that $Z^H B Z = I$.
If **jobz** = 'N', then **z** is not referenced.

info **INTEGER**.
If **info** = 0, the execution is successful.
If **info** = -*i*, the *i*th argument had an illegal value.
If **info** > 0, and
 if $i \leq n$, the algorithm failed to converge, and *i* off-diagonal elements of an intermediate tridiagonal did not converge to zero;
 if **info** = $n + i$, for $1 \leq i \leq n$, then **cpbstf/zpbstf** returned **info** = *i* and **B** is not positive-definite. The factorization of **B** could not be completed and no eigenvalues or eigenvectors were computed.

?sbgvd

Computes all eigenvalues and, optionally, eigenvectors of a real generalized symmetric definite eigenproblem with banded matrices. If eigenvectors are desired, it uses a divide and conquer method.

```
call ssbgvd ( jobz, uplo, n, ka, kb, ab, ldab, bb, ldbb, w, z, ldz,
             work, lwork, iwork, liwork, info )
call dsbgvd ( jobz, uplo, n, ka, kb, ab, ldab, bb, ldbb, w, z, ldz,
             work, lwork, iwork, liwork, info )
```

Discussion

This routine computes all the eigenvalues, and optionally, the eigenvectors of a real generalized symmetric-definite banded eigenproblem, of the form $Ax = \lambda Bx$. Here A and B are assumed to be symmetric and banded, and B is also positive definite. If eigenvectors are desired, it uses a divide and conquer algorithm.

Input Parameters

<i>jobz</i>	CHARACTER*1. Must be 'N' or 'V'. If <i>jobz</i> = 'N', then compute eigenvalues only. If <i>jobz</i> = 'V', then compute eigenvalues and eigenvectors.
<i>uplo</i>	CHARACTER*1. Must be 'U' or 'L'. If <i>uplo</i> = 'U', arrays <i>ab</i> and <i>bb</i> store the upper triangles of A and B ; If <i>uplo</i> = 'L', arrays <i>ab</i> and <i>bb</i> store the lower triangles of A and B .
<i>n</i>	INTEGER. The order of the matrices A and B ($n \geq 0$).
<i>ka</i>	INTEGER. The number of super- or sub-diagonals in A ($ka \geq 0$).

<i>kb</i>	INTEGER. The number of super- or sub-diagonals in <i>B</i> ($kb \geq 0$).
<i>ab,bb,work</i>	REAL for <i>ssbgvd</i> DOUBLE PRECISION for <i>dsbgvd</i> Arrays: <i>ab</i> (<i>ldab</i> ,*) is an array containing either upper or lower triangular part of the symmetric matrix <i>A</i> (as specified by <i>uplo</i>) in band storage format. The second dimension of the array <i>ab</i> must be at least $\max(1, n)$. <i>bb</i> (<i>ldbb</i> ,*) is an array containing either upper or lower triangular part of the symmetric matrix <i>B</i> (as specified by <i>uplo</i>) in band storage format. The second dimension of the array <i>bb</i> must be at least $\max(1, n)$. <i>work</i> (<i>lwork</i>) is a workspace array.
<i>ldab</i>	INTEGER. The first dimension of the array <i>ab</i> ; must be at least $ka+1$.
<i>ldbb</i>	INTEGER. The first dimension of the array <i>bb</i> ; must be at least $kb+1$.
<i>ldz</i>	INTEGER. The leading dimension of the output array <i>z</i> ; $ldz \geq 1$. If <i>jobz</i> = 'V', $ldz \geq \max(1, n)$.
<i>lwork</i>	INTEGER. The dimension of the array <i>work</i> . Constraints: If $n \leq 1$, $lwork \geq 1$; If <i>jobz</i> = 'N' and $n > 1$, $lwork \geq 3n$; If <i>jobz</i> = 'V' and $n > 1$, $lwork \geq 2n^2+5n+1$.
<i>iwork</i>	INTEGER. Workspace array, DIMENSION (<i>liwork</i>).
<i>liwork</i>	INTEGER. The dimension of the array <i>iwork</i> . Constraints: If $n \leq 1$, $liwork \geq 1$; If <i>jobz</i> = 'N' and $n > 1$, $liwork \geq 1$; If <i>jobz</i> = 'V' and $n > 1$, $liwork \geq 5n+3$.

Output Parameters

<i>ab</i>	On exit, the contents of <i>ab</i> are overwritten.
<i>bb</i>	On exit, contains the factor S from the split Cholesky factorization $B = S^T S$, as returned by <code>spbstf/dpbstf</code> .
<i>w, z</i>	<p>REAL for <code>ssbgvd</code> DOUBLE PRECISION for <code>dsbgvd</code></p> <p>Arrays: <i>w</i>(*), DIMENSION at least $\max(1, n)$. If <i>info</i> = 0, contains the eigenvalues in ascending order. <i>z</i>(<i>ldz</i>,*) . The second dimension of <i>z</i> must be at least $\max(1, n)$. If <i>jobz</i> = 'V', then if <i>info</i> = 0, <i>z</i> contains the matrix Z of eigenvectors, with the <i>i</i>-th column of <i>z</i> holding the eigenvector associated with <i>w</i>(<i>i</i>). The eigenvectors are normalized so that $Z^T B Z = I$. If <i>jobz</i> = 'N', then <i>z</i> is not referenced.</p>
<i>work</i> (1)	On exit, if <i>info</i> = 0, then <i>work</i> (1) returns the required minimal size of <i>lwork</i> .
<i>iwork</i> (1)	On exit, if <i>info</i> = 0, then <i>iwork</i> (1) returns the required minimal size of <i>liwork</i> .
<i>info</i>	<p>INTEGER.</p> <p>If <i>info</i> = 0, the execution is successful. If <i>info</i> = -<i>i</i>, the <i>i</i>th argument had an illegal value. If <i>info</i> > 0, and</p> <ul style="list-style-type: none"> if $i \leq n$, the algorithm failed to converge, and <i>i</i> off-diagonal elements of an intermediate tridiagonal did not converge to zero; if <i>info</i> = $n + i$, for $1 \leq i \leq n$, then <code>spbstf/dpbstf</code> returned <i>info</i> = <i>i</i> and B is not positive-definite. The factorization of B could not be completed and no eigenvalues or eigenvectors were computed.

?hbgvd

Computes all eigenvalues and, optionally, eigenvectors of a complex generalized Hermitian definite eigenproblem with banded matrices. If eigenvectors are desired, it uses a divide and conquer method.

```
call chbgvd ( jobz, uplo, n, ka, kb, ab, ldab, bb, ldbb, w, z, ldz,  
             work, lwork, rwork, lrwork, iwork, liwork, info )  
call zhbgvd ( jobz, uplo, n, ka, kb, ab, ldab, bb, ldbb, w, z, ldz,  
             work, lwork, rwork, lrwork, iwork, liwork, info )
```

Discussion

This routine computes all the eigenvalues, and optionally, the eigenvectors of a complex generalized Hermitian-definite banded eigenproblem, of the form $Ax = \lambda Bx$. Here A and B are assumed to be Hermitian and banded, and B is also positive definite. If eigenvectors are desired, it uses a divide and conquer algorithm.

Input Parameters

jobz CHARACTER*1. Must be 'N' or 'V'.
If **jobz** = 'N', then compute eigenvalues only.
If **jobz** = 'V', then compute eigenvalues and eigenvectors.

uplo CHARACTER*1. Must be 'U' or 'L'.
If **uplo** = 'U', arrays **ab** and **bb** store the upper triangles of A and B ;
If **uplo** = 'L', arrays **ab** and **bb** store the lower triangles of A and B .

n INTEGER. The order of the matrices A and B ($n \geq 0$).

ka INTEGER. The number of super- or sub-diagonals in A ($ka \geq 0$).

<i>kb</i>	INTEGER. The number of super- or sub-diagonals in <i>B</i> ($kb \geq 0$).
<i>ab,bb,work</i>	COMPLEX for <code>chbgvd</code> DOUBLE COMPLEX for <code>zhbgvd</code> Arrays: <i>ab</i> (<i>ldab</i> ,*) is an array containing either upper or lower triangular part of the Hermitian matrix <i>A</i> (as specified by <i>uplo</i>) in band storage format. The second dimension of the array <i>ab</i> must be at least $\max(1, n)$. <i>bb</i> (<i>ldb</i> ,*) is an array containing either upper or lower triangular part of the Hermitian matrix <i>B</i> (as specified by <i>uplo</i>) in band storage format. The second dimension of the array <i>bb</i> must be at least $\max(1, n)$. <i>work</i> (<i>lwork</i>) is a workspace array.
<i>ldab</i>	INTEGER. The first dimension of the array <i>ab</i> ; must be at least $ka+1$.
<i>ldb</i>	INTEGER. The first dimension of the array <i>bb</i> ; must be at least $kb+1$.
<i>ldz</i>	INTEGER. The leading dimension of the output array <i>z</i> ; $ldz \geq 1$. If <i>jobz</i> = 'V', $ldz \geq \max(1, n)$.
<i>lwork</i>	INTEGER. The dimension of the array <i>work</i> . Constraints: If $n \leq 1$, $lwork \geq 1$; If <i>jobz</i> = 'N' and $n > 1$, $lwork \geq n$; If <i>jobz</i> = 'V' and $n > 1$, $lwork \geq 2n^2$.
<i>rwork</i>	REAL for <code>chbgvd</code> DOUBLE PRECISION for <code>zhbgvd</code> . Workspace array, DIMENSION (<i>lwork</i>).
<i>lrwork</i>	INTEGER. The dimension of the array <i>rwork</i> .

Constraints:

If $n \leq 1$, $lrwork \geq 1$;

If $jobz = 'N'$ and $n > 1$, $lrwork \geq n$;

If $jobz = 'V'$ and $n > 1$, $lrwork \geq 2n^2 + 5n + 1$.

iwork INTEGER.
Workspace array, DIMENSION (*liwork*).

liwork INTEGER. The dimension of the array *iwork*.

Constraints:

If $n \leq 1$, $liwork \geq 1$;

If $jobz = 'N'$ and $n > 1$, $liwork \geq 1$;

If $jobz = 'V'$ and $n > 1$, $liwork \geq 5n + 3$.

Output Parameters

ab On exit, the contents of *ab* are overwritten.

bb On exit, contains the factor S from the split Cholesky factorization $B = S^H S$, as returned by *cpbstf/zpbstf*.

w REAL for *chbgvd*
DOUBLE PRECISION for *zhbgvd*.
Array, DIMENSION at least $\max(1, n)$.
If *info* = 0, contains the eigenvalues in ascending order.

z COMPLEX for *chbgvd*
DOUBLE COMPLEX for *zhbgvd*
Array *z(ldz,*)*. The second dimension of *z* must be at least $\max(1, n)$.
If $jobz = 'V'$, then if *info* = 0, *z* contains the matrix Z of eigenvectors, with the *i*-th column of *z* holding the eigenvector associated with $w(i)$. The eigenvectors are normalized so that $Z^H B Z = I$.
If $jobz = 'N'$, then *z* is not referenced.

work(1) On exit, if *info* = 0, then *work(1)* returns the required minimal size of *lwork*.

rwork(1) On exit, if *info* = 0, then *rwork(1)* returns the required minimal size of *lrwork*.

iwork(1) On exit, if *info* = 0, then *iwork(1)* returns the required minimal size of *liwork*.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th argument had an illegal value.
If *info* > 0, and
if $i \leq n$, the algorithm failed to converge, and *i* off-diagonal elements of an intermediate tridiagonal did not converge to zero;
if $info = n + i$, for $1 \leq i \leq n$, then *cpbstf/zpbstf* returned *info* = *i* and *B* is not positive-definite. The factorization of *B* could not be completed and no eigenvalues or eigenvectors were computed.

?sbgvx

Computes selected eigenvalues and, optionally, eigenvectors of a real generalized symmetric definite eigenproblem with banded matrices.

```
call ssbgvx ( jobz, range, uplo, n, ka, kb, ab, ldab, bb, ldbb, q,
             ldq, vl, vu, il, iu, abstol, m, w, z, ldz, work, iwork,
             ifail, info )
call dsbgvx ( jobz, range, uplo, n, ka, kb, ab, ldab, bb, ldbb, q,
             ldq, vl, vu, il, iu, abstol, m, w, z, ldz, work, iwork,
             ifail, info )
```

Discussion

This routine computes selected eigenvalues, and optionally, the eigenvectors of a real generalized symmetric-definite banded eigenproblem, of the form $Ax = \lambda Bx$. Here A and B are assumed to be symmetric and banded, and B is also positive definite.

Eigenvalues and eigenvectors can be selected by specifying either all eigenvalues, a range of values or a range of indices for the desired eigenvalues.

Input Parameters

jobz CHARACTER*1. Must be 'N' or 'V'.
If *jobz* = 'N', then compute eigenvalues only.
If *jobz* = 'V', then compute eigenvalues and eigenvectors.

range CHARACTER*1. Must be 'A' or 'V' or 'I'.
If *range* = 'A', the routine computes all eigenvalues.
If *range* = 'V', the routine computes eigenvalues λ_i in the half-open interval: $vl < \lambda_i \leq vu$.
If *range* = 'I', the routine computes eigenvalues with indices *il* to *iu*.

uplo CHARACTER*1. Must be 'U' or 'L'.

If *uplo* = 'U', arrays *ab* and *bb* store the upper triangles of *A* and *B*;
 If *uplo* = 'L', arrays *ab* and *bb* store the lower triangles of *A* and *B*.

n **INTEGER**. The order of the matrices *A* and *B* ($n \geq 0$).

ka **INTEGER**. The number of super- or sub-diagonals in *A* ($ka \geq 0$).

kb **INTEGER**. The number of super- or sub-diagonals in *B* ($kb \geq 0$).

ab,bb,work **REAL** for **ssbgvx**
DOUBLE PRECISION for **dsbgvx**
 Arrays:
ab (*ldab*, *) is an array containing either upper or lower triangular part of the symmetric matrix *A* (as specified by *uplo*) in band storage format.
 The second dimension of the array *ab* must be at least $\max(1, n)$.
bb (*ldbb*, *) is an array containing either upper or lower triangular part of the symmetric matrix *B* (as specified by *uplo*) in band storage format.
 The second dimension of the array *bb* must be at least $\max(1, n)$.
work(*) is a workspace array, **DIMENSION** at least $\max(1, 7n)$.

ldab **INTEGER**. The first dimension of the array *ab*; must be at least $ka+1$.

ldbb **INTEGER**. The first dimension of the array *bb*; must be at least $kb+1$.

vl, vu **REAL** for **ssbgvx**
DOUBLE PRECISION for **dsbgvx**.
 If *range* = 'V', the lower and upper bounds of the interval to be searched for eigenvalues.
 Constraint: $vl < vu$.
 If *range* = 'A' or 'I', *vl* and *vu* are not referenced.

il, iu **INTEGER.**
 If *range* = 'I', the indices in ascending order of the smallest and largest eigenvalues to be returned.
 Constraint: $1 \leq i_l \leq i_u \leq n$, if $n > 0$; $i_l = 1$ and $i_u = 0$ if $n = 0$.
 If *range* = 'A' or 'V', *il* and *iu* are not referenced.

abstol **REAL** for *ssbgvx*
DOUBLE PRECISION for *dsbgvx*.
 The absolute error tolerance for the eigenvalues.
 See *Application Notes* for more information.

ldz **INTEGER.** The leading dimension of the output array *z*;
 $ldz \geq 1$. If *jobz* = 'V', $ldz \geq \max(1, n)$.

ldq **INTEGER.** The leading dimension of the output array *q*;
 $ldq \geq 1$. If *jobz* = 'V', $ldq \geq \max(1, n)$.

iwork **INTEGER.**
 Workspace array, **DIMENSION** at least $\max(1, 5n)$.

Output Parameters

ab On exit, the contents of *ab* are overwritten.

bb On exit, contains the factor *S* from the split Cholesky factorization $B = S^T S$, as returned by *spbstf/dpbstf*.

m **INTEGER.** The total number of eigenvalues found,
 $0 \leq m \leq n$. If *range* = 'A', $m = n$, and if *range* = 'I',
 $m = i_u - i_l + 1$.

w, z, q **REAL** for *ssbgvx*
DOUBLE PRECISION for *dsbgvx*
 Arrays:
w(*), **DIMENSION** at least $\max(1, n)$.
 If *info* = 0, contains the eigenvalues in ascending order.
z(*ldz*,*). The second dimension of *z* must be at least $\max(1, n)$.
 If *jobz* = 'V', then if *info* = 0, *z* contains the matrix *Z* of eigenvectors, with the *i*-th column of *z* holding the eigenvector associated with *w*(*i*). The eigenvectors are

normalized so that $Z^T B Z = I$.

If `jobz = 'N'`, then `z` is not referenced.

`q(ldq,*)`. The second dimension of `q` must be at least $\max(1, n)$.

If `jobz = 'V'`, then `q` contains the n -by- n matrix used in the reduction of $Ax = \lambda Bx$ to standard form, that is, $Cx = \lambda x$ and consequently C to tridiagonal form.

If `jobz = 'N'`, then `q` is not referenced.

`ifail`

INTEGER.

Array, DIMENSION at least $\max(1, n)$.

If `jobz = 'V'`, then if `info = 0`, the first m elements of `ifail` are zero; if `info > 0`, the `ifail` contains the indices of the eigenvectors that failed to converge.

If `jobz = 'N'`, then `ifail` is not referenced.

`info`

INTEGER.

If `info = 0`, the execution is successful.

If `info = -i`, the i th argument had an illegal value.

If `info > 0`, and

if $i \leq n$, the algorithm failed to converge, and i off-diagonal elements of an intermediate tridiagonal did not converge to zero;

if $info = n + i$, for $1 \leq i \leq n$, then `spbstf/dpbstf` returned $info = i$ and B is not positive-definite. The factorization of B could not be completed and no eigenvalues or eigenvectors were computed.

Application Notes

An approximate eigenvalue is accepted as converged when it is determined to lie in an interval $[a,b]$ of width less than or equal to

$abstol + \epsilon * \max(|a|, |b|)$, where ϵ is the machine precision. If `abstol` is less than or equal to zero, then $\epsilon * \|T\|_1$ will be used in its place, where T is the tridiagonal matrix obtained by reducing A to tridiagonal form.

Eigenvalues will be computed most accurately when `abstol` is set to twice the underflow threshold $2 * ?lamch('S')$, not zero. If this routine returns with `info > 0`, indicating that some eigenvectors did not converge, try setting `abstol` to $2 * ?lamch('S')$.

?hbgvx

Computes selected eigenvalues and, optionally, eigenvectors of a complex generalized Hermitian definite eigenproblem with banded matrices.

```
call chbgvx ( jobz, range, uplo, n, ka, kb, ab, ldab, bb, ldbb, q,
             ldq, vl, vu, il, iu, abstol, m, w, z, ldz, work, rwork,
             iwork, ifail, info )
call zhbgvx ( jobz, range, uplo, n, ka, kb, ab, ldab, bb, ldbb, q,
             ldq, vl, vu, il, iu, abstol, m, w, z, ldz, work, rwork,
             iwork, ifail, info )
```

Discussion

This routine computes selected eigenvalues, and optionally, the eigenvectors of a complex generalized Hermitian-definite banded eigenproblem, of the form $Ax = \lambda Bx$. Here A and B are assumed to be Hermitian and banded, and B is also positive definite.

Eigenvalues and eigenvectors can be selected by specifying either all eigenvalues, a range of values or a range of indices for the desired eigenvalues.

Input Parameters

jobz CHARACTER*1. Must be 'N' or 'V'.
If *jobz* = 'N', then compute eigenvalues only.
If *jobz* = 'V', then compute eigenvalues and eigenvectors.

range CHARACTER*1. Must be 'A' or 'V' or 'I'.
If *range* = 'A', the routine computes all eigenvalues.
If *range* = 'V', the routine computes eigenvalues λ_i in the half-open interval: $vl < \lambda_i \leq vu$.
If *range* = 'I', the routine computes eigenvalues with indices *il* to *iu*.

uplo CHARACTER*1. Must be 'U' or 'L'.

If *uplo* = 'U', arrays *ab* and *bb* store the upper triangles of *A* and *B*;

If *uplo* = 'L', arrays *ab* and *bb* store the lower triangles of *A* and *B*.

<i>n</i>	INTEGER . The order of the matrices <i>A</i> and <i>B</i> ($n \geq 0$).
<i>ka</i>	INTEGER . The number of super- or sub-diagonals in <i>A</i> ($ka \geq 0$).
<i>kb</i>	INTEGER . The number of super- or sub-diagonals in <i>B</i> ($kb \geq 0$).
<i>ab, bb, work</i>	COMPLEX for chbgvx DOUBLE COMPLEX for zhbgvx Arrays: <i>ab</i> (<i>ldab</i> , *) is an array containing either upper or lower triangular part of the Hermitian matrix <i>A</i> (as specified by <i>uplo</i>) in band storage format. The second dimension of the array <i>ab</i> must be at least $\max(1, n)$. <i>bb</i> (<i>ldbb</i> , *) is an array containing either upper or lower triangular part of the Hermitian matrix <i>B</i> (as specified by <i>uplo</i>) in band storage format. The second dimension of the array <i>bb</i> must be at least $\max(1, n)$. <i>work</i> (*) is a workspace array, DIMENSION at least $\max(1, n)$.
<i>ldab</i>	INTEGER . The first dimension of the array <i>ab</i> ; must be at least $ka+1$.
<i>ldbb</i>	INTEGER . The first dimension of the array <i>bb</i> ; must be at least $kb+1$.
<i>vl, vu</i>	REAL for chbgvx DOUBLE PRECISION for zhbgvx . If <i>range</i> = 'V', the lower and upper bounds of the interval to be searched for eigenvalues. Constraint: $vl < vu$. If <i>range</i> = 'A' or 'I', <i>vl</i> and <i>vu</i> are not referenced.

<i>il, iu</i>	<p>INTEGER.</p> <p>If <i>range</i> = 'I', the indices in ascending order of the smallest and largest eigenvalues to be returned.</p> <p>Constraint: $1 \leq il \leq iu \leq n$, if $n > 0$; $il=1$ and $iu=0$ if $n = 0$.</p> <p>If <i>range</i> = 'A' or 'V', <i>il</i> and <i>iu</i> are not referenced.</p>
<i>abstol</i>	<p>REAL for <i>chbgvx</i></p> <p>DOUBLE PRECISION for <i>zhbgvx</i>.</p> <p>The absolute error tolerance for the eigenvalues.</p> <p>See <i>Application Notes</i> for more information.</p>
<i>ldz</i>	<p>INTEGER. The leading dimension of the output array <i>z</i>;</p> <p>$ldz \geq 1$. If <i>jobz</i> = 'V', $ldz \geq \max(1, n)$.</p>
<i>ldq</i>	<p>INTEGER. The leading dimension of the output array <i>q</i>;</p> <p>$ldq \geq 1$. If <i>jobz</i> = 'V', $ldq \geq \max(1, n)$.</p>
<i>rwork</i>	<p>REAL for <i>chbgvx</i></p> <p>DOUBLE PRECISION for <i>zhbgvx</i>.</p> <p>Workspace array, DIMENSION at least $\max(1, 7n)$.</p>
<i>iwork</i>	<p>INTEGER.</p> <p>Workspace array, DIMENSION at least $\max(1, 5n)$.</p>

Output Parameters

<i>ab</i>	On exit, the contents of <i>ab</i> are overwritten.
<i>bb</i>	On exit, contains the factor <i>S</i> from the split Cholesky factorization $B = S^H S$, as returned by <i>cpbstf/zpbstf</i> .
<i>m</i>	<p>INTEGER. The total number of eigenvalues found,</p> <p>$0 \leq m \leq n$. If <i>range</i> = 'A', $m = n$, and if <i>range</i> = 'I', $m = iu - il + 1$.</p>
<i>w</i>	<p>REAL for <i>chbgvx</i></p> <p>DOUBLE PRECISION for <i>zhbgvx</i>.</p> <p>Array <i>w</i>(*), DIMENSION at least $\max(1, n)$.</p> <p>If <i>info</i> = 0, contains the eigenvalues in ascending order.</p>
<i>z, q</i>	<p>COMPLEX for <i>chbgvx</i></p> <p>DOUBLE COMPLEX for <i>zhbgvx</i></p> <p>Arrays:</p>

$z(ldz, *)$. The second dimension of z must be at least $\max(1, n)$.

If $jobz = 'V'$, then if $info = 0$, z contains the matrix Z of eigenvectors, with the i -th column of z holding the eigenvector associated with $w(i)$. The eigenvectors are normalized so that $Z^H B Z = I$.

If $jobz = 'N'$, then z is not referenced.

$q(ldq, *)$. The second dimension of q must be at least $\max(1, n)$.

If $jobz = 'V'$, then q contains the n -by- n matrix used in the reduction of $Ax = \lambda Bx$ to standard form, that is, $Cx = \lambda x$ and consequently C to tridiagonal form.

If $jobz = 'N'$, then q is not referenced.

$ifail$

INTEGER.

Array, DIMENSION at least $\max(1, n)$.

If $jobz = 'V'$, then if $info = 0$, the first m elements of $ifail$ are zero; if $info > 0$, the $ifail$ contains the indices of the eigenvectors that failed to converge.

If $jobz = 'N'$, then $ifail$ is not referenced.

$info$

INTEGER.

If $info = 0$, the execution is successful.

If $info = -i$, the i th argument had an illegal value.

If $info > 0$, and

if $i \leq n$, the algorithm failed to converge, and i off-diagonal elements of an intermediate tridiagonal did not converge to zero;

if $info = n + i$, for $1 \leq i \leq n$, then $cpbstf/zpbstf$ returned $info = i$ and B is not positive-definite. The factorization of B could not be completed and no eigenvalues or eigenvectors were computed.

Application Notes

An approximate eigenvalue is accepted as converged when it is determined to lie in an interval $[a,b]$ of width less than or equal to

$abstol + \epsilon * \max(|a|, |b|)$, where ϵ is the machine precision. If $abstol$ is less than or equal to zero, then $\epsilon * \|T\|_1$ will be used in its place, where T

is the tridiagonal matrix obtained by reducing A to tridiagonal form. Eigenvalues will be computed most accurately when `abstol` is set to twice the underflow threshold $2 * \text{?lamch}('S')$, not zero. If this routine returns with `info` > 0 , indicating that some eigenvectors did not converge, try setting `abstol` to $2 * \text{?lamch}('S')$.

Generalized Nonsymmetric Eigenproblems

This section describes LAPACK driver routines used for solving generalized nonsymmetric eigenproblems. See also [computational routines](#) that can be called to solve these problems.

[Table 5-14](#) lists routines described in more detail below.

Table 5-14 Driver Routines for Solving Generalized Nonsymmetric Eigenproblems

Routine Name	Operation performed
?gges	Computes the generalized eigenvalues, Schur form, and the left and/or right Schur vectors for a pair of nonsymmetric matrices.
?ggesx	Computes the generalized eigenvalues, Schur form, and, optionally, the left and/or right matrices of Schur vectors .
?ggeev	Computes the generalized eigenvalues, and the left and/or right generalized eigenvectors for a pair of nonsymmetric matrices.
?ggeevx	Computes the generalized eigenvalues, and, optionally, the left and/or right generalized eigenvectors.

[?gges](#)

Computes the generalized eigenvalues, Schur form, and the left and/or right Schur vectors for a pair of nonsymmetric matrices.

```
call ssges ( jobvsl, jobvsr, sort, selctg, n, a, lda, b, ldb, sdim,
            alphas, alphas, beta, vsl, ldvsl, vsr, ldvsr, work,
            lwork, bwork, info )
call dsges ( jobvsl, jobvsr, sort, selctg, n, a, lda, b, ldb, sdim,
            alphas, alphas, beta, vsl, ldvsl, vsr, ldvsr, work,
            lwork, bwork, info )
call csges ( jobvsl, jobvsr, sort, selctg, n, a, lda, b, ldb, sdim,
            alpha, beta, vsl, ldvsl, vsr, ldvsr, work, lwork, rwork,
            bwork, info )
```

```
call zgges ( jobvsl, jobvsr, sort, selctg, n, a, lda, b, ldb, sdim,
            alpha, beta, vs1, ldvsl, vsr, ldvsr, work, lwork, rwork,
            bwork, info )
```

Discussion

This routine computes for a pair of n -by- n real/complex nonsymmetric matrices (A,B) , the generalized eigenvalues, the generalized real/complex Schur form (S,T) , optionally, the left and/or right matrices of Schur vectors ($vs1$ and vsr). This gives the generalized Schur factorization

$$(A,B) = (vs1 * S * vsr^H, vs1 * T * vsr^H)$$

Optionally, it also orders the eigenvalues so that a selected cluster of eigenvalues appears in the leading diagonal blocks of the upper quasi-triangular matrix S and the upper triangular matrix T . The leading columns of $vs1$ and vsr then form an orthonormal/unitary basis for the corresponding left and right eigenspaces (deflating subspaces).

(If only the generalized eigenvalues are needed, use the driver `?ggeev` instead, which is faster.)

A generalized eigenvalue for a pair of matrices (A,B) is a scalar w or a ratio $alpha / beta = w$, such that $A - w * B$ is singular. It is usually represented as the pair $(alpha, beta)$, as there is a reasonable interpretation for $beta=0$ or for both being zero.

A pair of matrices (S,T) is in generalized real Schur form if T is upper triangular with non-negative diagonal and S is block upper triangular with 1-by-1 and 2-by-2 blocks. 1-by-1 blocks correspond to real generalized eigenvalues, while 2-by-2 blocks of S will be "standardized" by making the corresponding elements of T have the form:

$$\begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}$$

and the pair of corresponding 2-by-2 blocks in S and T will have a complex conjugate pair of generalized eigenvalues.

A pair of matrices (S,T) is in generalized complex Schur form if S and T are upper triangular and, in addition, the diagonal of T are non-negative real numbers.

Input Parameters

- jobvsl* CHARACTER*1. Must be 'N' or 'V'.
 If *jobvsl* = 'N', then the left Schur vectors are not computed.
 If *jobvsl* = 'V', then the left Schur vectors are computed.
- jobvsr* CHARACTER*1. Must be 'N' or 'V'.
 If *jobvsr* = 'N', then the right Schur vectors are not computed.
 If *jobvsr* = 'V', then the right Schur vectors are computed.
- sort* CHARACTER*1. Must be 'N' or 'S'.
 Specifies whether or not to order the eigenvalues on the diagonal of the generalized Schur form.
 If *sort* = 'N', then eigenvalues are not ordered.
 If *sort* = 'S', eigenvalues are ordered (see *selctg*).
- selctg* LOGICAL FUNCTION of three REAL arguments for real flavors.
 LOGICAL FUNCTION of two COMPLEX arguments for complex flavors.
selctg must be declared EXTERNAL in the calling subroutine.
 If *sort* = 'S', *selctg* is used to select eigenvalues to sort to the top left of the Schur form.
 If *sort* = 'N', *selctg* is not referenced.
- For real flavors:*
 An eigenvalue ($\text{alphan}(j) + \text{alphai}(j)/\text{betan}(j)$) is selected if *selctg*(*alphan*(*j*), *alphai*(*j*), *betan*(*j*)) is true; that is, if either one of a complex conjugate pair of eigenvalues is selected, then both complex eigenvalues are selected.
 Note that in the ill-conditioned case, a selected complex eigenvalue may no longer satisfy *selctg*(*alphan*(*j*), *alphai*(*j*), *betan*(*j*)) = .TRUE. after ordering. In this case *info* is set to *n*+2.

For complex flavors:

An eigenvalue $\alpha(j) / \beta(j)$ is selected if $\text{selctg}(\alpha(j), \beta(j))$ is true.

Note that a selected complex eigenvalue may no longer satisfy $\text{selctg}(\alpha(j), \beta(j)) = \text{.TRUE.}$ after ordering, since ordering may change the value of complex eigenvalues (especially if the eigenvalue is ill-conditioned); in this case info is set to $n+2$ (see info below).

n INTEGER. The order of the matrices A , B , vs1 , and vsr ($n \geq 0$).

a, b, work REAL for sgges
DOUBLE PRECISION for dggcs
COMPLEX for cggcs
DOUBLE COMPLEX for zggcs .

Arrays:

$a(\text{lda}, *)$ is an array containing the n -by- n matrix A (first of the pair of matrices).

The second dimension of a must be at least $\max(1, n)$.

$b(\text{ldb}, *)$ is an array containing the n -by- n matrix B (second of the pair of matrices).

The second dimension of b must be at least $\max(1, n)$.

$\text{work}(\text{lwork})$ is a workspace array.

lda INTEGER. The first dimension of the array a .
Must be at least $\max(1, n)$.

ldb INTEGER. The first dimension of the array b .
Must be at least $\max(1, n)$.

$\text{ldvs1}, \text{ldvsr}$ INTEGER. The first dimensions of the output matrices vs1 and vsr , respectively. Constraints:
 $\text{ldvs1} \geq 1$. If $\text{jobvs1} = 'V'$, $\text{ldvs1} \geq \max(1, n)$.
 $\text{ldvsr} \geq 1$. If $\text{jobvsr} = 'V'$, $\text{ldvsr} \geq \max(1, n)$.

lwork INTEGER. The dimension of the array work .

$lwork \geq \max(1, 8n+16)$ for real flavors;
 $lwork \geq \max(1, 2n)$ for complex flavors.
 For good performance, $lwork$ must generally be larger.

$rwork$ REAL for `cgges`
 DOUBLE PRECISION for `zgges`
 Workspace array, DIMENSION at least $\max(1, 8n)$.
 This array is used in complex flavors only.

$bwork$ LOGICAL.
 Workspace array, DIMENSION at least $\max(1, n)$.
 Not referenced if `sort = 'N'`.

Output Parameters

a On exit, this array has been overwritten by its generalized Schur form S .

b On exit, this array has been overwritten by its generalized Schur form T .

$sdim$ INTEGER.
 If `sort = 'N'`, $sdim = 0$.
 If `sort = 'S'`, $sdim$ is equal to the number of eigenvalues (after sorting) for which `selctg` is true. Note that for real flavors complex conjugate pairs for which `selctg` is true for either eigenvalue count as 2.

$alphan, alphai$ REAL for `sgges`;
 DOUBLE PRECISION for `dgges`.
 Arrays, DIMENSION at least $\max(1, n)$ each. Contain values that form generalized eigenvalues in real flavors. See `beta`.

$alpha$ COMPLEX for `cgges`;
 DOUBLE COMPLEX for `zgges`.
 Array, DIMENSION at least $\max(1, n)$. Contain values that form generalized eigenvalues in complex flavors. See `beta`.

$beta$ REAL for `sgges`
 DOUBLE PRECISION for `dgges`
 COMPLEX for `cgges`

DOUBLE COMPLEX for *zgges*.

Array, *DIMENSION* at least $\max(1, n)$.

For real flavors:

On exit, $(\mathit{alpha}(j) + \mathit{alphai}(j)*i)/\mathit{beta}(j)$, $j=1, \dots, n$, will be the generalized eigenvalues.

$\mathit{alpha}(j) + \mathit{alphai}(j)*i$ and $\mathit{beta}(j)$, $j=1, \dots, n$ are the diagonals of the complex Schur form (S, T) that would result if the 2-by-2 diagonal blocks of the real generalized Schur form of (A, B) were further reduced to triangular form using complex unitary transformations. If $\mathit{alphai}(j)$ is zero, then the j -th eigenvalue is real; if positive, then the j -th and $(j+1)$ -st eigenvalues are a complex conjugate pair, with $\mathit{alphai}(j+1)$ negative.

For complex flavors:

On exit, $\mathit{alpha}(j)/\mathit{beta}(j)$, $j=1, \dots, n$, will be the generalized eigenvalues. $\mathit{alpha}(j)$, $j=1, \dots, n$, and $\mathit{beta}(j)$, $j=1, \dots, n$, are the diagonals of the complex Schur form (S, T) output by *cgges/zgges*. The $\mathit{beta}(j)$ will be non-negative real.

See also *Application Notes* below.

vsl, vsr

REAL for *sgges*

DOUBLE PRECISION for *dgges*

COMPLEX for *cgges*

DOUBLE COMPLEX for *zgges*.

Arrays:

*vsl(ldvsl, *)*, the second dimension of *vsl* must be at least $\max(1, n)$.

If *jobvsl* = 'V', this array will contain the left Schur vectors.

If *jobvsl* = 'N', *vsl* is not referenced.

*vsr(ldvsr, *)*, the second dimension of *vsr* must be at least $\max(1, n)$.

If *jobvsr* = 'V', this array will contain the right Schur vectors.

If *jobvsr* = 'N', *vsr* is not referenced.

work(1) On exit, if *info* = 0, then *work(1)* returns the required minimal size of *lwork*.

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.
 If *info* = *i*, and
i ≤ *n* :
 the *QZ* iteration failed. (*A,B*) is not in Schur form, but *alphan*(*j*), *alphai*(*j*) (for real flavors), or *alpha*(*j*) (for complex flavors), and *beta*(*j*), *j*=*info*+1,...,*n* should be correct.
i > *n* : errors that usually indicate LAPACK problems:
i = *n*+1: other than *QZ* iteration failed in ?*hgeqz*;
i = *n*+2: after reordering, roundoff changed values of some complex eigenvalues so that leading eigenvalues in the generalized Schur form no longer satisfy *selctg* = .TRUE.. This could also be caused due to scaling;
i = *n*+3: reordering failed in ?*tgse*n.

Application Notes

If you are in doubt how much workspace to supply for the array *work*, use a generous value of *lwork* for the first run. On exit, examine *work(1)* and use this value for subsequent runs.

The quotients *alphan*(*j*)/*beta*(*j*) and *alphai*(*j*)/*beta*(*j*) may easily overflow or underflow, and *beta*(*j*) may even be zero. Thus, you should avoid simply computing the ratio. However, *alphan* and *alphai* will be always less than and usually comparable with norm(*A*) in magnitude, and *beta* always less than and usually comparable with norm(*B*).

?ggesx

Computes the generalized eigenvalues, Schur form, and, optionally, the left and/or right matrices of Schur vectors.

```
call sggesx (jobvsl, jobvsr, sort, selctg, sense, n, a, lda, b, ldb,
            sdim, alphas, alphas, beta, vsl, ldvsl, vsr, ldvsr,
            rconde, rcondv, work, lwork, iwork, liwork, bwork, info )
call dggesx (jobvsl, jobvsr, sort, selctg, sense, n, a, lda, b, ldb,
            sdim, alphas, alphas, beta, vsl, ldvsl, vsr, ldvsr,
            rconde, rcondv, work, lwork, iwork, liwork, bwork, info )
call cggesx (jobvsl, jobvsr, sort, selctg, sense, n, a, lda, b, ldb,
            sdim, alpha, beta, vsl, ldvsl, vsr, ldvsr, rconde, rcondv,
            work, lwork, rwork, iwork, liwork, bwork, info )
call zggesx (jobvsl, jobvsr, sort, selctg, sense, n, a, lda, b, ldb,
            sdim, alpha, beta, vsl, ldvsl, vsr, ldvsr, rconde, rcondv,
            work, lwork, rwork, iwork, liwork, bwork, info )
```

Discussion

This routine computes for a pair of n -by- n real/complex nonsymmetric matrices (A,B) , the generalized eigenvalues, the generalized real/complex Schur form (S,T) , optionally, the left and/or right matrices of Schur vectors (vsl and vsr). This gives the generalized Schur factorization

$$(A,B) = (vsl * S * vsr^H, vsl * T * vsr^H)$$

Optionally, it also orders the eigenvalues so that a selected cluster of eigenvalues appears in the leading diagonal blocks of the upper quasi-triangular matrix S and the upper triangular matrix T ; computes a reciprocal condition number for the average of the selected eigenvalues ($rconde$); and computes a reciprocal condition number for the right and left deflating subspaces corresponding to the selected eigenvalues ($rcondv$). The leading columns of vsl and vsr then form an orthonormal/unitary basis for the corresponding left and right eigenspaces (deflating subspaces).

A generalized eigenvalue for a pair of matrices (A,B) is a scalar w or a ratio $\alpha / \beta = w$, such that $A - w*B$ is singular. It is usually represented as the pair (α, β) , as there is a reasonable interpretation for $\beta=0$ or for both being zero.

A pair of matrices (S,T) is in generalized real Schur form if T is upper triangular with non-negative diagonal and S is block upper triangular with 1-by-1 and 2-by-2 blocks. 1-by-1 blocks correspond to real generalized eigenvalues, while 2-by-2 blocks of S will be “standardized” by making the corresponding elements of T have the form:

$$\begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix}$$

and the pair of corresponding 2-by-2 blocks in S and T will have a complex conjugate pair of generalized eigenvalues.

A pair of matrices (S,T) is in generalized complex Schur form if S and T are upper triangular and, in addition, the diagonal of T are non-negative real numbers.

Input Parameters

<i>jobvs1</i>	<p>CHARACTER*1. Must be 'N' or 'V'.</p> <p>If <i>jobvs1</i> = 'N', then the left Schur vectors are not computed.</p> <p>If <i>jobvs1</i> = 'V', then the left Schur vectors are computed.</p>
<i>jobvsr</i>	<p>CHARACTER*1. Must be 'N' or 'V'.</p> <p>If <i>jobvsr</i> = 'N', then the right Schur vectors are not computed.</p> <p>If <i>jobvsr</i> = 'V', then the right Schur vectors are computed.</p>
<i>sort</i>	<p>CHARACTER*1. Must be 'N' or 'S'.</p> <p>Specifies whether or not to order the eigenvalues on the diagonal of the generalized Schur form.</p> <p>If <i>sort</i> = 'N', then eigenvalues are not ordered.</p> <p>If <i>sort</i> = 'S', eigenvalues are ordered (see <i>selctg</i>).</p>

selctg LOGICAL FUNCTION of three REAL arguments for real flavors.
LOGICAL FUNCTION of two COMPLEX arguments for complex flavors.
selctg must be declared EXTERNAL in the calling subroutine.
If *sort* = 'S', *selctg* is used to select eigenvalues to sort to the top left of the Schur form.
If *sort* = 'N', *selctg* is not referenced.
For real flavors:
An eigenvalue $(\text{alphan}(j) + \text{alphai}(j))/\text{betan}(j)$ is selected if *selctg*(*alphan*(*j*), *alphai*(*j*), *betan*(*j*)) is true; that is, if either one of a complex conjugate pair of eigenvalues is selected, then both complex eigenvalues are selected.
Note that in the ill-conditioned case, a selected complex eigenvalue may no longer satisfy *selctg*(*alphan*(*j*), *alphai*(*j*), *betan*(*j*)) = .TRUE. after ordering. In this case *info* is set to *n*+2.
For complex flavors:
An eigenvalue $\text{alpha}(j) / \text{betan}(j)$ is selected if *selctg*(*alpha*(*j*), *betan*(*j*)) is true.
Note that a selected complex eigenvalue may no longer satisfy *selctg*(*alpha*(*j*), *betan*(*j*)) = .TRUE. after ordering, since ordering may change the value of complex eigenvalues (especially if the eigenvalue is ill-conditioned); in this case *info* is set to *n*+2 (see *info* below).

sense CHARACTER*1. Must be 'N', 'E', 'V', or 'B'.
Determines which reciprocal condition number are computed.
If *sense* = 'N', none are computed;
If *sense* = 'E', computed for average of selected eigenvalues only;
If *sense* = 'V', computed for selected deflating subspaces only;

	If <i>sense</i> = 'B', computed for both. If <i>sense</i> is 'E', 'V', or 'B', then <i>sort</i> must equal 'S'.
<i>n</i>	INTEGER. The order of the matrices <i>A</i> , <i>B</i> , <i>vs1</i> , and <i>vsr</i> ($n \geq 0$).
<i>a</i> , <i>b</i> , <i>work</i>	REAL for <i>sggesx</i> DOUBLE PRECISION for <i>dggexx</i> COMPLEX for <i>cggexx</i> DOUBLE COMPLEX for <i>zggexx</i> . Arrays: <i>a</i> (<i>lda</i> ,*) is an array containing the <i>n</i> -by- <i>n</i> matrix <i>A</i> (first of the pair of matrices). The second dimension of <i>a</i> must be at least $\max(1, n)$. <i>b</i> (<i>ldb</i> ,*) is an array containing the <i>n</i> -by- <i>n</i> matrix <i>B</i> (second of the pair of matrices). The second dimension of <i>b</i> must be at least $\max(1, n)$. <i>work</i> (<i>lwork</i>) is a workspace array.
<i>lda</i>	INTEGER. The first dimension of the array <i>a</i> . Must be at least $\max(1, n)$.
<i>ldb</i>	INTEGER. The first dimension of the array <i>b</i> . Must be at least $\max(1, n)$.
<i>ldvs1</i> , <i>ldvsr</i>	INTEGER. The first dimensions of the output matrices <i>vs1</i> and <i>vsr</i> , respectively. Constraints: <i>ldvs1</i> ≥ 1 . If <i>jobvs1</i> = 'V', <i>ldvs1</i> $\geq \max(1, n)$. <i>ldvsr</i> ≥ 1 . If <i>jobvsr</i> = 'V', <i>ldvsr</i> $\geq \max(1, n)$.
<i>lwork</i>	INTEGER. The dimension of the array <i>work</i> . For real flavors: <i>lwork</i> $\geq \max(1, 8(n+1)+16)$; if <i>sense</i> = 'E', 'V', or 'B', then <i>lwork</i> $\geq \max(8(n+1)+16, 2*sdim*(n-sdim))$. For complex flavors: <i>lwork</i> $\geq \max(1, 2n)$; if <i>sense</i> = 'E', 'V', or 'B', then <i>lwork</i> $\geq \max(2n, 2*sdim*(n-sdim))$.

For good performance, *lwork* must generally be larger.

<i>rwork</i>	REAL for <i>cggesx</i> DOUBLE PRECISION for <i>zggesx</i> Workspace array, DIMENSION at least $\max(1, 8n)$. This array is used in complex flavors only.
<i>iwork</i>	INTEGER. Workspace array, DIMENSION (<i>liwork</i>). Not referenced if <i>sense</i> = 'N'.
<i>liwork</i>	INTEGER. The dimension of the array <i>iwork</i> . <i>liwork</i> \geq <i>n</i> +6 for real flavors; <i>liwork</i> \geq <i>n</i> +2 for complex flavors.
<i>bwork</i>	LOGICAL. Workspace array, DIMENSION at least $\max(1, n)$. Not referenced if <i>sort</i> = 'N'.

Output Parameters

<i>a</i>	On exit, this array has been overwritten by its generalized Schur form <i>S</i> .
<i>b</i>	On exit, this array has been overwritten by its generalized Schur form <i>T</i> .
<i>sdim</i>	INTEGER. If <i>sort</i> = 'N', <i>sdim</i> = 0. If <i>sort</i> = 'S', <i>sdim</i> is equal to the number of eigenvalues (after sorting) for which <i>selctg</i> is true. Note that for real flavors complex conjugate pairs for which <i>selctg</i> is true for either eigenvalue count as 2.
<i>alphan</i> , <i>alpha</i> _{<i>i</i>}	REAL for <i>sggesx</i> ; DOUBLE PRECISION for <i>dggesx</i> . Arrays, DIMENSION at least $\max(1, n)$ each. Contain values that form generalized eigenvalues in real flavors. See <i>beta</i> .

alpha **COMPLEX** for *cggesx*;
DOUBLE COMPLEX for *zggesx*.
 Array, **DIMENSION** at least $\max(1, n)$. Contain values that form generalized eigenvalues in complex flavors. See *beta*.

beta **REAL** for *sggesx*
DOUBLE PRECISION for *dggesx*
COMPLEX for *cggesx*
DOUBLE COMPLEX for *zggesx*.
 Array, **DIMENSION** at least $\max(1, n)$.
For real flavors:
 On exit, $(\mathit{alphar}(j) + \mathit{alphai}(j)*i)/\mathit{beta}(j)$, $j=1, \dots, n$, will be the generalized eigenvalues.
 $\mathit{alphar}(j) + \mathit{alphai}(j)*i$ and $\mathit{beta}(j)$, $j=1, \dots, n$ are the diagonals of the complex Schur form (S, T) that would result if the 2-by-2 diagonal blocks of the real generalized Schur form of (A, B) were further reduced to triangular form using complex unitary transformations. If $\mathit{alphai}(j)$ is zero, then the j -th eigenvalue is real; if positive, then the j -th and $(j+1)$ -st eigenvalues are a complex conjugate pair, with $\mathit{alphai}(j+1)$ negative.
For complex flavors:
 On exit, $\mathit{alpha}(j)/\mathit{beta}(j)$, $j=1, \dots, n$, will be the generalized eigenvalues. $\mathit{alpha}(j)$, $j=1, \dots, n$, and $\mathit{beta}(j)$, $j=1, \dots, n$, are the diagonals of the complex Schur form (S, T) output by *cggesx/zggesx*. The $\mathit{beta}(j)$ will be non-negative real.
 See also *Application Notes* below.

vsl, vsr **REAL** for *sggesx*
DOUBLE PRECISION for *dggesx*
COMPLEX for *cggesx*
DOUBLE COMPLEX for *zggesx*.
 Arrays:
 $\mathit{vsl}(\mathit{ldvsl}, *)$, the second dimension of *vsl* must be at least $\max(1, n)$.

If *jobvsl* = 'V', this array will contain the left Schur vectors.

If *jobvsl* = 'N', *vsl* is not referenced.

vsr(*ldvsr*, *), the second dimension of *vsr* must be at least $\max(1, n)$.

If *jobvsr* = 'V', this array will contain the right Schur vectors.

If *jobvsr* = 'N', *vsr* is not referenced.

rconde, *rcondv* REAL for single precision flavors
DOUBLE PRECISION for double precision flavors.
Arrays, DIMENSION (2) each

If *sense* = 'E' or 'B', *rconde*(1) and *rconde*(2) contain the reciprocal condition numbers for the average of the selected eigenvalues.
Not referenced if *sense* = 'N' or 'V'.

If *sense* = 'V' or 'B', *rcondv*(1) and *rcondv*(2) contain the reciprocal condition numbers for the selected deflating subspaces.
Not referenced if *sense* = 'N' or 'E'.

work(1) On exit, if *info* = 0, then *work*(1) returns the required minimal size of *lwork*.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i*, and
i ≤ *n* :
the QZ iteration failed. (A,B) is not in Schur form, but *alphar*(*j*), *alphai*(*j*) (for real flavors), or *alpha*(*j*) (for complex flavors), and *beta*(*j*), *j*=*info*+1,...,*n* should be correct.
i > *n* : errors that usually indicate LAPACK problems:
i = *n*+1: other than QZ iteration failed in ?hgeqz;

$i = n+2$: after reordering, roundoff changed values of some complex eigenvalues so that leading eigenvalues in the generalized Schur form no longer satisfy `selctg = .TRUE..` This could also be caused due to scaling;

$i = n+3$: reordering failed in `?tgsen`.

Application Notes

If you are in doubt how much workspace to supply for the array `work`, use a generous value of `lwork` for the first run. On exit, examine `work(1)` and use this value for subsequent runs.

The quotients `alphan(j)/beta(j)` and `alphai(j)/beta(j)` may easily overflow or underflow, and `beta(j)` may even be zero. Thus, you should avoid simply computing the ratio. However, `alphan` and `alphai` will be always less than and usually comparable with `norm(A)` in magnitude, and `beta` always less than and usually comparable with `norm(B)`.

?ggev

Computes the generalized eigenvalues, and the left and/or right generalized eigenvectors for a pair of nonsymmetric matrices.

```
call sggev ( jobvl, jobvr, n, a, lda, b, ldb, alphas, alphas, beta,
             vl, ldvl, vr, ldvr, work, lwork, info )
call dggev ( jobvl, jobvr, n, a, lda, b, ldb, alphas, alphas, beta,
             vl, ldvl, vr, ldvr, work, lwork, info )
call cggev ( jobvl, jobvr, n, a, lda, b, ldb, alpha, beta,
             vl, ldvl, vr, ldvr, work, lwork, rwork, info )
call zggev ( jobvl, jobvr, n, a, lda, b, ldb, alpha, beta,
             vl, ldvl, vr, ldvr, work, lwork, rwork, info )
```

Discussion

This routine computes for a pair of n -by- n real/complex nonsymmetric matrices (A,B) , the generalized eigenvalues, and optionally, the left and/or right generalized eigenvectors.

A generalized eigenvalue for a pair of matrices (A,B) is a scalar λ or a ratio $alpha / beta = \lambda$, such that $A - \lambda * B$ is singular. It is usually represented as the pair $(alpha, beta)$, as there is a reasonable interpretation for $beta=0$ and even for both being zero.

The right generalized eigenvector $v(j)$ corresponding to the generalized eigenvalue $\lambda(j)$ of (A,B) satisfies

$$A * v(j) = \lambda(j) * B * v(j) .$$

The left generalized eigenvector $u(j)$ corresponding to the generalized eigenvalue $\lambda(j)$ of (A,B) satisfies

$$u(j)^H * A = \lambda(j) * u(j)^H * B$$

where $u(j)^H$ denotes the conjugate transpose of $u(j)$.

Input Parameters

<i>jobvl</i>	<p>CHARACTER*1. Must be 'N' or 'V'.</p> <p>If <i>jobvl</i> = 'N', the left generalized eigenvectors are not computed;</p> <p>If <i>jobvl</i> = 'V', the left generalized eigenvectors are computed.</p>
<i>jobvr</i>	<p>CHARACTER*1. Must be 'N' or 'V'.</p> <p>If <i>jobvr</i> = 'N', the right generalized eigenvectors are not computed;</p> <p>If <i>jobvr</i> = 'V', the right generalized eigenvectors are computed.</p>
<i>n</i>	<p>INTEGER. The order of the matrices <i>A</i>, <i>B</i>, <i>vl</i>, and <i>vr</i> ($n \geq 0$).</p>
<i>a, b, work</i>	<p>REAL for <i>sggev</i></p> <p>DOUBLE PRECISION for <i>dggev</i></p> <p>COMPLEX for <i>cggev</i></p> <p>DOUBLE COMPLEX for <i>zggev</i>.</p> <p>Arrays:</p> <p><i>a(lda,*)</i> is an array containing the <i>n</i>-by-<i>n</i> matrix <i>A</i> (first of the pair of matrices). The second dimension of <i>a</i> must be at least $\max(1, n)$.</p> <p><i>b(ldb,*)</i> is an array containing the <i>n</i>-by-<i>n</i> matrix <i>B</i> (second of the pair of matrices). The second dimension of <i>b</i> must be at least $\max(1, n)$.</p> <p><i>work(lwork)</i> is a workspace array.</p>
<i>lda</i>	<p>INTEGER. The first dimension of the array <i>a</i>. Must be at least $\max(1, n)$.</p>
<i>ldb</i>	<p>INTEGER. The first dimension of the array <i>b</i>. Must be at least $\max(1, n)$.</p>
<i>ldvl, ldvr</i>	<p>INTEGER. The first dimensions of the output matrices <i>vl</i> and <i>vr</i>, respectively. Constraints:</p> <p><i>ldvl</i> ≥ 1. If <i>jobvl</i> = 'V', <i>ldvl</i> $\geq \max(1, n)$.</p> <p><i>ldvr</i> ≥ 1. If <i>jobvr</i> = 'V', <i>ldvr</i> $\geq \max(1, n)$.</p>
<i>lwork</i>	<p>INTEGER. The dimension of the array <i>work</i>.</p>

$lwork \geq \max(1, 8n+16)$ for real flavors;
 $lwork \geq \max(1, 2n)$ for complex flavors.
 For good performance, $lwork$ must generally be larger.

$rwork$ REAL for `cggev`
 DOUBLE PRECISION for `zggev`
 Workspace array, DIMENSION at least $\max(1, 8n)$.
 This array is used in complex flavors only.

Output Parameters

a, b On exit, these arrays have been overwritten.

$alphan, alphai$ REAL for `sggev`;
 DOUBLE PRECISION for `dggev`.
 Arrays, DIMENSION at least $\max(1, n)$ each. Contain values that form generalized eigenvalues in real flavors.
 See $beta$.

$alpha$ COMPLEX for `cggev`;
 DOUBLE COMPLEX for `zggev`.
 Array, DIMENSION at least $\max(1, n)$. Contain values that form generalized eigenvalues in complex flavors.
 See $beta$.

$beta$ REAL for `sggev`
 DOUBLE PRECISION for `dggev`
 COMPLEX for `cggev`
 DOUBLE COMPLEX for `zggev`.
 Array, DIMENSION at least $\max(1, n)$.
 For real flavors:
 On exit, $(alphan(j) + alphai(j)*i)/beta(j)$, $j=1, \dots, n$, will be the generalized eigenvalues.
 If $alphai(j)$ is zero, then the j -th eigenvalue is real; if positive, then the j -th and $(j+1)$ -st eigenvalues are a complex conjugate pair, with $alphai(j+1)$ negative.
 For complex flavors:
 On exit, $alpha(j)/beta(j)$, $j=1, \dots, n$, will be the generalized eigenvalues.
 See also *Application Notes* below.

`v1, vr`REAL for `sggev`DOUBLE PRECISION for `dggev`COMPLEX for `cggev`DOUBLE COMPLEX for `zggev`.

Arrays:

`v1(ldv1, *)`; the second dimension of `v1` must be at least $\max(1, n)$.

If `jobv1 = 'V'`, the left generalized eigenvectors $u(j)$ are stored one after another in the columns of `v1`, in the same order as their eigenvalues. Each eigenvector will be scaled so the largest component have $\text{abs}(\text{Re}) + \text{abs}(\text{Im}) = 1$. If `jobv1 = 'N'`, `v1` is not referenced.

For real flavors:

If the j -th eigenvalue is real, then $u(j) = v1(:,j)$, the j -th column of `v1`. If the j -th and $(j+1)$ -st eigenvalues form a complex conjugate pair, then $u(j) = v1(:,j) + i*v1(:,j+1)$ and $u(j+1) = v1(:,j) - i*v1(:,j+1)$, where $i = \sqrt{-1}$.

For complex flavors: $u(j) = v1(:,j)$, the j -th column of `v1`.`vr(ldvr, *)`; the second dimension of `vr` must be at least $\max(1, n)$.

If `jobv1 = 'V'`, the right generalized eigenvectors $v(j)$ are stored one after another in the columns of `vr`, in the same order as their eigenvalues. Each eigenvector will be scaled so the largest component have $\text{abs}(\text{Re}) + \text{abs}(\text{Im}) = 1$. If `jobv1 = 'N'`, `vr` is not referenced.

For real flavors:

If the j -th eigenvalue is real, then $v(j) = vr(:,j)$, the j -th column of `vr`. If the j -th and $(j+1)$ -st eigenvalues form a complex conjugate pair, then $v(j) = vr(:,j) + i*vr(:,j+1)$ and $v(j+1) = vr(:,j) - i*vr(:,j+1)$.

For complex flavors: $v(j) = vr(:,j)$, the j -th column of `vr`.`work(1)`

On exit, if `info = 0`, then `work(1)` returns the required minimal size of `lwork`.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = *i*, and
 $i \leq n$:
the *QZ* iteration failed. No eigenvectors have been calculated, but *alphar*(*j*), *alphai*(*j*) (for real flavors), or *alpha*(*j*) (for complex flavors), and *beta*(*j*), $j = info + 1, \dots, n$ should be correct.
 $i > n$: errors that usually indicate LAPACK problems:
 $i = n + 1$: other than *QZ* iteration failed in ?hgeqz;
 $i = n + 2$: error return from ?tgevc.

Application Notes

If you are in doubt how much workspace to supply for the array *work*, use a generous value of *lwork* for the first run. On exit, examine *work*(1) and use this value for subsequent runs.

The quotients *alphar*(*j*)/*beta*(*j*) and *alphai*(*j*)/*beta*(*j*) may easily overflow or underflow, and *beta*(*j*) may even be zero. Thus, you should avoid simply computing the ratio. However, *alphar* and *alphai* (for real flavors) or *alpha* (for complex flavors) will be always less than and usually comparable with $\text{norm}(A)$ in magnitude, and *beta* always less than and usually comparable with $\text{norm}(B)$.

?ggev

Computes the generalized eigenvalues, and, optionally, the left and/or right generalized eigenvectors.

```
call sggev ( balanc, jobvl, jobvr, sense, n, a, lda, b, ldb,
             alphas, alphas, beta, vl, ldvl, vr, ldvr, ilo, ihi,
             lscale, rscale, abnrm, bbnrm, rconde, rcondv, work,
             lwork, iwork, bwork, info)
call dggev ( balanc, jobvl, jobvr, sense, n, a, lda, b, ldb,
             alphas, alphas, beta, vl, ldvl, vr, ldvr, ilo, ihi,
             lscale, rscale, abnrm, bbnrm, rconde, rcondv, work,
             lwork, iwork, bwork, info)
call cggev ( balanc, jobvl, jobvr, sense, n, a, lda, b, ldb,
             alpha, beta, vl, ldvl, vr, ldvr, ilo, ihi,
             lscale, rscale, abnrm, bbnrm, rconde, rcondv, work,
             lwork, rwork, iwork, bwork, info)
call zggev ( balanc, jobvl, jobvr, sense, n, a, lda, b, ldb,
             alpha, beta, vl, ldvl, vr, ldvr, ilo, ihi,
             lscale, rscale, abnrm, bbnrm, rconde, rcondv, work,
             lwork, rwork, iwork, bwork, info)
```

Discussion

This routine computes for a pair of n -by- n real/complex nonsymmetric matrices (A,B) , the generalized eigenvalues, and optionally, the left and/or right generalized eigenvectors.

Optionally also, it computes a balancing transformation to improve the conditioning of the eigenvalues and eigenvectors (*ilo*, *ihi*, *lscale*, *rscale*, *abnrm*, and *bbnrm*), reciprocal condition numbers for the eigenvalues (*rconde*), and reciprocal condition numbers for the right eigenvectors (*rcondv*).

A generalized eigenvalue for a pair of matrices (A,B) is a scalar λ or a ratio $alpha / beta = \lambda$, such that $A - \lambda * B$ is singular. It is usually represented as the pair $(alpha, beta)$, as there is a reasonable interpretation for $beta=0$ and

even for both being zero.

The right generalized eigenvector $v(j)$ corresponding to the generalized eigenvalue $\lambda(j)$ of (A,B) satisfies

$$A * v(j) = \lambda(j) * B * v(j).$$

The left generalized eigenvector $u(j)$ corresponding to the generalized eigenvalue $\lambda(j)$ of (A,B) satisfies

$$u(j)^H * A = \lambda(j) * u(j)^H * B$$

where $u(j)^H$ denotes the conjugate transpose of $u(j)$.

Input Parameters

- balanc* CHARACTER*1. Must be 'N', 'P', 'S', or 'B'. Specifies the balance option to be performed.
- If *balanc* = 'N', do not diagonally scale or permute;
 If *balanc* = 'P', permute only;
 If *balanc* = 'S', scale only;
 If *balanc* = 'B', both permute and scale.
- Computed reciprocal condition numbers will be for the matrices after balancing and/or permuting. Permuting does not change condition numbers (in exact arithmetic), but balancing does.
- jobvl* CHARACTER*1. Must be 'N' or 'V'.
 If *jobvl* = 'N', the left generalized eigenvectors are not computed;
 If *jobvl* = 'V', the left generalized eigenvectors are computed.
- jobvr* CHARACTER*1. Must be 'N' or 'V'.
 If *jobvr* = 'N', the right generalized eigenvectors are not computed;
 If *jobvr* = 'V', the right generalized eigenvectors are computed.
- sense* CHARACTER*1. Must be 'N', 'E', 'V', or 'B'.
 Determines which reciprocal condition number are computed.

If *sense* = 'N', none are computed;
 If *sense* = 'E', computed for eigenvalues only;
 If *sense* = 'V', computed for eigenvectors only;
 If *sense* = 'B', computed for eigenvalues and eigenvectors.

n **INTEGER**. The order of the matrices *A*, *B*, *v1*, and *vr* ($n \geq 0$).

a, *b*, *work* **REAL** for *sggevx*
 DOUBLE PRECISION for *dggevx*
 COMPLEX for *cggevx*
 DOUBLE COMPLEX for *zggevx*.

Arrays:
a(*lda*,*) is an array containing the *n*-by-*n* matrix *A* (first of the pair of matrices).
 The second dimension of *a* must be at least $\max(1, n)$.
b(*ldb*,*) is an array containing the *n*-by-*n* matrix *B* (second of the pair of matrices).
 The second dimension of *b* must be at least $\max(1, n)$.
work(*lwork*) is a workspace array.

lda **INTEGER**. The first dimension of the array *a*.
 Must be at least $\max(1, n)$.

ldb **INTEGER**. The first dimension of the array *b*.
 Must be at least $\max(1, n)$.

ldv1, *ldvr* **INTEGER**. The first dimensions of the output matrices *v1* and *vr*, respectively. Constraints:
ldv1 ≥ 1 . If *jobv1* = 'V', *ldv1* $\geq \max(1, n)$.
ldvr ≥ 1 . If *jobvr* = 'V', *ldvr* $\geq \max(1, n)$.

lwork **INTEGER**. The dimension of the array *work*.
 For real flavors:
lwork $\geq \max(1, 6n)$;
 if *sense* = 'E', *lwork* $\geq 12n$;
 if *sense* = 'V', or 'B', *lwork* $\geq 2n^2 + 12n + 16$.
 For complex flavors:

$lwork \geq \max(1, 2n)$;
 if $sense = 'N'$, or $'E'$, $lwork \geq 2n$;
 if $sense = 'V'$, or $'B'$, $lwork \geq 2n^2 + 2n$.

rwork REAL for *cggev*
 DOUBLE PRECISION for *zggev*
 Workspace array, DIMENSION at least $\max(1, 6n)$.
 This array is used in complex flavors only.

iwork INTEGER.
 Workspace array, DIMENSION at least $(n+6)$ for real
 flavors and at least $(n+2)$ for complex flavors.
 Not referenced if $sense = 'E'$.

bwork LOGICAL.
 Workspace array, DIMENSION at least $\max(1, n)$.
 Not referenced if $sense = 'N'$.

Output Parameters

a, *b* On exit, these arrays have been overwritten.
 If $jobvl = 'V'$ or $jobvr = 'V'$ or both, then *a* contains
 the first part of the real Schur form of the "balanced"
 versions of the input *A* and *B*, and *b* contains its second
 part.

alphan, *alpha* REAL for *sggev*;
 DOUBLE PRECISION for *dggev*.
 Arrays, DIMENSION at least $\max(1, n)$ each. Contain
 values that form generalized eigenvalues in real flavors.
 See *beta*.

alpha COMPLEX for *cggev*;
 DOUBLE COMPLEX for *zggev*.
 Array, DIMENSION at least $\max(1, n)$. Contain values
 that form generalized eigenvalues in complex flavors.
 See *beta*.

beta REAL for *sggev*
 DOUBLE PRECISION for *dggev*
 COMPLEX for *cggev*
 DOUBLE COMPLEX for *zggev*.

Array, **DIMENSION** at least $\max(1, n)$.

For real flavors:

On exit, $(\mathit{alphar}(j) + \mathit{alphai}(j)*i)/\mathit{beta}(j)$, $j=1, \dots, n$, will be the generalized eigenvalues.

If $\mathit{alphai}(j)$ is zero, then the j -th eigenvalue is real; if positive, then the j -th and $(j+1)$ -st eigenvalues are a complex conjugate pair, with $\mathit{alphai}(j+1)$ negative.

For complex flavors:

On exit, $\mathit{alpha}(j)/\mathit{beta}(j)$, $j=1, \dots, n$, will be the generalized eigenvalues.

See also *Application Notes* below.

v1, *vr*

REAL for **sggev**

DOUBLE PRECISION for **dggev**

COMPLEX for **cggev**

DOUBLE COMPLEX for **zggev**.

Arrays:

$\mathit{v1}(\mathit{ldv1}, *)$; the second dimension of *v1* must be at least $\max(1, n)$.

If $\mathit{jobv1} = 'V'$, the left generalized eigenvectors $u(j)$ are stored one after another in the columns of *v1*, in the same order as their eigenvalues. Each eigenvector will be scaled so the largest component have $\text{abs}(\text{Re}) + \text{abs}(\text{Im}) = 1$. If $\mathit{jobv1} = 'N'$, *v1* is not referenced.

For real flavors:

If the j -th eigenvalue is real, then $u(j) = \mathit{v1}(:, j)$, the j -th column of *v1*. If the j -th and $(j+1)$ -st eigenvalues form a complex conjugate pair, then $u(j) = \mathit{v1}(:, j) + i*\mathit{v1}(:, j+1)$ and $u(j+1) = \mathit{v1}(:, j) - i*\mathit{v1}(:, j+1)$, where $i = \sqrt{-1}$.

For complex flavors:

$u(j) = \mathit{v1}(:, j)$, the j -th column of *v1*.

$\mathit{vr}(\mathit{ldvr}, *)$; the second dimension of *vr* must be at least $\max(1, n)$.

If $\mathit{jobvr} = 'V'$, the right generalized eigenvectors $v(j)$ are stored one after another in the columns of *vr*, in the same order as their eigenvalues. Each eigenvector will

be scaled so the largest component have $\text{abs}(\text{Re}) + \text{abs}(\text{Im}) = 1$. If *jobvr* = 'N', *vr* is not referenced.

For real flavors:

If the *j*-th eigenvalue is real, then $v(j) = \text{vr}(:,j)$, the *j*-th column of *vr*. If the *j*-th and (*j*+1)-st eigenvalues form a complex conjugate pair, then $v(j) = \text{vr}(:,j) + i*\text{vr}(:,j+1)$ and $v(j+1) = \text{vr}(:,j) - i*\text{vr}(:,j+1)$.

For complex flavors:

$v(j) = \text{vr}(:,j)$, the *j*-th column of *vr*.

ilo, ihi

INTEGER.

ilo and *ihi* are integer values such that on exit $A(i,j) = 0$ and $B(i,j) = 0$ if $i > j$ and $j = 1, \dots, ilo-1$ or $i = ihi+1, \dots, n$.

If *balanc* = 'N' or 'S', *ilo* = 1 and *ihi* = *n*.

lscale, rscale

REAL for single-precision flavors

DOUBLE PRECISION for double-precision flavors.

Arrays, DIMENSION at least $\max(1, n)$ each.

lscale contains details of the permutations and scaling factors applied to the left side of *A* and *B*.

If *PL(j)* is the index of the row interchanged with row *j*, and *DL(j)* is the scaling factor applied to row *j*, then

$$\begin{aligned} \text{lscale}(j) &= PL(j), & \text{for } j = 1, \dots, ilo-1 \\ &= DL(j), & \text{for } j = ilo, \dots, ihi \\ &= PL(j) & \text{for } j = ihi+1, \dots, n. \end{aligned}$$

The order in which the interchanges are made is *n* to *ihi*+1, then 1 to *ilo*-1.

rscale contains details of the permutations and scaling factors applied to the right side of *A* and *B*.

If *PR(j)* is the index of the column interchanged with column *j*, and *DR(j)* is the scaling factor applied to column *j*, then

$$\begin{aligned} \text{rscale}(j) &= PR(j), & \text{for } j = 1, \dots, ilo-1 \\ &= DR(j), & \text{for } j = ilo, \dots, ihi \\ &= PR(j) & \text{for } j = ihi+1, \dots, n. \end{aligned}$$

The order in which the interchanges are made is n to $i_{hi}+1$, then 1 to $i_{lo}-1$.

abnrm, bbnrm REAL for single-precision flavors
DOUBLE PRECISION for double-precision flavors.

The one-norms of the balanced matrices A and B , respectively.

rconde, rcondv REAL for single precision flavors
DOUBLE PRECISION for double precision flavors.
Arrays, DIMENSION at least $\max(1, n)$ each.

If *sense* = 'E', or 'B', *rconde* contains the reciprocal condition numbers of the selected eigenvalues, stored in consecutive elements of the array. For a complex conjugate pair of eigenvalues two consecutive elements of *rconde* are set to the same value. Thus *rconde*(j), *rcondv*(j), and the j -th columns of *vl* and *vr* all correspond to the same eigenpair (but not in general the j -th eigenpair, unless all eigenpairs are selected).
If *sense* = 'V', *rconde* is not referenced.

If *sense* = 'V', or 'B', *rcondv* contains the estimated reciprocal condition numbers of the selected eigenvectors, stored in consecutive elements of the array. For a complex eigenvector two consecutive elements of *rcondv* are set to the same value. If the eigenvalues cannot be reordered to compute *rcondv*(j), *rcondv*(j) is set to 0; this can only occur when the true value would be very small anyway.

If *sense* = 'E', *rcondv* is not referenced.

work(1) On exit, if *info* = 0, then *work(1)* returns the required minimal size of *lwork*.

info INTEGER.

If *info* = 0, the execution is successful.

If *info* = $-i$, the i th parameter had an illegal value.

If *info* = i , and

$i \leq n$:

the *QZ* iteration failed. No eigenvectors have been calculated, but *alphar*(*j*), *alphai*(*j*) (for real flavors), or *alpha*(*j*) (for complex flavors), and *beta*(*j*), $j = \text{info} + 1, \dots, n$ should be correct.

i > *n* : errors that usually indicate LAPACK problems:

i = *n*+1: other than *QZ* iteration failed in ?hgeqz;

i = *n*+2: error return from ?tgevc.

Application Notes

If you are in doubt how much workspace to supply for the array *work*, use a generous value of *lwork* for the first run. On exit, examine *work*(1) and use this value for subsequent runs.

The quotients *alphar*(*j*)/*beta*(*j*) and *alphai*(*j*)/*beta*(*j*) may easily overflow or underflow, and *beta*(*j*) may even be zero. Thus, you should avoid simply computing the ratio. However, *alphar* and *alphai* (for real flavors) or *alpha* (for complex flavors) will be always less than and usually comparable with $\text{norm}(A)$ in magnitude, and *beta* always less than and usually comparable with $\text{norm}(B)$.

References

- [LUG] E. Anderson, Z. Bai et al. *LAPACK User's Guide*. Third edition, SIAM, Philadelphia, 1999.
- [Golub96] G. Golub, C. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, third edition, 1996.

LAPACK Auxiliary Routines

6

This chapter describes the Intel[®] Math Kernel Library implementation of LAPACK auxiliary routines. The library includes auxiliary routines for both real and complex data.

Routine naming conventions, mathematical notation, and matrix storage schemes used for LAPACK auxiliary routines are the same as for the driver and computational routines described in previous chapters.

Routines are listed in alphabetical order of the routine or function group name (which always begins with `-?`).

?lacgv

Conjugates a complex vector.

```
call clacgv (n, x, incx)
call zlacgv (n, x, incx)
```

Discussion

This routine conjugates a complex vector `x` of length `n` and increment `incx` (see [Vector Arguments in BLAS](#) in Appendix A).

Input Parameters

<code>n</code>	<code>INTEGER</code> . The length of the vector <code>x</code> ($n \geq 0$).
<code>x</code>	<code>COMPLEX</code> for <code>clacgv</code> <code>COMPLEX*16</code> for <code>zlacgv</code> . Array, dimension $(1+(n-1) * incx)$. Contains the vector of length <code>n</code> to be conjugated.

incx **INTEGER**. The spacing between successive elements of *x*.

Output Parameters

x On exit, overwritten with `conjg(x)`.

?lacrm

Multiplies a complex matrix by a square real matrix.

```
call clacrm (m, n, a, lda, b, ldb, c, ldc, rwork)
call zlacrm (m, n, a, lda, b, ldb, c, ldc, rwork)
```

Discussion

This routine performs a simple matrix-matrix multiplication of the form

$$C = A * B,$$

where *A* is *m*-by-*n* and complex, *B* is *n*-by-*n* and real, *C* is *m*-by-*n* and complex.

Input Parameters

m **INTEGER**. The number of rows of the matrix *A* and of the matrix *C* (*m* ≥ 0).

n **INTEGER**. The number of columns and rows of the matrix *B* and the number of columns of the matrix *C* (*n* ≥ 0).

a **COMPLEX** for `clacrm`
COMPLEX*16 for `zlacrm`
 Array, **DIMENSION** (*lda*, *n*). Contains the *m*-by-*n* matrix *A*.

lda **INTEGER**. The leading dimension of the array *a*, *lda* ≥ max(1, *m*).

b REAL for `clacrm`
 DOUBLE PRECISION for `zlacrm`
 Array, DIMENSION (*ldb*, *n*). Contains the *n*-by-*n* matrix *B*.

ldb INTEGER. The leading dimension of the array *b*,
ldb ≥ max(1, *n*).

ldc INTEGER. The leading dimension of the output array *c*,
ldc ≥ max(1, *n*).

rwork REAL for `clacrm`
 DOUBLE PRECISION for `zlacrm`
 Workspace array, DIMENSION (2**m***n*).

Output Parameters

c COMPLEX for `clacrm`
 COMPLEX*16 for `zlacrm`
 Array, DIMENSION (*ldc*, *n*). Contains the *m*-by-*n* matrix *C*.

?lacrt

Performs a linear transformation of a pair of complex vectors.

```
call clacrt (n, cx, incx, cy, incy, c, s)
call zlacrt (n, cx, incx, cy, incy, c, s)
```

Discussion

This routine performs the following transformation

$$\begin{pmatrix} c & s \\ -s & c \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \Rightarrow \begin{pmatrix} x \\ y \end{pmatrix},$$

where *c*, *s* are complex scalars and *x*, *y* are complex vectors.

Input Parameters

<i>n</i>	INTEGER. The number of elements in the vectors <i>cx</i> and <i>cy</i> ($n \geq 0$).
<i>cx, cy</i>	COMPLEX for <code>clacrt</code> COMPLEX*16 for <code>zlacrt</code> Arrays, dimension (<i>n</i>). Contain input vectors <i>x</i> and <i>y</i> , respectively.
<i>incx</i>	INTEGER. The increment between successive elements of <i>cx</i> .
<i>incy</i>	INTEGER. The increment between successive elements of <i>cy</i> .
<i>c, s</i>	COMPLEX for <code>clacrt</code> COMPLEX*16 for <code>zlacrt</code> Complex scalars that define the transform matrix

$$\begin{bmatrix} c & s \\ -s & c \end{bmatrix}$$

Output Parameters

<i>cx</i>	On exit, overwritten with $c*x + s*y$.
<i>cy</i>	On exit, overwritten with $-s*x + c*y$.

?laesy

Computes the eigenvalues and eigenvectors of a 2-by-2 complex symmetric matrix, and checks that the norm of the matrix of eigenvectors is larger than a threshold value.

```
call claesy (a, b, c, rt1, rt2, evscal, cs1, sn1)
call zlaesy (a, b, c, rt1, rt2, evscal, cs1, sn1)
```

Discussion

This routine performs the eigendecomposition of a 2-by-2 symmetric matrix

$$\begin{bmatrix} a & b \\ b & c \end{bmatrix},$$

provided the norm of the matrix of eigenvectors is larger than some threshold value.

rt1 is the eigenvalue of larger absolute value, and *rt2* of smaller absolute value. If the eigenvectors are computed, then on return (*cs1*, *sn1*) is the unit eigenvector for *rt1*, hence

$$\begin{bmatrix} cs1 & sn1 \\ -sn1 & cs1 \end{bmatrix} \cdot \begin{bmatrix} a & b \\ b & c \end{bmatrix} \cdot \begin{bmatrix} cs1 & -sn1 \\ sn1 & cs1 \end{bmatrix} = \begin{bmatrix} rt1 & 0 \\ 0 & rt2 \end{bmatrix}$$

Input Parameters

a, *b*, *c* **COMPLEX** for **claesy**
 COMPLEX*16 for **zlaesy**
 Elements of the input matrix.

Output Parameters

rt1, *rt2* **COMPLEX** for **claesy**
 COMPLEX*16 for **zlaesy**
 Eigenvalues of larger and smaller modulus, respectively.

evscal **COMPLEX** for **claesy**
 COMPLEX*16 for **zlaesy**
 The complex value by which the eigenvector matrix was scaled to make it orthonormal. If *evscal* is zero, the eigenvectors were not computed. This means one of two things: the 2-by-2 matrix could not be diagonalized, or the norm of the matrix of eigenvectors before scaling was larger than the threshold value **thresh** (set to 0.1E0).

cs1, sn1 COMPLEX for *claesy*
 COMPLEX*16 for *zlaesy*

If *evscal* is not zero, then (*cs1, sn1*) is the unit right eigenvector for *rt1*.

?rot

Applies a plane rotation with real cosine and complex sine to a pair of complex vectors.

```
call crot (n, cx, incx, cy, incy, c, s)
call zrot (n, cx, incx, cy, incy, c, s)
```

Discussion

This routine applies a plane rotation, where the cosine (*c*) is real and the sine (*s*) is complex, and the vectors *cx* and *cy* are complex. This routine has its real equivalents in BLAS (see [?rot](#) in Chapter 2).

Input Parameters

<i>n</i>	INTEGER. The number of elements in the vectors <i>cx</i> and <i>cy</i> .
<i>cx, cy</i>	COMPLEX for <i>crot</i> COMPLEX*16 for <i>zrot</i> Arrays of dimension (<i>n</i>), contain input vectors <i>x</i> and <i>y</i> , respectively.
<i>incx</i>	INTEGER. The increment between successive elements of <i>cx</i> .
<i>incy</i>	INTEGER. The increment between successive elements of <i>cy</i> .
<i>c</i>	REAL for <i>crot</i> DOUBLE PRECISION for <i>zrot</i>

s **COMPLEX** for **crot**
COMPLEX*16 for **zrot**
 Values that define a rotation

$$\begin{bmatrix} c & s \\ -\text{conjg}(s) & c \end{bmatrix}$$

where $c*c + s*\text{conjg}(s) = 1.0$.

Output Parameters

cx On exit, overwritten with $c*x + s*y$.
cy On exit, overwritten with $-\text{conjg}(s)*x + c*y$.

?spmv

Computes a matrix-vector product for complex vectors using a complex symmetric packed matrix.

```
call cspmv ( uplo, n, alpha, ap, x, incx, beta, y, incy )
call zspmv ( uplo, n, alpha, ap, x, incx, beta, y, incy )
```

Discussion

These routines perform a matrix-vector operation defined as

$y := \text{alpha}*a*x + \text{beta}*y,$

where:

alpha and *beta* are complex scalars,

x and *y* are *n*-element complex vectors

a is an *n*-by-*n* complex symmetric matrix, supplied in packed form.

These routines have their real equivalents in BLAS (see [?spmv](#) in Chapter 2).

Input Parameters

<i>uplo</i>	<p>CHARACTER*1. Specifies whether the upper or lower triangular part of the matrix <i>a</i> is supplied in the packed array <i>ap</i>, as follows:</p> <p>If <i>uplo</i> = 'U' or 'u', the upper triangular part of the matrix <i>a</i> is supplied in the array <i>ap</i>.</p> <p>If <i>uplo</i> = 'L' or 'l', the lower triangular part of the matrix <i>a</i> is supplied in the array <i>ap</i>.</p>
<i>n</i>	<p>INTEGER. Specifies the order of the matrix <i>a</i>. The value of <i>n</i> must be at least zero.</p>
<i>alpha, beta</i>	<p>COMPLEX for <i>cspmv</i> COMPLEX*16 for <i>zspmv</i></p> <p>Specify complex scalars <i>alpha</i> and <i>beta</i>. When <i>beta</i> is supplied as zero, then <i>y</i> need not be set on input.</p>
<i>ap</i>	<p>COMPLEX for <i>cspmv</i> COMPLEX*16 for <i>zspmv</i></p> <p>Array, DIMENSION at least $((n*(n+1))/2)$. Before entry, with <i>uplo</i> = 'U' or 'u', the array <i>ap</i> must contain the upper triangular part of the symmetric matrix packed sequentially, column-by-column, so that <i>ap</i>(1) contains <i>a</i>(1, 1), <i>ap</i>(2) and <i>ap</i>(3) contain <i>a</i>(1, 2) and <i>a</i>(2, 2) respectively, and so on. Before entry, with <i>uplo</i> = 'L' or 'l', the array <i>ap</i> must contain the lower triangular part of the symmetric matrix packed sequentially, column-by-column, so that <i>ap</i>(1) contains <i>a</i>(1, 1), <i>ap</i>(2) and <i>ap</i>(3) contain <i>a</i>(2, 1) and <i>a</i>(3, 1) respectively, and so on.</p>
<i>x</i>	<p>COMPLEX for <i>cspmv</i> COMPLEX*16 for <i>zspmv</i></p> <p>Array, DIMENSION at least $(1 + (n - 1)*abs(incx))$. Before entry, the incremented array <i>x</i> must contain the <i>n</i>-element vector <i>x</i>.</p>
<i>incx</i>	<p>INTEGER. Specifies the increment for the elements of <i>x</i>. The value of <i>incx</i> must not be zero.</p>

y **COMPLEX** for `cspmv`
COMPLEX*16 for `zspmv`
 Array, **DIMENSION** at least $(1 + (n - 1) * \text{abs}(incy))$.
 Before entry, the incremented array *y* must contain the
n-element vector *y*.

incy **INTEGER**. Specifies the increment for the elements of *y*.
 The value of *incy* must not be zero.

Output Parameters

y Overwritten by the updated vector *y*.

?spr

*Performs the symmetrical rank-1 update
of a complex symmetric packed matrix.*

```
call cspr( uplo, n, alpha, x, incx, ap )
call zspr( uplo, n, alpha, x, incx, ap )
```

Discussion

The `?spr` routines perform a matrix-vector operation defined as

$$a := \text{alpha} * x * \text{conjg}(x') + a,$$

where:

alpha is a complex scalar

x is an *n*-element complex vector

a is an *n*-by-*n* complex symmetric matrix, supplied in packed form.

These routines have their real equivalents in BLAS (see `?spr` in Chapter 2).

Input Parameters

uplo **CHARACTER*1**. Specifies whether the upper or lower
triangular part of the matrix *a* is supplied in the packed
array *ap*, as follows:

If `uplo = 'U'` or `'u'`, the upper triangular part of the matrix `a` is supplied in the array `ap`.

If `uplo = 'L'` or `'l'`, the lower triangular part of the matrix `a` is supplied in the array `ap`.

`n` **INTEGER**. Specifies the order of the matrix `a`. The value of `n` must be at least zero.

`alpha` **COMPLEX** for `cspr`
COMPLEX*16 for `zspr`

Specifies the scalar `alpha`.

`x` **COMPLEX** for `cspr`
COMPLEX*16 for `zspr`

Array, **DIMENSION** at least $(1 + (n - 1) * \text{abs}(\text{incx}))$. Before entry, the incremented array `x` must contain the `n`-element vector `x`.

`incx` **INTEGER**. Specifies the increment for the elements of `x`. The value of `incx` must not be zero.

`ap` **COMPLEX** for `cspr`
COMPLEX*16 for `zspr`

Array, **DIMENSION** at least $((n * (n + 1)) / 2)$. Before entry, with `uplo = 'U'` or `'u'`, the array `ap` must contain the upper triangular part of the symmetric matrix packed sequentially, column-by-column, so that `ap(1)` contains `a(1,1)`, `ap(2)` and `ap(3)` contain `a(1,2)` and `a(2,2)` respectively, and so on.

Before entry, with `uplo = 'L'` or `'l'`, the array `ap` must contain the lower triangular part of the symmetric matrix packed sequentially, column-by-column, so that `ap(1)` contains `a(1,1)`, `ap(2)` and `ap(3)` contain `a(2,1)` and `a(3,1)` respectively, and so on.

Note that the imaginary parts of the diagonal elements need not be set, they are assumed to be zero, and on exit they are set to zero.

Output Parameters

ap With *uplo* = 'U' or 'u', overwritten by the upper triangular part of the updated matrix.
With *uplo* = 'L' or 'l', overwritten by the lower triangular part of the updated matrix.

?symv

Computes a matrix-vector product for a complex symmetric matrix.

```
call csymv ( uplo, n, alpha, a, lda, x, incx, beta, y, incy )
call zsymv ( uplo, n, alpha, a, lda, x, incx, beta, y, incy )
```

Discussion

These routines perform the matrix-vector operation defined as

$$y := \alpha * a * x + \beta * y,$$

where:

alpha and *beta* are complex scalars

x and *y* are *n*-element complex vectors

a is an *n*-by-*n* symmetric complex matrix.

These routines have their real equivalents in BLAS (see [?symv](#) in Chapter 2).

Input Parameters

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the array *a* is to be referenced, as follows:

	<p>If <code>uplo = 'U'</code> or <code>'u'</code>, the upper triangular part of the array <code>a</code> is to be referenced .</p> <p>If <code>uplo = 'L'</code> or <code>'l'</code>, the lower triangular part of the array <code>a</code> is to be referenced.</p>
<code>n</code>	<p>INTEGER. Specifies the order of the matrix <code>a</code>. The value of <code>n</code> must be at least zero.</p>
<code>alpha, beta</code>	<p>COMPLEX for <code>csymv</code> COMPLEX*16 for <code>zsymv</code></p> <p>Specify the scalars <code>alpha</code> and <code>beta</code>. When <code>beta</code> is supplied as zero, then <code>y</code> need not be set on input.</p>
<code>a</code>	<p>COMPLEX for <code>csymv</code> COMPLEX*16 for <code>zsymv</code></p> <p>Array, DIMENSION (<code>lda, n</code>). Before entry with <code>uplo = 'U'</code> or <code>'u'</code>, the leading <code>n</code>-by-<code>n</code> upper triangular part of the array <code>a</code> must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of <code>a</code> is not referenced. Before entry with <code>uplo = 'L'</code> or <code>'l'</code>, the leading <code>n</code>-by-<code>n</code> lower triangular part of the array <code>a</code> must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of <code>a</code> is not referenced.</p>
<code>lda</code>	<p>INTEGER. Specifies the first dimension of <code>a</code> as declared in the calling (sub)program. The value of <code>lda</code> must be at least <code>max(1, n)</code>.</p>
<code>x</code>	<p>COMPLEX for <code>csymv</code> COMPLEX*16 for <code>zsymv</code></p> <p>Array, DIMENSION at least <code>(1 + (n - 1) * abs(incx))</code>. Before entry, the incremented array <code>x</code> must contain the <code>n</code>-element vector <code>x</code>.</p>
<code>incx</code>	<p>INTEGER. Specifies the increment for the elements of <code>x</code>. The value of <code>incx</code> must not be zero.</p>
<code>y</code>	<p>COMPLEX for <code>csymv</code> COMPLEX*16 for <code>zsymv</code></p>

Array, `DIMENSION` at least $(1 + (n - 1) * \text{abs}(incy))$.
Before entry, the incremented array `y` must contain the
`n`-element vector `y`.

`incy` **INTEGER**. Specifies the increment for the elements of `y`.
The value of `incy` must not be zero.

Output Parameters

`y` Overwritten by the updated vector `y`.

?syr

*Performs the symmetric rank-1 update
of a complex symmetric matrix.*

```
call csyr( uplo, n, alpha, x, incx, a, lda )
call zsyr( uplo, n, alpha, x, incx, a, lda )
```

Discussion

These routines perform the symmetric rank 1 operation defined as

$a := \text{alpha} * x * x' + a,$

where:

`alpha` is a complex scalar

`x` is an `n`-element complex vector

`a` is an `n`-by-`n` complex symmetric matrix.

These routines have their real equivalents in BLAS (see [?syr](#) in Chapter 2).

Input Parameters

`uplo` **CHARACTER*1**. Specifies whether the upper or lower
triangular part of the array `a` is to be referenced, as
follows:

	<p>If <code>uplo = 'U'</code> or <code>'u'</code>, the upper triangular part of the array <code>a</code> is to be referenced .</p> <p>If <code>uplo = 'L'</code> or <code>'l'</code>, the lower triangular part of the array <code>a</code> is to be referenced.</p>
<code>n</code>	<p>INTEGER. Specifies the order of the matrix <code>a</code>. The value of <code>n</code> must be at least zero.</p>
<code>alpha</code>	<p>COMPLEX for <code>csyr</code> COMPLEX*16 for <code>zsyr</code></p> <p>Specifies the scalar <code>alpha</code>.</p>
<code>x</code>	<p>COMPLEX for <code>csyr</code> COMPLEX*16 for <code>zsyr</code></p> <p>Array, DIMENSION at least $(1 + (n - 1) * \text{abs}(\text{incx}))$. Before entry, the incremented array <code>x</code> must contain the <code>n</code>-element vector <code>x</code>.</p>
<code>incx</code>	<p>INTEGER. Specifies the increment for the elements of <code>x</code>. The value of <code>incx</code> must not be zero.</p>
<code>a</code>	<p>COMPLEX for <code>csyr</code> COMPLEX*16 for <code>zsyr</code></p> <p>Array, DIMENSION (<code>lda, n</code>). Before entry with <code>uplo = 'U'</code> or <code>'u'</code>, the leading <code>n</code>-by-<code>n</code> upper triangular part of the array <code>a</code> must contain the upper triangular part of the symmetric matrix and the strictly lower triangular part of <code>a</code> is not referenced.</p> <p>Before entry with <code>uplo = 'L'</code> or <code>'l'</code>, the leading <code>n</code>-by-<code>n</code> lower triangular part of the array <code>a</code> must contain the lower triangular part of the symmetric matrix and the strictly upper triangular part of <code>a</code> is not referenced.</p>
<code>lda</code>	<p>INTEGER. Specifies the first dimension of <code>a</code> as declared in the calling (sub)program. The value of <code>lda</code> must be at least $\max(1, n)$.</p>

Output Parameters

a With *uplo* = 'U' or 'u', the upper triangular part of the array *a* is overwritten by the upper triangular part of the updated matrix.

With *uplo* = 'L' or 'l', the lower triangular part of the array *a* is overwritten by the lower triangular part of the updated matrix.

i?max1

Finds the index of the vector element whose real part has maximum absolute value.

```
index = icmax1 ( n, cx, incx )  
index = izmax1 ( n, cx, incx )
```

Discussion

Given a complex vector *cx*, the *i?max1* functions return the index of the vector element whose real part has maximum absolute value. These functions are based on the BLAS functions *icamax/izamax*, but using the absolute value of the real part. They are designed for use with *clacon/zlacon*.

Input Parameters

n **INTEGER**. Specifies the number of elements in the vector *cx*.

cx **COMPLEX** for *icmax1*
COMPLEX*16 for *izmax1*
Array, **DIMENSION** at least $(1+(n-1)*abs(incx))$.
Contains the input vector.

incx **INTEGER**. Specifies the spacing between successive elements of *cx*.

Output Parameters

index **INTEGER**. Contains the index of the vector element whose real part has maximum absolute value.

ilaenv

Environmental enquiry function which returns values for tuning algorithmic performance.

```
value = ilaenv ( ispec, name, opts, n1, n2, n3, n4 )
```

Discussion

Enquiry function *ilaenv* is called from the LAPACK routines to choose problem-dependent parameters for the local environment. See *ispec* for a description of the parameters.

This version provides a set of parameters which should give good, but not optimal, performance on many of the currently available computers. Users are encouraged to modify this subroutine to set the tuning parameters for their particular machine using the option and problem size information in the arguments.

This routine will not function correctly if it is converted to all lower case. Converting it to all upper case is allowed.

Input Parameters

ispec **INTEGER**. Specifies the parameter to be returned as the value of *ilaenv*:
= 1: the optimal blocksize; if this value is 1, an unblocked algorithm will give the best performance.

- = 2: the minimum block size for which the block routine should be used; if the usable block size is less than this value, an unblocked routine should be used.
 - = 3: the crossover point (in a block routine, for N less than this value, an unblocked routine should be used)
 - = 4: the number of shifts, used in the nonsymmetric eigenvalue routines
 - = 5: the minimum column dimension for blocking to be used; rectangular blocks must have dimension at least k by m , where k is given by `ilaenv(2,...)` and m by `ilaenv(5,...)`
 - = 6: the crossover point for the SVD (when reducing an m by n matrix to bidiagonal form, if $\max(m,n)/\min(m,n)$ exceeds this value, a QR factorization is used first to reduce the matrix to a triangular form.)
 - = 7: the number of processors
 - = 8: the crossover point for the multishift QR and QZ methods for nonsymmetric eigenvalue problems.
 - = 9: maximum size of the subproblems at the bottom of the computation tree in the divide-and-conquer algorithm (used by `?gelsd` and `?gesdd`)
 - = 10: IEEE NaN arithmetic can be trusted not to trap
 - = 11: infinity arithmetic can be trusted not to trap
- `name` CHARACTER*(*). The name of the calling subroutine, in either upper case or lower case.
- `opts` CHARACTER*(*). The character options to the subroutine `name`, concatenated into a single character string. For example, `uplo = 'U'`, `trans = 'T'`, and `diag = 'N'` for a triangular routine would be specified as `opts = 'UTN'`.
- `n1, n2, n3, n4` INTEGER. Problem dimensions for the subroutine `name`; these may not all be required.

Output Parameters

value `INTEGER`.
If *value* ≥ 0 : the value of the parameters specified by *ispec*;
If *value* = $-k < 0$: the *k*-th argument had an illegal value.

Application Notes

The following conventions have been used when calling `ilaenv` from the LAPACK routines:

1) *opts* is a concatenation of all of the character options to subroutine *name*, in the same order that they appear in the argument list for *name*, even if they are not used in determining the value of the parameter specified by *ispec*.

2) The problem dimensions *n1*, *n2*, *n3*, *n4* are specified in the order that they appear in the argument list for *name*. *n1* is used first, *n2* second, and so on, and unused problem dimensions are passed a value of -1.

3) The parameter value returned by `ilaenv` is checked for validity in the calling subroutine. For example, `ilaenv` is used to retrieve the optimal blocksize for `strtri` as follows:

```
nb = ilaenv( 1, 'strtri', uplo // diag, n, -1, -1, -1 )  
if( nb.le.1 ) nb = max( 1, n )
```

lsame

*Tests two characters for equality
regardless of case.*

```
val = lsame ( ca, cb )
```

Discussion

This logical function returns `.TRUE.` if *ca* is the same letter as *cb* regardless of case.

Input Parameters

ca, cb **CHARACTER*1**. Specify the single characters to be compared.

Output Parameters

val **LOGICAL**. Result of the comparison.

Isamen

Tests two character strings for equality regardless of case.

```
val = lsamen ( n, ca, cb )
```

Discussion

This logical function tests if the first *n* letters of the string *ca* are the same as the first *n* letters of *cb*, regardless of case. The function `lsamen` returns `.TRUE.` if *ca* and *cb* are equivalent except for case and `.FALSE.` otherwise. `lsamen` also returns `.FALSE.` if `len(ca)` or `len(cb)` is less than *n*.

Input Parameters

n **INTEGER**. The number of characters in *ca* and *cb* to be compared.

ca, cb **CHARACTER*(*)**. Specify two character strings of length at least *n* to be compared. Only the first *n* characters of each string will be accessed.

Output Parameters

val **LOGICAL**. Result of the comparison.

?sum1

Forms the 1-norm of the complex vector using the true absolute value.

```
res = ssum1 ( n, cx, incx )
res = dsum1 ( n, cx, incx )
```

Discussion

Given a complex vector `cx`, `ssum1/dsum1` functions take the sum of the absolute values of vector elements and return a single/double precision result, respectively. These functions are based on `scasum/dzasum` from Level 1 BLAS, but use the true absolute value and were designed for use with `clacon/zlacon`.

Input Parameters

<code>n</code>	INTEGER. Specifies the number of elements in the vector <code>cx</code> .
<code>cx</code>	COMPLEX for <code>ssum1</code> COMPLEX*16 for <code>dsum1</code> Array, DIMENSION at least $(1+(n-1)*abs(incx))$. Contains the input vector whose elements will be summed.
<code>incx</code>	INTEGER. Specifies the spacing between successive elements of <code>cx</code> (<code>incx > 0</code>).

Output Parameters

<code>res</code>	REAL for <code>ssum1</code> DOUBLE PRECISION for <code>dsum1</code> Contains the sum of absolute values.
------------------	--

?gbtf2

Computes the *LU* factorization of a general band matrix using the unblocked version of the algorithm.

```
call sgbtf2 ( m, n, kl, ku, ab, ldab, ipiv, info )
call dgbtf2 ( m, n, kl, ku, ab, ldab, ipiv, info )
call cgbtf2 ( m, n, kl, ku, ab, ldab, ipiv, info )
call zgbtf2 ( m, n, kl, ku, ab, ldab, ipiv, info )
```

Discussion

The routine forms the *LU* factorization of a general real/complex *m* by *n* band matrix *A* with *kl* sub-diagonals and *ku* super-diagonals. The routine uses partial pivoting with row interchanges and implements the unblocked version of the algorithm, calling Level 2 BLAS. See also [?gbtrf](#).

Input Parameters

m **INTEGER**. The number of rows of the matrix *A* ($m \geq 0$).

n **INTEGER**. The number of columns in *A* ($n \geq 0$).

kl **INTEGER**. The number of sub-diagonals within the band of *A* ($kl \geq 0$).

ku **INTEGER**. The number of super-diagonals within the band of *A* ($ku \geq 0$).

ab **REAL** for *sgbtf2*
DOUBLE PRECISION for *dgbtf2*
COMPLEX for *cgbtf2*
COMPLEX*16 for *zgbtf2*.
Array, **DIMENSION** (*ldab*, *).
The array *ab* contains the matrix *A* in band storage (see [Matrix Arguments](#)).

The second dimension of *ab* must be at least $\max(1, n)$.

ldab **INTEGER**. The first dimension of the array *ab*.
($ldab \geq 2kl + ku + 1$)

Output Parameters

- ab* Overwritten by details of the factorization. The diagonal and $kl + ku$ super-diagonals of U are stored in the first $1 + kl + ku$ rows of *ab*. The multipliers used during the factorization are stored in the next kl rows.
- ipiv* INTEGER.
Array, DIMENSION at least $\max(1, \min(m, n))$.
The pivot indices: row i was interchanged with row $ipiv(i)$.
- info* INTEGER. If $info = 0$, the execution is successful.
If $info = -i$, the i th parameter had an illegal value.
If $info = i$, u_{ii} is 0. The factorization has been completed, but U is exactly singular. Division by 0 will occur if you use the factor U for solving a system of linear equations.

?gebd2

Reduces a general matrix to bidiagonal form using an unblocked algorithm.

```
call sgebd2 ( m, n, a, lda, d, e, tauq, taup, work, info )
call dgebd2 ( m, n, a, lda, d, e, tauq, taup, work, info )
call cgebd2 ( m, n, a, lda, d, e, tauq, taup, work, info )
call zgebd2 ( m, n, a, lda, d, e, tauq, taup, work, info )
```

Discussion

The routine reduces a general m -by- n matrix A to upper or lower bidiagonal form B by an orthogonal (unitary) transformation: $Q^T A P = B$

If $m \geq n$, B is upper bidiagonal; if $m < n$, B is lower bidiagonal.

The routine does not form the matrices Q and P explicitly, but represents them as products of elementary reflectors. If $m \geq n$,

$$Q = H(1)H(2)\dots H(n) \quad \text{and} \quad P = G(1)G(2)\dots G(n-1)$$

If $m < n$,

$$Q = H(1)H(2)\dots H(m-1) \quad \text{and} \quad P = G(1)G(2)\dots G(m)$$

Each $H(i)$ and $G(i)$ has the form

$$H(i) = I - \tau_{i,q} v v^T \quad \text{and} \quad G(i) = I - \tau_{i,p} u u^T$$

where $\tau_{i,q}$ and $\tau_{i,p}$ are scalars (real for `sgebd2/dgebd2`, complex for `cgebd2/zgebd2`), and v and u are vectors (real for `sgebd2/dgebd2`, complex for `cgebd2/zgebd2`).

Input Parameters

m **INTEGER.** The number of rows in the matrix A ($m \geq 0$).

n **INTEGER.** The number of columns in A ($n \geq 0$).

a , $work$ **REAL** for `sgebd2`
DOUBLE PRECISION for `dgebd2`
COMPLEX for `cgebd2`
COMPLEX*16 for `zgebd2`.

Arrays:

$a(lda, *)$ contains the m -by- n general matrix A to be reduced. The second dimension of a must be at least $\max(1, n)$.

$work(*)$ is a workspace array, the dimension of $work$ must be at least $\max(1, m, n)$.

lda **INTEGER**. The first dimension of a ; at least $\max(1, m)$.

Output Parameters

a If $m \geq n$, the diagonal and first super-diagonal of a are overwritten with the upper bidiagonal matrix B . Elements below the diagonal, with the array $tauq$, represent the orthogonal/unitary matrix Q as a product of elementary reflectors, and elements above the first superdiagonal, with the array $taup$, represent the orthogonal/unitary matrix P as a product of elementary reflectors.

If $m < n$, the diagonal and first sub-diagonal of a are overwritten by the lower bidiagonal matrix B . Elements below the first subdiagonal, with the array $tauq$, represent the orthogonal/unitary matrix Q as a product of elementary reflectors, and elements above the diagonal, with the array $taup$, represent the orthogonal/unitary matrix P as a product of elementary reflectors.

d **REAL** for single-precision flavors
DOUBLE PRECISION for double-precision flavors.
 Array, **DIMENSION** at least $\max(1, \min(m, n))$.
 Contains the diagonal elements of the bidiagonal matrix B : $d(i) = a(i, i)$.

e **REAL** for single-precision flavors
DOUBLE PRECISION for double-precision flavors.
 Array, **DIMENSION** at least $\max(1, \min(m, n) - 1)$.
 Contains the off-diagonal elements of the bidiagonal matrix B :

If $m \geq n$, $e(i) = a(i, i+1)$ for $i = 1, 2, \dots, n-1$;
If $m < n$, $e(i) = a(i+1, i)$ for $i = 1, 2, \dots, m-1$.

tauq, taup

REAL for sgebd2

DOUBLE PRECISION for dgebd2

COMPLEX for cgebd2

COMPLEX*16 for zgebd2.

Arrays, DIMENSION at least $\max(1, \min(m, n))$.

Contain scalar factors of the elementary reflectors which represent orthogonal/unitary matrices Q and P , respectively.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

?gehd2

Reduces a general square matrix to upper Hessenberg form using an unblocked algorithm.

```
call sgehd2 ( n, ilo, ihi, a, lda, tau, work, info )
call dgehd2 ( n, ilo, ihi, a, lda, tau, work, info )
call cgehd2 ( n, ilo, ihi, a, lda, tau, work, info )
call zgehd2 ( n, ilo, ihi, a, lda, tau, work, info )
```

Discussion

The routine reduces a real/complex general matrix A to upper Hessenberg form H by an orthogonal or unitary similarity transformation $Q' A Q = H$.

The routine does not form the matrix Q explicitly. Instead, Q is represented as a product of *elementary reflectors*.

Input Parameters

n **INTEGER.** The order of the matrix A ($n \geq 0$).

ilo, ihi **INTEGER.** It is assumed that A is already upper triangular in rows and columns $1:ilo - 1$ and $ihi+1:n$. If A has been output by ?gebal, then ilo and ihi must contain the values returned by that routine. Otherwise they should be set to $ilo = 1$ and $ihi = n$. Constraint: $1 \leq ilo \leq ihi \leq \max(1, n)$.

$a, work$ **REAL** for sgehd2
DOUBLE PRECISION for dgehd2
COMPLEX for cgehd2
COMPLEX*16 for zgehd2.
Arrays:
 $a(lda, *)$ contains the n -by- n matrix A to be reduced. The second dimension of a must be at least $\max(1, n)$.
 $work(n)$ is a workspace array.

lda **INTEGER**. The first dimension of *a*; at least $\max(1, n)$.

Output Parameters

a On exit, the upper triangle and the first subdiagonal of *A* are overwritten with the upper Hessenberg matrix *H* and the elements below the first subdiagonal, with the array *tau*, represent the orthogonal/unitary matrix *Q* as a product of elementary reflectors. See *Application Notes* below.

tau **REAL** for *sgehd2*
DOUBLE PRECISION for *dgehd2*
COMPLEX for *cgehd2*
COMPLEX*16 for *zgehd2*.
 Array, **DIMENSION** at least $\max(1, n-1)$.
 Contains the scalar factors of elementary reflectors. See *Application Notes* below.

info **INTEGER**.
 If *info* = 0, the execution is successful.
 If *info* = -*i*, the *i*th parameter had an illegal value.

Application Notes

The matrix *Q* is represented as a product of (*ihi* - *ilo*) elementary reflectors

$$Q = H(ilo) H(ilo + 1) \dots H(ihi - 1)$$

Each *H*(*i*) has the form

$$H(i) = I - tau * v * v'$$

where *tau* is a real/complex scalar, and *v* is a real/complex vector with $v(1:i) = 0$, $v(i+1) = 1$ and $v(ihi+1:n) = 0$.

On exit, $v(i+2:ihi)$ is stored in $a(i+2:ihi, i)$ and *tau* in *tau*(*i*).

The contents of *a* are illustrated by the following example, with $n = 7$, $ilo = 2$ and $ihi = 6$:

on entry

$$\begin{bmatrix} a & a & a & a & a & a & a \\ & a & a & a & a & a & a \\ & & a & a & a & a & a \\ & & & a & a & a & a \\ & & & & a & a & a \\ & & & & & a & a \\ & & & & & & a \end{bmatrix}$$

on exit

$$\begin{bmatrix} a & a & h & h & h & h & a \\ & a & h & h & h & h & a \\ & & h & h & h & h & h \\ & & & v_2 & h & h & h \\ & & & & v_2 & v_3 & h \\ & & & & & v_2 & v_3 & v_4 & h & h & h \\ & & & & & & & & & & a \end{bmatrix}$$

where a denotes an element of the original matrix A , h denotes a modified element of the upper Hessenberg matrix H , and v_i denotes an element of the vector defining $H(i)$.

?gelq2

Computes the LQ factorization of a general rectangular matrix using an unblocked algorithm.

```
call sgelq2 ( m, n, a, lda, tau, work, info )
call dgelq2 ( m, n, a, lda, tau, work, info )
call cgelq2 ( m, n, a, lda, tau, work, info )
call zgelq2 ( m, n, a, lda, tau, work, info )
```

Discussion

The routine computes an LQ factorization of a real/complex m by n matrix A as $A = LQ$.

The routine does not form the matrix Q explicitly. Instead, Q is represented as a product of $\min(m, n)$ elementary reflectors :

$Q = H(k) \dots H(2) H(1)$ (or $Q = H(k)' \dots H(2)' H(1)'$ for complex flavors), where $k = \min(m, n)$

Each $H(i)$ has the form

$$H(i) = I - \tau v v'$$

where τ is a real/complex scalar stored in $\tau(i)$, and v is a real/complex vector with $v(1:i-1) = 0$ and $v(i) = 1$.

On exit, $v(i+1:n)$ is stored in $a(i, i+1:n)$.

Input Parameters

m **INTEGER**. The number of rows in the matrix A ($m \geq 0$).

n **INTEGER**. The number of columns in A ($n \geq 0$).

$a, work$ **REAL** for `sgelq2`
 DOUBLE PRECISION for `dgelq2`
 COMPLEX for `cgelq2`
 COMPLEX*16 for `zgelq2`.

Arrays:

$a(lda, *)$ contains the m -by- n matrix A .

The second dimension of a must be at least $\max(1, n)$.

$work(m)$ is a workspace array.

lda **INTEGER**. The first dimension of a ; at least $\max(1, m)$.

Output Parameters

a Overwritten by the factorization data as follows:
on exit, the elements on and below the diagonal of the array a contain the m -by- $\min(n, m)$ lower trapezoidal matrix L (L is lower triangular if $n \geq m$); the elements above the diagonal, with the array tau , represent the orthogonal/unitary matrix Q as a product of $\min(n, m)$ elementary reflectors.

tau **REAL** for `sgelq2`
DOUBLE PRECISION for `dgelq2`
COMPLEX for `cgelq2`
COMPLEX*16 for `zgelq2`.
Array, **DIMENSION** at least $\max(1, \min(m, n))$.
Contains scalar factors of the elementary reflectors.

$info$ **INTEGER**.
If $info = 0$, the execution is successful.
If $info = -i$, the i th parameter had an illegal value.

?geql2

Computes the *QL* factorization of a general rectangular matrix using an unblocked algorithm.

```
call sgeql2 ( m, n, a, lda, tau, work, info )
call dgeql2 ( m, n, a, lda, tau, work, info )
call cgeql2 ( m, n, a, lda, tau, work, info )
call zgeql2 ( m, n, a, lda, tau, work, info )
```

Discussion

The routine computes a *QL* factorization of a real/complex *m* by *n* matrix *A* as $A = QL$.

The routine does not form the matrix *Q* explicitly. Instead, *Q* is represented as a product of $\min(m, n)$ elementary reflectors :

$$Q = H(k) \dots H(2) H(1), \text{ where } k = \min(m, n)$$

Each *H*(*i*) has the form

$$H(i) = I - \tau v v'$$

where *tau* is a real/complex scalar stored in *tau*(*i*), and *v* is a real/complex vector with $v(m-k+i+1:m) = 0$ and $v(m-k+i) = 1$.

On exit, $v(1:m-k+i-1)$ is stored in *a*(1:*m-k+i-1*, *n-k+i*).

Input Parameters

m **INTEGER**. The number of rows in the matrix *A* ($m \geq 0$).

n **INTEGER**. The number of columns in *A* ($n \geq 0$).

a, *work* **REAL** for *sgeql2*
 DOUBLE PRECISION for *dgeql2*
 COMPLEX for *cgeql2*
 COMPLEX*16 for *zgeql2*.

Arrays:

$a(lda, *)$ contains the m -by- n matrix A .

The second dimension of a must be at least $\max(1, n)$.

$work(m)$ is a workspace array.

lda **INTEGER**. The first dimension of a ; at least $\max(1, m)$.

Output Parameters

a Overwritten by the factorization data as follows:
on exit, if $m \geq n$, the lower triangle of the subarray $a(m-n+1:m, 1:n)$ contains the n -by- n lower triangular matrix L ;
if $m < n$, the elements on and below the $(n-m)$ th superdiagonal contain the m -by- n lower trapezoidal matrix L ; the remaining elements, with the array τ , represent the orthogonal/unitary matrix Q as a product of elementary reflectors.

τ **REAL** for `sgeql2`
DOUBLE PRECISION for `dgeql2`
COMPLEX for `cgeql2`
COMPLEX*16 for `zgeql2`.
Array, **DIMENSION** at least $\max(1, \min(m, n))$.
Contains scalar factors of the elementary reflectors.

$info$ **INTEGER**.
If $info = 0$, the execution is successful.
If $info = -i$, the i th parameter had an illegal value.

?geqr2

Computes the *QR* factorization of a general rectangular matrix using an unblocked algorithm.

```
call sgeqr2 ( m, n, a, lda, tau, work, info )
call dgeqr2 ( m, n, a, lda, tau, work, info )
call cgeqr2 ( m, n, a, lda, tau, work, info )
call zgeqr2 ( m, n, a, lda, tau, work, info )
```

Discussion

The routine computes a *QR* factorization of a real/complex m by n matrix A as $A = QR$.

The routine does not form the matrix Q explicitly. Instead, Q is represented as a product of $\min(m, n)$ elementary reflectors :

$$Q = H(1)H(2) \dots H(k), \text{ where } k = \min(m, n)$$

Each $H(i)$ has the form

$$H(i) = I - \tau v v'$$

where τ is a real/complex scalar stored in $\tau(i)$, and v is a real/complex vector with $v(1:i-1) = 0$ and $v(i) = 1$.

On exit, $v(i+1:m)$ is stored in $a(i+1:m, i)$.

Input Parameters

m **INTEGER**. The number of rows in the matrix A ($m \geq 0$).

n **INTEGER**. The number of columns in A ($n \geq 0$).

$a, work$ **REAL** for *sgeqr2*
 DOUBLE PRECISION for *dgeqr2*
 COMPLEX for *cgeqr2*
 COMPLEX*16 for *zgeqr2*.

Arrays:

$a(lda, *)$ contains the m -by- n matrix A .

The second dimension of a must be at least $\max(1, n)$.

$work(n)$ is a workspace array.

lda **INTEGER**. The first dimension of a ; at least $\max(1, m)$.

Output Parameters

a Overwritten by the factorization data as follows:
on exit, the elements on and above the diagonal of the array a contain the $\min(n,m)$ -by- n upper trapezoidal matrix R (R is upper triangular if $m \geq n$); the elements below the diagonal, with the array tau , represent the orthogonal/unitary matrix Q as a product of elementary reflectors.

tau **REAL** for `sgeqr2`
DOUBLE PRECISION for `dgeqr2`
COMPLEX for `cgeqr2`
COMPLEX*16 for `zgeqr2`.
Array, **DIMENSION** at least $\max(1, \min(m, n))$.
Contains scalar factors of the elementary reflectors.

$info$ **INTEGER**.
If $info = 0$, the execution is successful.
If $info = -i$, the i th parameter had an illegal value.

?gerq2

Computes the RQ factorization of a general rectangular matrix using an unblocked algorithm.

```
call sgerq2 ( m, n, a, lda, tau, work, info )
call dgerq2 ( m, n, a, lda, tau, work, info )
call cgerq2 ( m, n, a, lda, tau, work, info )
call zgerq2 ( m, n, a, lda, tau, work, info )
```

Discussion

The routine computes a RQ factorization of a real/complex m by n matrix A as $A = RQ$.

The routine does not form the matrix Q explicitly. Instead, Q is represented as a product of $\min(m, n)$ elementary reflectors :

$$Q = H(1)H(2) \dots H(k), \text{ where } k = \min(m, n)$$

Each $H(i)$ has the form

$$H(i) = I - \tau v v'$$

where τ is a real/complex scalar stored in $\tau(i)$, and v is a real/complex vector with $v(n-k+i+1:n) = 0$ and $v(n-k+i) = 1$.

On exit, $v(1:n-k+i-1)$ is stored in $a(m-k+i, 1:n-k+i-1)$.

Input Parameters

m **INTEGER**. The number of rows in the matrix A ($m \geq 0$).

n **INTEGER**. The number of columns in A ($n \geq 0$).

$a, work$ **REAL** for `sgerq2`
 DOUBLE PRECISION for `dgerq2`
 COMPLEX for `cgerq2`
 COMPLEX*16 for `zgerq2`.

Arrays:

$a(lda, *)$ contains the m -by- n matrix A .

The second dimension of a must be at least $\max(1, n)$.

$work(m)$ is a workspace array.

lda **INTEGER**. The first dimension of a ; at least $\max(1, m)$.

Output Parameters

a Overwritten by the factorization data as follows:
on exit, if $m \leq n$, the upper triangle of the subarray
 $a(1:m, n-m+1:n)$ contains the m -by- m upper triangular
matrix R ;
if $m > n$, the elements on and above the $(m-n)$ th
subdiagonal contain the m -by- n upper trapezoidal
matrix R ; the remaining elements, with the array tau ,
represent the orthogonal/unitary matrix Q as a product
of elementary reflectors.

tau **REAL** for `sgerq2`
DOUBLE PRECISION for `dgerq2`
COMPLEX for `cgerq2`
COMPLEX*16 for `zgerq2`.
Array, **DIMENSION** at least $\max(1, \min(m, n))$.
Contains scalar factors of the elementary reflectors.

$info$ **INTEGER**.
If $info = 0$, the execution is successful.
If $info = -i$, the i th parameter had an illegal value.

?gesc2

Solves a system of linear equations using the *LU* factorization with complete pivoting computed by ?getc2.

```
call sgetc2 ( n, a, lda, rhs, ipiv, jpiv, scale )
call dgetc2 ( n, a, lda, rhs, ipiv, jpiv, scale )
call cgetc2 ( n, a, lda, rhs, ipiv, jpiv, scale )
call zgetc2 ( n, a, lda, rhs, ipiv, jpiv, scale )
```

Discussion

This routine solves a system of linear equations

$$AX = scale * RHS$$

with a general *n*-by-*n* matrix *A* using the *LU* factorization with complete pivoting computed by ?getc2.

Input Parameters

n **INTEGER**. The order of the matrix *A*.

a, *rhs* **REAL** for sgetc2
 DOUBLE PRECISION for dgetc2
 COMPLEX for cgetc2
 COMPLEX*16 for zgetc2.

Arrays:

a(*lda*,*) contains the *LU* part of the factorization of the *n*-by-*n* matrix *A* computed by ?getc2:

$$A = P L U Q.$$

The second dimension of *a* must be at least max(1, *n*);

rhs(*n*) contains on entry the right hand side vector for the system of equations.

lda **INTEGER**. The first dimension of *a*; at least max(1, *n*).

ipiv INTEGER.
Array, DIMENSION at least max(1,*n*).
The pivot indices: for $1 \leq i \leq n$, row *i* of the matrix has been interchanged with row *ipiv*(*i*).

jpiv INTEGER.
Array, DIMENSION at least max(1,*n*).
The pivot indices: for $1 \leq j \leq n$, column *j* of the matrix has been interchanged with column *jpiv*(*j*).

Output Parameters

rhs On exit, overwritten with the solution vector *X*.

scale REAL for *sgesc2/cgesc2*
DOUBLE PRECISION for *dgesc2/zgesc2*
Contains the scale factor. *scale* is chosen in the range $0 \leq scale \leq 1$ to prevent overflow in the solution.

?getc2

Computes the *LU* factorization with complete pivoting of the general *n*-by-*n* matrix.

```
call sgetc2 ( n, a, lda, ipiv, jpiv, info )
call dgetc2 ( n, a, lda, ipiv, jpiv, info )
call cgetc2 ( n, a, lda, ipiv, jpiv, info )
call zgetc2 ( n, a, lda, ipiv, jpiv, info )
```

Discussion

This routine computes an *LU* factorization with complete pivoting of the *n*-by-*n* matrix *A*. The factorization has the form $A = P * L * U * Q$, where *P* and *Q* are permutation matrices, *L* is lower triangular with unit diagonal elements and *U* is upper triangular.

Input Parameters

n **INTEGER**. The order of the matrix *A* ($n \geq 0$).

a **REAL** for *sgetc2*
DOUBLE PRECISION for *dgetc2*
COMPLEX for *cgetc2*
COMPLEX*16 for *zgetc2*.
Array *a*(*lda*,*) contains the *n*-by-*n* matrix *A* to be factored.
The second dimension of *a* must be at least $\max(1, n)$;

lda **INTEGER**. The first dimension of *a*; at least $\max(1, n)$.

Output Parameters

a On exit, the factors *L* and *U* from the factorization $A = P*L*U*Q$; the unit diagonal elements of *L* are not stored. If *U*(*k*, *k*) appears to be less than **smin**, *U*(*k*, *k*) is given the value of **smin**, i.e., giving a nonsingular perturbed system.

<i>ipiv</i>	<p>INTEGER.</p> <p>Array, DIMENSION at least $\max(1,n)$.</p> <p>The pivot indices: for $1 \leq i \leq n$, row <i>i</i> of the matrix has been interchanged with row <i>ipiv(i)</i>.</p>
<i>jpiv</i>	<p>INTEGER.</p> <p>Array, DIMENSION at least $\max(1,n)$.</p> <p>The pivot indices: for $1 \leq j \leq n$, column <i>j</i> of the matrix has been interchanged with column <i>jpiv(j)</i>.</p>
<i>info</i>	<p>INTEGER.</p> <p>If <i>info</i> = 0, the execution is successful.</p> <p>If <i>info</i> = <i>k</i> > 0, <i>U(k, k)</i> is likely to produce overflow if we try to solve for <i>x</i> in $Ax = b$. So <i>U</i> is perturbed to avoid the overflow.</p>

?getf2

Computes the LU factorization of a general m by n matrix using partial pivoting with row interchanges (unblocked algorithm).

```
call sgetf2 ( m, n, a, lda, ipiv, info )
call dgetf2 ( m, n, a, lda, ipiv, info )
call cgetf2 ( m, n, a, lda, ipiv, info )
call zgetf2 ( m, n, a, lda, ipiv, info )
```

Discussion

The routine computes the *LU* factorization of a general *m*-by-*n* matrix *A* using partial pivoting with row interchanges. The factorization has the form

$$A = PLU$$

where P is a permutation matrix, L is lower triangular with unit diagonal elements (lower trapezoidal if $m > n$) and U is upper triangular (upper trapezoidal if $m < n$).

Input Parameters

m **INTEGER**. The number of rows in the matrix A ($m \geq 0$).

n **INTEGER**. The number of columns in A ($n \geq 0$).

a **REAL** for *sgetf2*
DOUBLE PRECISION for *dgetf2*
COMPLEX for *cgetf2*
COMPLEX*16 for *zgetf2*.
 Array, **DIMENSION** (*lda*, *). Contains the matrix A to be factored. The second dimension of *a* must be at least $\max(1, n)$.

lda **INTEGER**. The first dimension of *a*; at least $\max(1, m)$.

Output Parameters

a Overwritten by L and U . The unit diagonal elements of L are not stored.

ipiv **INTEGER**.
 Array, **DIMENSION** at least $\max(1, \min(m, n))$.
 The pivot indices: for $1 \leq i \leq n$, row i was interchanged with row *ipiv*(i).

info **INTEGER**. If *info*=0, the execution is successful.
 If *info* = - i , the i th parameter had an illegal value.
 If *info* = $i > 0$, u_{ii} is 0. The factorization has been completed, but U is exactly singular. Division by 0 will occur if you use the factor U for solving a system of linear equations.

?gtts2

Solves a system of linear equations with a tridiagonal matrix using the LU factorization computed by ?gttrf.

```
call sgtts2 (itrans, n, nrhs, dl, d, du, du2, ipiv, b, ldb)
call dgtts2 (itrans, n, nrhs, dl, d, du, du2, ipiv, b, ldb)
call cgtts2 (itrans, n, nrhs, dl, d, du, du2, ipiv, b, ldb)
call zgtts2 (itrans, n, nrhs, dl, d, du, du2, ipiv, b, ldb)
```

Discussion

This routine solves for X one of the following systems of linear equations with multiple right hand sides:

$AX = B$ $A^T X = B$ or $A^H X = B$ (for complex matrices only),
with a tridiagonal matrix A using the LU factorization computed by ?gttrf.

Input Parameters

itrans **INTEGER**. Must be 0, 1, or 2.
Indicates the form of the equations being solved:
If *itrans* = 0, then $AX = B$ (no transpose).
If *itrans* = 1, then $A^T X = B$ (transpose).
If *itrans* = 2, then $A^H X = B$ (conjugate transpose).

n **INTEGER**. The order of the matrix A ($n \geq 0$).

nrhs **INTEGER**. The number of right-hand sides, i.e., the number of columns in B ($nrhs \geq 0$).

dl, d, du, du2, b **REAL** for sgtts2
DOUBLE PRECISION for dgtts2
COMPLEX for cgtts2
COMPLEX*16 for zgtts2.
Arrays: *dl*($n - 1$), *d*(n), *du*($n - 1$), *du2*($n - 2$),
b(*ldb*, *nrhs*).
The array *dl* contains the ($n - 1$) multipliers that define

the matrix L from the LU factorization of A .

The array d contains the n diagonal elements of the upper triangular matrix U from the LU factorization of A .

The array du contains the $(n - 1)$ elements of the first super-diagonal of U .

The array $du2$ contains the $(n - 2)$ elements of the second super-diagonal of U .

The array b contains the matrix B whose columns are the right-hand sides for the systems of equations.

ldb **INTEGER**. The leading dimension of b ; must be $ldb \geq \max(1, n)$.

$ipiv$ **INTEGER**.
Array, **DIMENSION** (n).
The pivot indices array, as returned by [?gttrf](#).

Output Parameters

b Overwritten by the solution matrix X .

?labad

Returns the square root of the underflow and overflow thresholds if the exponent-range is very large.

```
call slabad ( small, large )
call dlabad ( small, large )
```

Discussion

This routine takes as input the values computed by [slamch/dlamch](#) for underflow and overflow, and returns the square root of each of these values if the log of *large* is sufficiently large. This subroutine is intended to identify machines with a large exponent range, such as the Crays, and

redefine the underflow and overflow limits to be the square roots of the values computed by `?lamch`. This subroutine is needed because `?lamch` does not compensate for poor arithmetic in the upper half of the exponent range, as is found on a Cray.

Input Parameters

small REAL for `slabad`
 DOUBLE PRECISION for `dlabad`.
 The underflow threshold as computed by `?lamch`.

large REAL for `slabad`
 DOUBLE PRECISION for `dlabad`.
 The overflow threshold as computed by `?lamch`.

Output Parameters

small On exit, if $\log_{10}(\textit{large})$ is sufficiently large, the square root of *small*, otherwise unchanged.

large On exit, if $\log_{10}(\textit{large})$ is sufficiently large, the square root of *large*, otherwise unchanged.

?labrd

Reduces the first nb rows and columns of a general matrix to a bidiagonal form.

```
call slabrd ( m, n, nb, a, lda, d, e, tauq, taup, x, ldx, y, ldy )
call dlabrd ( m, n, nb, a, lda, d, e, tauq, taup, x, ldx, y, ldy )
call clabrd ( m, n, nb, a, lda, d, e, tauq, taup, x, ldx, y, ldy )
call zlabrd ( m, n, nb, a, lda, d, e, tauq, taup, x, ldx, y, ldy )
```

Discussion

The routine reduces the first nb rows and columns of a general m -by- n matrix A to upper or lower bidiagonal form by an orthogonal/unitary transformation $Q' A P$, and returns the matrices X and Y which are needed to apply the transformation to the unreduced part of A .

If $m \geq n$, A is reduced to upper bidiagonal form; if $m < n$, to lower bidiagonal form.

The matrices Q and P are represented as products of elementary reflectors:
 $Q = H(1) H(2) \dots H(nb)$ and $P = G(1) G(2) \dots G(nb)$

Each $H(i)$ and $G(i)$ has the form

$$H(i) = I - \tau_{i,q} v v' \quad \text{and} \quad G(i) = I - \tau_{i,p} u u'$$

where $\tau_{i,q}$ and $\tau_{i,p}$ are scalars, and v and u are vectors.

The elements of the vectors v and u together form the m -by- nb matrix V and the nb -by- n matrix U' which are needed, with X and Y , to apply the transformation to the unreduced part of the matrix, using a block update of the form: $A := A - V * Y' - X * U'$.

This is an auxiliary routine called by `?gebrd`.

Input Parameters

m **INTEGER**. The number of rows in the matrix A ($m \geq 0$).

n **INTEGER**. The number of columns in A ($n \geq 0$).

<i>nb</i>	INTEGER. The number of leading rows and columns of <i>A</i> to be reduced.
<i>a</i>	REAL for slabrd DOUBLE PRECISION for dlabrd COMPLEX for clabrd COMPLEX*16 for zlabrd . Array <i>a</i> (<i>lda</i> , *) contains the matrix <i>A</i> to be reduced. The second dimension of <i>a</i> must be at least max(1, <i>n</i>).
<i>lda</i>	INTEGER. The first dimension of <i>a</i> ; at least max(1, <i>m</i>).
<i>ldx</i>	INTEGER. The first dimension of the output array <i>x</i> ; must beat least max(1, <i>m</i>).
<i>ldy</i>	INTEGER. The first dimension of the output array <i>y</i> ; must beat least max(1, <i>n</i>).

Output Parameters

<i>a</i>	On exit, the first <i>nb</i> rows and columns of the matrix are overwritten; the rest of the array is unchanged. If $m \geq n$, elements on and below the diagonal in the first <i>nb</i> columns, with the array <i>tauq</i> , represent the orthogonal/unitary matrix <i>Q</i> as a product of elementary reflectors; and elements above the diagonal in the first <i>nb</i> rows, with the array <i>taup</i> , represent the orthogonal/unitary matrix <i>P</i> as a product of elementary reflectors. If $m < n$, elements below the diagonal in the first <i>nb</i> columns, with the array <i>tauq</i> , represent the orthogonal/unitary matrix <i>Q</i> as a product of elementary reflectors, and elements on and above the diagonal in the first <i>nb</i> rows, with the array <i>taup</i> , represent the orthogonal/unitary matrix <i>P</i> as a product of elementary reflectors.
<i>d</i> , <i>e</i>	REAL for single-precision flavors DOUBLE PRECISION for double-precision flavors. Arrays, DIMENSION (<i>nb</i>) each.

The array d contains the diagonal elements of the first nb rows and columns of the reduced matrix:

$$d(i) = a(i,i).$$

The array e contains the off-diagonal elements of the first nb rows and columns of the reduced matrix.

$tauq, taup$ REAL for `slabrd`
 DOUBLE PRECISION for `dlabrd`
 COMPLEX for `clabrd`
 COMPLEX*16 for `zlabrd`.

Arrays, DIMENSION (nb) each.
 Contain scalar factors of the elementary reflectors which represent the orthogonal/unitary matrices Q and P , respectively.

x, y REAL for `slabrd`
 DOUBLE PRECISION for `dlabrd`
 COMPLEX for `clabrd`
 COMPLEX*16 for `zlabrd`.

Arrays, dimension $x(ldx, nb)$, $y(ldy, nb)$.
 The array x contains the m -by- nb matrix X required to update the unreduced part of A .
 The array y contains the n -by- nb matrix Y required to update the unreduced part of A .

Application Notes

If $m \geq n$, then for the elementary reflectors $H(i)$ and $G(i)$,

$v(1:i-1) = 0$, $v(i) = 1$, and $v(i:m)$ is stored on exit in $a(i:m, i)$;
 $u(1:i) = 0$, $u(i+1) = 1$, and $u(i+1:n)$ is stored on exit in $a(i, i+1:n)$;
 $tauq$ is stored in $tauq(i)$ and $taup$ in $taup(i)$.

If $m < n$,

$v(1:i) = 0$, $v(i+1) = 1$, and $v(i+1:m)$ is stored on exit in $a(i+2:m, i)$;
 $u(1:i-1) = 0$, $u(i) = 1$, and $u(i:n)$ is stored on exit in $a(i, i+1:n)$;
 $tauq$ is stored in $tauq(i)$ and $taup$ in $taup(i)$.

The contents of a on exit are illustrated by the following examples with $nb = 2$:

$m = 6, n = 5 (m > n)$

$$\begin{bmatrix} 1 & 1 & u_1 & u_1 & u_1 \\ v_1 & 1 & 1 & u_2 & u_2 \\ v_1 & v_2 & a & a & a \\ v_1 & v_2 & a & a & a \\ v_1 & v_2 & a & a & a \\ v_1 & v_2 & a & a & a \end{bmatrix}$$

$m = 5, n = 6 (m < n)$

$$\begin{bmatrix} 1 & u_1 & u_1 & u_1 & u_1 & u_1 \\ 1 & 1 & u_2 & u_2 & u_2 & u_2 \\ v_1 & 1 & a & a & a & a \\ v_1 & v_2 & a & a & a & a \\ v_1 & v_2 & a & a & a & a \end{bmatrix}$$

where a denotes an element of the original matrix which is unchanged, v_i denotes an element of the vector defining $H(i)$, and u_i an element of the vector defining $G(i)$.

?lacon

Estimates the 1-norm of a square matrix, using reverse communication for evaluating matrix-vector products.

```
call slacon ( n, v, x, isgn, est, kase )
call dlacon ( n, v, x, isgn, est, kase )
call clacon ( n, v, x, est, kase )
call zlacon ( n, v, x, est, kase )
```

Discussion

This routine estimates the 1-norm of a square, real/complex matrix A . Reverse communication is used for evaluating matrix-vector products.

Input Parameters

n **INTEGER.** The order of the matrix A ($n \geq 1$).

<i>v</i> , <i>x</i>	REAL for <i>slacon</i> DOUBLE PRECISION for <i>dlacon</i> COMPLEX for <i>clacon</i> COMPLEX*16 for <i>zlacon</i> . Arrays, DIMENSION (<i>n</i>) each. <i>v</i> is a workspace array. <i>x</i> is used as input after an intermediate return.
<i>isgn</i>	INTEGER. Workspace array, DIMENSION (<i>n</i>), used with real flavors only.
<i>kase</i>	INTEGER. On the initial call to <i>?lacon</i> , <i>kase</i> should be 0.

Output Parameters

<i>est</i>	REAL for <i>slacon/clacon</i> DOUBLE PRECISION for <i>dlacon/zlacon</i> An estimate (a lower bound) for norm(<i>A</i>).
<i>kase</i>	On an intermediate return, <i>kase</i> will be 1 or 2, indicating whether <i>x</i> should be overwritten by $A * x$ or $A' * x$. On the final return from <i>?lacon</i> , <i>kase</i> will again be 0.
<i>v</i>	On the final return, $v = A * w$, where $est = \text{norm}(v) / \text{norm}(w)$ (<i>w</i> is not returned).
<i>x</i>	On an intermediate return, <i>x</i> should be overwritten by $A * x$, if <i>kase</i> = 1, $A' * x$, if <i>kase</i> = 2, (where for complex flavors <i>A'</i> is the conjugate transpose of <i>A</i>), and <i>?lacon</i> must be re-called with all the other parameters unchanged.

?lacpy

Copies all or part of one two-dimensional array to another.

```
call slacpy ( uplo, m, n, a, lda, b, ldb )
call dlacpy ( uplo, m, n, a, lda, b, ldb )
call clacpy ( uplo, m, n, a, lda, b, ldb )
call zlacpy ( uplo, m, n, a, lda, b, ldb )
```

Discussion

This routine copies all or part of a two-dimensional matrix A to another matrix B .

Input Parameters

<i>uplo</i>	CHARACTER*1. Specifies the part of the matrix A to be copied to B . If <i>uplo</i> = 'U', the upper triangular part of A is copied. If <i>uplo</i> = 'L', the lower triangular part of A is copied. Otherwise, all of the matrix A is copied.
<i>m</i>	INTEGER. The number of rows in the matrix A ($m \geq 0$).
<i>n</i>	INTEGER. The number of columns in A ($n \geq 0$).
<i>a</i>	REAL for <i>slacpy</i> DOUBLE PRECISION for <i>dlacpy</i> COMPLEX for <i>clacpy</i> COMPLEX*16 for <i>zlacpy</i> . Array $a(lda, *)$, contains the m -by- n matrix A . The second dimension of a must be at least $\max(1, n)$. If <i>uplo</i> = 'U', only the upper triangle or trapezoid is accessed; if <i>uplo</i> = 'L', only the lower triangle or trapezoid is accessed.
<i>lda</i>	INTEGER. The first dimension of a ; $lda \geq \max(1, m)$.

ldb **INTEGER**. The first dimension of the output array *b*;
ldb ≥ max(1, *m*).

Output Parameters

b **REAL** for `slacpy`
DOUBLE PRECISION for `dlacpy`
COMPLEX for `clacpy`
COMPLEX*16 for `zlacpy`.
 Array *b*(*ldb*, *), contains the *m*-by-*n* matrix *B*.
 The second dimension of *b* must be at least max(1,*n*).
 On exit, *B* = *A* in the locations specified by `uplo`.

?ladiv

Performs complex division in real arithmetic, avoiding unnecessary overflow.

```
call sladiv ( a, b, c, d, p, q )
call dladiv ( a, b, c, d, p, q )
res = cladiv ( x, y )
res = zladiv ( x, y )
```

Discussion

The routines `sladiv/dladiv` perform complex division in real arithmetic as

$$p + iq = \frac{a + ib}{c + id}$$

Complex functions `cladiv/zladiv` compute the result as

$$res = x / y ,$$

where *x* and *y* are complex. The computation of *x* / *y* will not overflow on an intermediary step unless the results overflows.

Input Parameters

<i>a, b, c, d</i>	REAL for <code>sldiv</code> DOUBLE PRECISION for <code>dladiv</code> The scalars <i>a</i> , <i>b</i> , <i>c</i> , and <i>d</i> in the above expression (for real flavors only).
<i>x, y</i>	COMPLEX for <code>cladiv</code> COMPLEX*16 for <code>zladiv</code> The complex scalars <i>x</i> and <i>y</i> (for complex flavors only).

Output Parameters

<i>p, q</i>	REAL for <code>sldiv</code> DOUBLE PRECISION for <code>dladiv</code> The scalars <i>p</i> and <i>q</i> in the above expression (for real flavors only).
<i>res</i>	COMPLEX for <code>cladiv</code> DOUBLE COMPLEX for <code>zladiv</code> Contains the result of division x / y .

?lae2

Computes the eigenvalues of a 2-by-2 symmetric matrix.

```
call sla2 ( a, b, c, rt1, rt2 )
call dla2 ( a, b, c, rt1, rt2 )
```

Discussion

The routines `s1a2/d1ae2` compute the eigenvalues of a 2-by-2 symmetric matrix

$$\begin{bmatrix} a & b \\ b & c \end{bmatrix}$$

On return, *rt1* is the eigenvalue of larger absolute value, and *rt2* is the eigenvalue of smaller absolute value.

Input Parameters

a, b, c REAL for `s1ae2`
 DOUBLE PRECISION for `d1ae2`
The elements *a*, *b*, and *c* of the 2-by-2 matrix above.

Output Parameters

rt1, rt2 REAL for `s1ae2`
 DOUBLE PRECISION for `d1ae2`
The computed eigenvalues of larger and smaller absolute value, respectively.

Application Notes

rt1 is accurate to a few ulps barring over/underflow. *rt2* may be inaccurate if there is massive cancellation in the determinant $a*c-b*b$; higher precision or correctly rounded or correctly truncated arithmetic would be needed to compute *rt2* accurately in all cases.

Overflow is possible only if *rt1* is within a factor of 5 of overflow. Underflow is harmless if the input data is 0 or exceeds $underflow_threshold / macheps$.

?laebz

Computes the number of eigenvalues of a real symmetric tridiagonal matrix which are less than or equal to a given value, and performs other tasks required by the routine ?stebz.

```
call slaebz( ijob, nitmax, n, mmax, minp, nbmin, abstol,
            reltol, pivmin, d, e, e2, nval, ab, c, mout, nab,
            work, iwork, info )
call dlaebz( ijob, nitmax, n, mmax, minp, nbmin, abstol,
            reltol, pivmin, d, e, e2, nval, ab, c, mout, nab,
            work, iwork, info )
```

Discussion

The routine ?laebz contains the iteration loops which compute and use the function $N(w)$, which is the count of eigenvalues of a symmetric tridiagonal matrix T less than or equal to its argument w . It performs a choice of two types of loops:

ijob =1, followed by

ijob =2: It takes as input a list of intervals and returns a list of sufficiently small intervals whose union contains the same eigenvalues as the union of the original intervals. The input intervals are $(ab(j,1), ab(j,2)]$, $j=1, \dots, minp$. The output interval $(ab(j,1), ab(j,2)]$ will contain eigenvalues $nab(j,1)+1, \dots, nab(j,2)$, where $1 \leq mout$.

ijob =3: It performs a binary search in each input interval $(ab(j,1), ab(j,2)]$ for a point $w(j)$ such that $N(w(j))=nval(j)$, and uses $c(j)$ as the starting point of the search. If such a $w(j)$ is found, then on output $ab(j,1)=ab(j,2)=w$. If no such $w(j)$ is found, then on output $(ab(j,1), ab(j,2)]$ will be a small interval containing the point where $N(w)$ jumps through $nval(j)$, unless that point lies outside the initial interval.

Note that the intervals are in all cases half-open intervals, that is, of the form $(a,b]$, which includes b but not a .

To avoid underflow, the matrix should be scaled so that its largest element is no greater than $overflow^{**}(1/2) * underflow^{**}(1/4)$ in absolute value. To assure the most accurate computation of small eigenvalues, the matrix should be scaled to be not much smaller than that, either.

Note: the arguments are, in general, **not** checked for unreasonable values.

Input Parameters

- ijob* **INTEGER**. Specifies what is to be done:
 = 1: Compute *nab* for the initial intervals.
 = 2: Perform bisection iteration to find eigenvalues of T .
 = 3: Perform bisection iteration to invert $N(w)$, i.e., to find a point which has a specified number of eigenvalues of T to its left.
 Other values will cause *?laebz* to return with *info*=-1.
- nitmax* **INTEGER**.
 The maximum number of "levels" of bisection to be performed, i.e., an interval of width W will not be made smaller than $2^{(-nitmax)} * W$. If not all intervals have converged after *nitmax* iterations, then *info* is set to the number of non-converged intervals.
- n* **INTEGER**.
 The dimension n of the tridiagonal matrix T . It must be at least 1.
- mmax* **INTEGER**.
 The maximum number of intervals. If more than *mmax* intervals are generated, then *?laebz* will quit with *info*=*mmax*+1.
- minp* **INTEGER**.
 The initial number of intervals. It may not be greater than *mmax*.

<i>nbmin</i>	<p>INTEGER.</p> <p>The smallest number of intervals that should be processed using a vector loop. If zero, then only the scalar loop will be used.</p>
<i>abstol</i>	<p>REAL for <i>slaebz</i> DOUBLE PRECISION for <i>dlaebz</i>.</p> <p>The minimum (absolute) width of an interval. When an interval is narrower than <i>abstol</i>, or than <i>reltol</i> times the larger (in magnitude) endpoint, then it is considered to be sufficiently small, i.e., converged. This must be at least zero.</p>
<i>reltol</i>	<p>REAL for <i>slaebz</i> DOUBLE PRECISION for <i>dlaebz</i>.</p> <p>The minimum relative width of an interval. When an interval is narrower than <i>abstol</i>, or than <i>reltol</i> times the larger (in magnitude) endpoint, then it is considered to be sufficiently small, i.e., converged. Note: this should always be at least <i>radix*machine epsilon</i>.</p>
<i>pivmin</i>	<p>REAL for <i>slaebz</i> DOUBLE PRECISION for <i>dlaebz</i>.</p> <p>The minimum absolute value of a "pivot" in the Sturm sequence loop. This must be at least $\max e(j)*2 * safe_min$ and at least <i>safe_min</i>, where <i>safe_min</i> is at least the smallest number that can divide one without overflow.</p>
<i>d, e, e2</i>	<p>REAL for <i>slaebz</i> DOUBLE PRECISION for <i>dlaebz</i>.</p> <p>Arrays, dimension (<i>n</i>) each.</p> <p>The array <i>d</i> contains the diagonal elements of the tridiagonal matrix <i>T</i>.</p> <p>The array <i>e</i> contains the off-diagonal elements of the tridiagonal matrix <i>T</i> in positions 1 through <i>n</i>-1. <i>e</i>(<i>n</i>) is arbitrary.</p> <p>The array <i>e2</i> contains the squares of the off-diagonal elements of the tridiagonal matrix <i>T</i>. <i>e2</i>(<i>n</i>) is ignored.</p>

nval INTEGER.
Array, dimension (*minp*).
If *ijob*=1 or 2, not referenced.
If *ijob*=3, the desired values of $N(w)$.

ab REAL for *slaebz*
DOUBLE PRECISION for *dlaebz*.
Array, dimension (*mmax*,2)
The endpoints of the intervals. *ab*(*j*,1) is *a*(*j*), the left endpoint of the *j*-th interval, and *ab*(*j*,2) is *b*(*j*), the right endpoint of the *j*-th interval.

c REAL for *slaebz*
DOUBLE PRECISION for *dlaebz*.
Array, dimension (*mmax*)
If *ijob*=1, ignored.
If *ijob*=2, workspace.
If *ijob*=3, then on input *c*(*j*) should be initialized to the first search point in the binary search.

nab INTEGER.
Array, dimension (*mmax*,2)
If *ijob*=2, then on input, *nab*(*i*,*j*) should be set. It must satisfy the condition:
 $N(ab(i,1)) \leq nab(i,1) \leq nab(i,2) \leq N(ab(i,2))$, which means that in interval *i* only eigenvalues *nab*(*i*,1)+1,...,*nab*(*i*,2) will be considered. Usually, *nab*(*i*,*j*)= $N(ab(i,j))$, from a previous call to *?laebz* with *ijob*=1.
If *ijob*=3, normally, *nab* should be set to some distinctive value(s) before *?laebz* is called.

work REAL for *slaebz*
DOUBLE PRECISION for *dlaebz*.
Workspace array, dimension (*mmax*).

iwork INTEGER.
Workspace array, dimension (*mmax*).

Output Parameters

<i>nval</i>	The elements of <i>nval</i> will be reordered to correspond with the intervals in <i>ab</i> . Thus, <i>nval</i> (j) on output will not, in general be the same as <i>nval</i> (j) on input, but it will correspond with the interval (<i>ab</i> (j,1), <i>ab</i> (j,2)] on output.
<i>ab</i>	The input intervals will, in general, be modified, split, and reordered by the calculation.
<i>mout</i>	INTEGER. If <i>ijob</i> =1, the number of eigenvalues in the intervals. If <i>ijob</i> =2 or 3, the number of intervals output. If <i>ijob</i> =3, <i>mout</i> will equal <i>minp</i> .
<i>nab</i>	If <i>ijob</i> =1, then on output <i>nab</i> (i,j) will be set to $N(ab(i,j))$. If <i>ijob</i> =2, then on output, <i>nab</i> (i,j) will contain $\max(na(k), \min(nb(k), N(ab(i,j))))$, where <i>k</i> is the index of the input interval that the output interval (<i>ab</i> (j,1), <i>ab</i> (j,2)] came from, and <i>na</i> (<i>k</i>) and <i>nb</i> (<i>k</i>) are the the input values of <i>nab</i> (<i>k</i> ,1) and <i>nab</i> (<i>k</i> ,2). If <i>ijob</i> =3, then on output, <i>nab</i> (i,j) contains $N(ab(i,j))$, unless $N(w) > nval(i)$ for all search points <i>w</i> , in which case <i>nab</i> (i,1) will not be modified, i.e., the output value will be the same as the input value (modulo reorderings, see <i>nval</i> and <i>ab</i>), or unless $N(w) < nval(i)$ for all search points <i>w</i> , in which case <i>nab</i> (i,2) will not be modified.
<i>info</i>	INTEGER. 0: All intervals converged. 1-- <i>mmax</i> : The last <i>info</i> intervals did not converge. <i>mmax</i> +1: More than <i>mmax</i> intervals were generated.

Application Notes

This routine is intended to be called only by other LAPACK routines, thus the interface is less user-friendly. It is intended for two purposes:

(a) finding eigenvalues. In this case, `?1aebz` should have one or more initial intervals set up in `ab`, and `?1aebz` should be called with `ijob=1`. This sets up `nab`, and also counts the eigenvalues. Intervals with no eigenvalues would usually be thrown out at this point. Also, if not all the eigenvalues in an interval `i` are desired, `nab(i,1)` can be increased or `nab(i,2)` decreased. For example, set `nab(i,1)=nab(i,2)-1` to get the largest eigenvalue. `?1aebz` is then called with `ijob=2` and `mmax` no smaller than the value of `mout` returned by the call with `ijob=1`. After this (`ijob=2`) call, eigenvalues `nab(i,1)+1` through `nab(i,2)` are approximately `ab(i,1)` (or `ab(i,2)`) to the tolerance specified by `abstol` and `reltol`.

(b) finding an interval $(a',b']$ containing eigenvalues $w(f), \dots, w(l)$. In this case, start with a Gershgorin interval (a,b) . Set up `ab` to contain 2 search intervals, both initially (a,b) . One `nval` element should contain `f-1` and the other should contain 1, while `c` should contain `a` and `b`, respectively. `nab(i,1)` should be -1 and `nab(i,2)` should be `n+1`, to flag an error if the desired interval does not lie in (a,b) . `?1aebz` is then called with `ijob=3`. On exit, if $w(f-1) < w(f)$, then one of the intervals -- `j` -- will have `ab(j,1)=ab(j,2)` and `nab(j,1)=nab(j,2)=f-1`, while if, to the specified tolerance, $w(f-k) \approx w(f+r)$, $k > 0$ and $r \geq 0$, then the interval will have $N(ab(j,1))=nab(j,1)=f-k$ and $N(ab(j,2))=nab(j,2)=f+r$. The cases $w(l) < w(l+1)$ and $w(l-r) \approx w(l+k)$ are handled similarly.

?laed0

Used by ?stedc. Computes all eigenvalues and corresponding eigenvectors of an unreduced symmetric tridiagonal matrix using the divide and conquer method.

```
call slaed0(icompg, qsiz, n, d, e, q, ldq, qstore, ldqs,
           work, iwork, info)
call dlaed0(icompg, qsiz, n, d, e, q, ldq, qstore, ldqs,
           work, iwork, info)
call claed0(qsiz, n, d, e, q, ldq, qstore, ldqs, rwork,
           iwork, info)
call zlaed0(qsiz, n, d, e, q, ldq, qstore, ldqs, rwork,
           iwork, info)
```

Discussion

Real flavors of this routine compute all eigenvalues and (optionally) corresponding eigenvectors of a symmetric tridiagonal matrix using the divide and conquer method.

Complex flavors `claed0/zlaed0` compute all eigenvalues of a symmetric tridiagonal matrix which is one diagonal block of those from reducing a dense or band Hermitian matrix and corresponding eigenvectors of the dense or band matrix.

Input Parameters

icompg **INTEGER.** Used with real flavors only.
 If *icompg* = 0, compute eigenvalues only.
 If *icompg* = 1, compute eigenvectors of original dense symmetric matrix also. On entry, the array *q* must contain the orthogonal matrix used to reduce the original matrix to tridiagonal form.
 If *icompg* = 2, compute eigenvalues and eigenvectors of the tridiagonal matrix.

qsiz **INTEGER.**

The dimension of the orthogonal/unitary matrix used to reduce the full matrix to tridiagonal form; $qsiz \geq n$ (for real flavors, $qsiz \geq n$ if $icompg = 1$).

n **INTEGER**. The dimension of the symmetric tridiagonal matrix ($n \geq 0$).

d, e, rwork **REAL** for single-precision flavors
DOUBLE PRECISION for double-precision flavors.
 Arrays:
d*(*)** contains the main diagonal of the tridiagonal matrix. The dimension of ***d must be at least $\max(1, n)$.
e*(*)** contains the off-diagonal elements of the tridiagonal matrix. The dimension of ***e must be at least $\max(1, n-1)$.
rwork*(*)** is a workspace array used in complex flavors only. The dimension of ***rwork must be at least $(1 + 3n + 2n \lg(n) + 3n^2)$, where $\lg(n)$ = smallest integer k such that $2^k \geq n$.

q, qstore **REAL** for **slaed0**
DOUBLE PRECISION for **dlaed0**
COMPLEX for **claed0**
COMPLEX*16 for **zlaed0**.
 Arrays: ***q*(*ldq*, *)**, ***qstore*(*ldqs*, *)**. The second dimension of these arrays must be at least $\max(1, n)$.
For real flavors:
 If $icompg = 0$, array ***q*** is not referenced.
 If $icompg = 1$, on entry, ***q*** is a subset of the columns of the orthogonal matrix used to reduce the full matrix to tridiagonal form corresponding to the subset of the full matrix which is being decomposed at this time.
 If $icompg = 2$, on entry, ***q*** will be the identity matrix.
 The array ***qstore*** is a workspace array referenced only when $icompg = 1$. Used to store parts of the eigenvector matrix when the updating matrix multiplies take place.

For complex flavors:

On entry, q must contain an $qsiz$ -by- n matrix whose columns are unitarily orthonormal. It is a part of the unitary matrix that reduces the full dense Hermitian matrix to a (reducible) symmetric tridiagonal matrix. The array $qstore$ is a workspace array used to store parts of the eigenvector matrix when the updating matrix multiplies take place.

ldq **INTEGER**. The first dimension of the array q ;
 $ldq \geq \max(1, n)$.

$ldqs$ **INTEGER**. The first dimension of the array $qstore$;
 $ldqs \geq \max(1, n)$.

$work$ **REAL** for `slaed0`
DOUBLE PRECISION for `dlaed0`.
 Workspace array, used in real flavors only.
 If $icompg = 0$ or 1 , the dimension of $work$ must be at least $(1 + 3n + 2n \lg(n) + 2n^2)$, where $\lg(n) =$ smallest integer k such that $2^k \geq n$.
 If $icompg = 2$, the dimension of $work$ must be at least $(4n + n^2)$.

$iwork$ **INTEGER**.
 Workspace array.
 For real flavors, if $icompg = 0$ or 1 , and for complex flavors, the dimension of $iwork$ must be at least $(6 + 6n + 5n \lg(n))$,
 For real flavors, If $icompg = 2$, the dimension of $iwork$ must be at least $(3 + 5n)$.

Output Parameters

d On exit, contains eigenvalues in ascending order.

e On exit, the array has been destroyed.

q If $icompg = 2$, on exit, q contains the eigenvectors of the tridiagonal matrix.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

If *info* = *i* > 0, the algorithm failed to compute an eigenvalue while working on the submatrix lying in rows and columns *i*/(*n*+1) through mod(*i*, *n*+1).

?laed1

Used by `sstedc/dstedc`. Computes the updated eigensystem of a diagonal matrix after modification by a rank-one symmetric matrix. Used when the original matrix is tridiagonal.

```
call slaed1( n, d, q, ldq, indxq, rho, cutpnt, work,  
            iwork, info)  
call dlaed1( n, d, q, ldq, indxq, rho, cutpnt, work,  
            iwork, info)
```

Discussion

The routine `?laed1` computes the updated eigensystem of a diagonal matrix after modification by a rank-one symmetric matrix. This routine is used only for the eigenproblem which requires all eigenvalues and eigenvectors of a tridiagonal matrix. `?laed7` handles the case in which eigenvalues only or eigenvalues and eigenvectors of a full symmetric matrix (which was reduced to tridiagonal form) are desired.

$$T = Q(\text{in}) (D(\text{in}) + \text{rho} * Z * Z') Q'(\text{in}) = Q(\text{out}) * D(\text{out}) * Q'(\text{out})$$

where $Z = Q'u$, u is a vector of length n with ones in the `cutpnt` and `(cutpnt + 1)`-th elements and zeros elsewhere. The eigenvectors of the original matrix are stored in Q , and the eigenvalues are in D . The algorithm consists of three stages:

The first stage consists of deflating the size of the problem when there are multiple eigenvalues or if there is a zero in the Z vector. For each such occurrence the dimension of the secular equation problem is reduced by one. This stage is performed by the routine `?laed2`.

The second stage consists of calculating the updated eigenvalues. This is done by finding the roots of the secular equation via the routine `?laed4` (as called by `?laed3`). This routine also calculates the eigenvectors of the current problem.

The final stage consists of computing the updated eigenvectors directly using the updated eigenvalues. The eigenvectors for the current problem are multiplied with the eigenvectors from the overall problem.

Input Parameters

- n* **INTEGER**. The dimension of the symmetric tridiagonal matrix ($n \geq 0$).
- d*, *q*, *work* **REAL** for `slaed1`
DOUBLE PRECISION for `dlaed1`.
 Arrays:
d(*) contains the eigenvalues of the rank-1-perturbed matrix. The dimension of *d* must be at least $\max(1, n)$.
q(*ldq*, *) contains the eigenvectors of the rank-1-perturbed matrix. The second dimension of *q* must be at least $\max(1, n)$.
work(*) is a workspace array, dimension at least $(4n+n^2)$.
- ldq* **INTEGER**. The first dimension of the array *q*;
 $ldq \geq \max(1, n)$.
- indxq* **INTEGER**. Array, dimension (*n*).
 On entry, the permutation which separately sorts the two subproblems in *d* into ascending order.
- rho* **REAL** for `slaed1`
DOUBLE PRECISION for `dlaed1`.
 The subdiagonal entry used to create the rank-1 modification.
- cutpnt* **INTEGER**.
 The location of the last eigenvalue in the leading sub-matrix. $\min(1, n) \leq cutpnt \leq n/2$.
- iwork* **INTEGER**. Workspace array, dimension $(4n)$.

Output Parameters

- d* On exit, contains the eigenvalues of the repaired matrix.

<i>q</i>	On exit, <i>q</i> contains the eigenvectors of the repaired tridiagonal matrix.
<i>indxq</i>	On exit, contains the permutation which will reintegrate the subproblems back into sorted order, that is, $d(\text{indxq}(i = 1, n))$ will be in ascending order.
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = - <i>i</i> , the <i>i</i> th parameter had an illegal value. If <i>info</i> = 1, an eigenvalue did not converge.

?laed2

Used by `sstedc/dstedc`. Merges eigenvalues and deflates secular equation.
 Used when the original matrix is tridiagonal.

```
call slaed2( k, n, n1, d, q, ldq, indxq, rho, z, dlamda,
            w, q2, indx, indx, indxp, coltyp, info)
call dlaed2( k, n, n1, d, q, ldq, indxq, rho, z, dlamda,
            w, q2, indx, indx, indxp, coltyp, info)
```

Discussion

The routine `?laed2` merges the two sets of eigenvalues together into a single sorted set. Then it tries to deflate the size of the problem. There are two ways in which deflation can occur: when two or more eigenvalues are close together or if there is a tiny entry in the `z` vector. For each such occurrence the order of the related secular equation problem is reduced by one.

Input Parameters

`k` **INTEGER**. The number of non-deflated eigenvalues, and the order of the related secular equation ($0 \leq k \leq n$).

`n` **INTEGER**. The dimension of the symmetric tridiagonal matrix ($n \geq 0$).

`n1` **INTEGER**. The location of the last eigenvalue in the leading sub-matrix; $\min(1, n) \leq n1 \leq n/2$.

`d, q, z` **REAL** for `slaed2`
DOUBLE PRECISION for `dlaed2`.
 Arrays:
`d(*)` contains the eigenvalues of the two submatrices to be combined. The dimension of `d` must be at least $\max(1, n)$.

$q(ldq, *)$ contains the eigenvectors of the two submatrices in the two square blocks with corners at $(1,1)$, $(n1,n1)$ and $(n1+1,n1+1)$, (n,n) . The second dimension of q must be at least $\max(1, n)$.

$z(*)$ contains the updating vector (the last row of the first sub-eigenvector matrix and the first row of the second sub-eigenvector matrix).

ldq **INTEGER**. The first dimension of the array q ;
 $ldq \geq \max(1, n)$.

indxq **INTEGER**. Array, dimension (n) .
 On entry, the permutation which separately sorts the two subproblems in d into ascending order. Note that elements in the second half of this permutation must first have $n1$ added to their values.

rho **REAL** for **slaed2**
DOUBLE PRECISION for **dlaed2**.
 On entry, the off-diagonal element associated with the rank-1 cut which originally split the two submatrices which are now being recombined.

indx, indxp **INTEGER**.
 Workspace arrays, dimension (n) each.
 Array *indx* contains the permutation used to sort the contents of *dlamda* into ascending order.
 Array *indxp* contains the permutation used to place deflated values of d at the end of the array.
indxp(1: k) points to the nondeflated d -values and
indxp($k+1:n$) points to the deflated eigenvalues.

coltyp **INTEGER**. Workspace array, dimension (n) .
 During execution, a label which will indicate which of the following types a column in the $q2$ matrix is:
 1 : non-zero in the upper half only;
 2 : dense;
 3 : non-zero in the lower half only;
 4 : deflated.

Output Parameters

- d* On exit, *d* contains the trailing (*n-k*) updated eigenvalues (those which were deflated) sorted into increasing order.
- q* On exit, *q* contains the trailing (*n-k*) updated eigenvectors (those which were deflated) in its last *n-k* columns.
- indxq* Destroyed on exit.
- rho* On exit, *rho* has been modified to the value required by `?laed3`.
- dlamda*, *w*, *q2* REAL for `slaed2`
DOUBLE PRECISION for `dlaed2`.
Arrays: *dlamda*(*n*), *w*(*n*), *q2*($n1^2+(n-n1)^2$).
The array *dlamda* contains a copy of the first *k* eigenvalues which will be used by `?laed3` to form the secular equation.
The array *w* contains the first *k* values of the final deflation-altered *z*-vector which will be passed to `?laed3`.
The array *q2* contains a copy of the first *k* eigenvectors which will be used by `?laed3` in a matrix multiply (`sgemm/dgemm`) to solve for the new eigenvectors.
- indx* INTEGER. Array, dimension (*n*).
The permutation used to arrange the columns of the deflated *q* matrix into three groups: the first group contains non-zero elements only at and above *n1*, the second contains non-zero elements only below *n1*, and the third is dense.
- coltyp* On exit, *coltyp*(*i*) is the number of columns of type *i*, for *i*=1 to 4 only (see the definition of types in the description of *coltyp* in *Input Parameters*).
- info* INTEGER.
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.

?laed3

Used by `sstedc/dstedc`. Finds the roots of the secular equation and updates the eigenvectors. Used when the original matrix is tridiagonal.

```
call slaed3( k, n, n1, d, q, ldq, rho, dlamda, q2, indx,
            ctot, w, s, info)
call dlaed3( k, n, n1, d, q, ldq, rho, dlamda, q2, indx,
            ctot, w, s, info)
```

Discussion

The routine `?laed3` finds the roots of the secular equation, as defined by the values in `d`, `w`, and `rho`, between 1 and `k`. It makes the appropriate calls to `?laed4` and then updates the eigenvectors by multiplying the matrix of eigenvectors of the pair of eigensystems being combined by the matrix of eigenvectors of the `k`-by-`k` system which is solved here.

This code makes very mild assumptions about floating point arithmetic. It will work on machines with a guard digit in add/subtract, or on those binary machines without guard digits which subtract like the Cray X-MP, Cray Y-MP, Cray C-90, or Cray-2. It could conceivably fail on hexadecimal or decimal machines without guard digits, but none are known.

Input Parameters

`k` **INTEGER**. The number of terms in the rational function to be solved by `?laed4` ($k \geq 0$).

`n` **INTEGER**. The number of rows and columns in the `q` matrix. $n \geq k$ (deflation may result in $n > k$).

`n1` **INTEGER**. The location of the last eigenvalue in the leading sub-matrix; $\min(1, n) \leq n1 \leq n/2$.

`q` **REAL** for `slaed3`
DOUBLE PRECISION for `dlaed3`.
 Array `q(ldq, *)`. The second dimension of `q` must be

at least $\max(1, n)$.

Initially, the first k columns of this array are used as workspace.

- ldq* **INTEGER**. The first dimension of the array *q*;
 $ldq \geq \max(1, n)$.
- rho* **REAL** for **slaed3**
DOUBLE PRECISION for **dlaed3**.
 The value of the parameter in the rank one update equation. $rho \geq 0$ required.
- dlambda, q2, w* **REAL** for **slaed3**
DOUBLE PRECISION for **dlaed3**.
 Arrays: *dlambda(k)*, *q2(ldq2, *)*, *w(k)*.
 The first k elements of the array *dlambda* contain the old roots of the deflated updating problem. These are the poles of the secular equation.
 The first k columns of the array *q2* contain the non-deflated eigenvectors for the split problem. The second dimension of *q2* must be at least $\max(1, n)$.
 The first k elements of the array *w* contain the components of the deflation-adjusted updating vector.
- indx* **INTEGER**. Array, dimension (n).
 The permutation used to arrange the columns of the deflated *q* matrix into three groups (see **?laed2**). The rows of the eigenvectors found by **?laed4** must be likewise permuted before the matrix multiply can take place.
- ctot* **INTEGER**. Array, dimension (4).
 A count of the total number of the various types of columns in *q*, as described in *indx*. The fourth column type is any column which has been deflated.
- s* **REAL** for **slaed3**
DOUBLE PRECISION for **dlaed3**.
 Workspace array, dimension $(n1+1)*k$.

Will contain the eigenvectors of the repaired matrix which will be multiplied by the previously accumulated eigenvectors to update the system.

Output Parameters

<i>d</i>	<p>REAL for <code>slaed3</code> DOUBLE PRECISION for <code>dlaed3</code>. Array, dimension at least $\max(1, n)$. $d(i)$ contains the updated eigenvalues for $1 \leq i \leq k$.</p>
<i>q</i>	<p>On exit, the columns 1 to k of q contain the updated eigenvectors.</p>
<i>dlambda</i>	<p>May be changed on output by having lowest order bit set to zero on Cray X-MP, Cray Y-MP, Cray-2, or Cray C-90, as described above.</p>
<i>w</i>	<p>Destroyed on exit.</p>
<i>info</i>	<p>INTEGER. If $info = 0$, the execution is successful. If $info = -i$, the ith parameter had an illegal value. If $info = 1$, an eigenvalue did not converge.</p>

?laed4

Used by `sstedc/dstedc`. Finds a single root of the secular equation.

```
call slaed4 ( n, i, d, z, delta, rho, dlam, info )
call dlaed4 ( n, i, d, z, delta, rho, dlam, info )
```

Discussion

This subroutine computes the i -th updated eigenvalue of a symmetric rank-one modification to a diagonal matrix whose elements are given in the array d , and that

$$D(i) < D(j) \text{ for } i < j$$

and that $\rho > 0$. This is arranged by the calling routine, and is no loss in generality. The rank-one modified system is thus

$$\text{diag}(D) + \rho * Z * \text{transpose}(Z).$$

where we assume the Euclidean norm of Z is 1.

The method consists of approximating the rational functions in the secular equation by simpler interpolating rational functions.

Input Parameters

n **INTEGER**. The length of all arrays.

i **INTEGER**. The index of the eigenvalue to be computed;
 $1 \leq i \leq n$.

d, z **REAL** for `slaed4`
DOUBLE PRECISION for `dlaed4`
Arrays, dimension (n) each.
The array d contains the original eigenvalues. It is assumed that they are in order, $d(i) < d(j)$ for $i < j$.
The array z contains the components of the updating vector Z .

rho REAL for `slaed4`
DOUBLE PRECISION for `dlaed4`
The scalar in the symmetric updating formula.

Output Parameters

delta REAL for `slaed4`
DOUBLE PRECISION for `dlaed4`
Array, dimension (*n*).
If $n \neq 1$, *delta* contains $(d(j) - \lambda_i)$ in its *j*-th component. If $n = 1$, then *delta*(1) = 1. The vector *delta* contains the information necessary to construct the eigenvectors.

dlam REAL for `slaed4`
DOUBLE PRECISION for `dlaed4`
The computed λ_i , the *i*-th updated eigenvalue.

info INTEGER.
If *info* = 0, the execution is successful.
If *info* = 1, the updating process failed.

?laed5

Used by `sstedc/dstedc`.

Solves the 2-by-2 secular equation.

```
call slaed5 ( i, d, z, delta, rho, dlam )
call dlaed5 ( i, d, z, delta, rho, dlam )
```

Discussion

This subroutine computes the *i*-th eigenvalue of a symmetric rank-one modification of a 2-by-2 diagonal matrix

$$\text{diag}(D) + \text{rho} * Z * \text{transpose}(Z).$$

The diagonal elements in the array *D* are assumed to satisfy

$$D(i) < D(j) \text{ for } i < j.$$

We also assume $\rho > 0$ and that the Euclidean norm of the vector Z is one.

Input Parameters

- i* **INTEGER**. The index of the eigenvalue to be computed;
 $1 \leq i \leq 2$.
- d*, *z* **REAL** for `slaed5`
DOUBLE PRECISION for `dlaed5`
 Arrays, dimension (2) each.
 The array *d* contains the original eigenvalues. It is
 assumed that $d(1) < d(2)$.
 The array *z* contains the components of the updating
 vector.
- rho* **REAL** for `slaed5`
DOUBLE PRECISION for `dlaed5`
 The scalar in the symmetric updating formula.

Output Parameters

- delta* **REAL** for `slaed5`
DOUBLE PRECISION for `dlaed5`
 Array, dimension (2).
 The vector *delta* contains the information necessary to
 construct the eigenvectors.
- diam* **REAL** for `slaed5`
DOUBLE PRECISION for `dlaed5`
 The computed λ_i , the *i*-th updated eigenvalue.

?laed6

Used by `sstedc/dstedc`.
 Computes one Newton step in solution
 of the secular equation.

```
call slaed6(kniter, orgati, rho, d, z, finit, tau, info)
```

`call dlaed6(kniter, orgati, rho, d, z, finit, tau, info)`

Discussion

This routine computes the positive or negative root (closest to the origin) of

$$f(x) = \text{rho} + \frac{z(1)}{d(1) - x} + \frac{z(2)}{d(2) - x} + \frac{z(3)}{d(3) - x}$$

It is assumed that if `orgati = .TRUE.` the root is between `d(2)` and `d(3)`; otherwise it is between `d(1)` and `d(2)`.

This routine will be called by `?laed4` when necessary. In most cases, the root sought is the smallest in magnitude, though it might not be in some extremely rare situations.

Input Parameters

<code>kniter</code>	INTEGER. Refer to <code>?laed4</code> for its significance.
<code>orgati</code>	LOGICAL. If <code>orgati = .TRUE.</code> , the needed root is between <code>d(2)</code> and <code>d(3)</code> ; otherwise it is between <code>d(1)</code> and <code>d(2)</code> . See <code>?laed4</code> for further details.
<code>rho</code>	REAL for <code>slaed6</code> DOUBLE PRECISION for <code>dlaed6</code> Refer to the equation for $f(x)$ above.
<code>d, z</code>	REAL for <code>slaed6</code> DOUBLE PRECISION for <code>dlaed6</code> Arrays, dimension (3) each. The array <code>d</code> satisfies <code>d(1) < d(2) < d(3)</code> . Each of the elements in the array <code>z</code> must be positive.
<code>finit</code>	REAL for <code>slaed6</code> DOUBLE PRECISION for <code>dlaed6</code> The value of $f(x)$ at 0. It is more accurate than the one evaluated inside this routine (if someone wants to do so).

Output Parameters

tau REAL for *slaed6*
 DOUBLE PRECISION for *dlaed6*
 The root of the equation for $f(x)$.

info INTEGER.
 If *info* = 0, the execution is successful.
 If *info* = 1, failure to converge.

?laed7

Used by ?stedc. Computes the updated eigensystem of a diagonal matrix after modification by a rank-one symmetric matrix. Used when the original matrix is dense.

```
call slaed7( icompg, n, qsiz, tlvs, curlvl, curpbm, d, q, ldq,
            indxq, rho, cutpnt, qstore, qptr, prmptr, perm, givptr, givcol,
            givnum, work, iwork, info )
call dlaed7( icompg, n, qsiz, tlvs, curlvl, curpbm, d, q, ldq,
            indxq, rho, cutpnt, qstore, qptr, prmptr, perm, givptr, givcol,
            givnum, work, iwork, info )
call claed7( n, cutpnt, qsiz, tlvs, curlvl, curpbm, d, q, ldq, rho,
            indxq, qstore, qptr, prmptr, perm, givptr, givcol, givnum,
            work, rwork, iwork, info )
call zlaed7( n, cutpnt, qsiz, tlvs, curlvl, curpbm, d, q, ldq, rho,
            indxq, qstore, qptr, prmptr, perm, givptr, givcol, givnum,
            work, rwork, iwork, info )
```

Discussion

The routine ?laed7 computes the updated eigensystem of a diagonal matrix after modification by a rank-one symmetric matrix. This routine is used only for the eigenproblem which requires all eigenvalues and optionally eigenvectors of a dense symmetric/Hermitian matrix that has been reduced to tridiagonal form. For real flavors, slaed1/dlaed1 handles the case in which all eigenvalues and eigenvectors of a symmetric tridiagonal matrix are desired.

$$T = Q(\text{in}) (D(\text{in}) + \text{rho} * Z * Z') Q'(\text{in}) = Q(\text{out}) * D(\text{out}) * Q'(\text{out})$$

where $Z = Q'u$, u is a vector of length n with ones in the `cutpnt` and `(cutpnt + 1)`-th elements and zeros elsewhere. The eigenvectors of the original matrix are stored in Q , and the eigenvalues are in D . The algorithm consists of three stages:

The first stage consists of deflating the size of the problem when there are multiple eigenvalues or if there is a zero in the Z vector. For each such occurrence the dimension of the secular equation problem is reduced by one. This stage is performed by the routine `s1aed8/d1aed8` (for real flavors) or by the routine `s1aed2/d1aed2` (for complex flavors).

The second stage consists of calculating the updated eigenvalues. This is done by finding the roots of the secular equation via the routine `?1aed4` (as called by `?1aed9` or `?1aed3`). This routine also calculates the eigenvectors of the current problem.

The final stage consists of computing the updated eigenvectors directly using the updated eigenvalues. The eigenvectors for the current problem are multiplied with the eigenvectors from the overall problem.

Input Parameters

<code>icompq</code>	INTEGER. Used with real flavors only. If <code>icompq = 0</code> , compute eigenvalues only. If <code>icompq = 1</code> , compute eigenvectors of original dense symmetric matrix also. On entry, the array <code>q</code> must contain the orthogonal matrix used to reduce the original matrix to tridiagonal form.
<code>n</code>	INTEGER. The dimension of the symmetric tridiagonal matrix ($n \geq 0$).
<code>cutpnt</code>	INTEGER. The location of the last eigenvalue in the leading sub-matrix. $\min(1, n) \leq \text{cutpnt} \leq n$.
<code>qsiz</code>	INTEGER. The dimension of the orthogonal/unitary matrix used to reduce the full matrix to tridiagonal form; $qsiz \geq n$ (for real flavors, $qsiz \geq n$ if <code>icompq = 1</code>).
<code>tlvls</code>	INTEGER. The total number of merging levels in the overall divide and conquer tree.
<code>curlvl</code>	INTEGER. The current level in the overall merge routine, $0 \leq \text{curlvl} \leq \text{tlvls}$.
<code>curpbm</code>	INTEGER. The current problem in the current level in the overall merge routine (counting from upper left to lower right).

<i>d</i>	<p>REAL for <code>slaed7/claed7</code> DOUBLE PRECISION for <code>dlaed7/zlaed7</code>.</p> <p>Array, dimension at least $\max(1, n)$. Array <i>d</i>(*) contains the eigenvalues of the rank-1-perturbed matrix.</p>
<i>q, work</i>	<p>REAL for <code>slaed7</code> DOUBLE PRECISION for <code>dlaed7</code> COMPLEX for <code>claed7</code> COMPLEX*16 for <code>zlaed7</code>.</p> <p>Arrays: <i>q</i>(<i>ldq</i>, *) contains the the eigenvectors of the rank-1-perturbed matrix. The second dimension of <i>q</i> must be at least $\max(1, n)$.</p> <p><i>work</i>(*) is a workspace array, dimension at least $(3n+qsiz*n)$ for real flavors and at least $(qsiz*n)$ for complex flavors.</p>
<i>ldq</i>	<p>INTEGER. The first dimension of the array <i>q</i>; $ldq \geq \max(1, n)$.</p>
<i>rho</i>	<p>REAL for <code>slaed7/claed7</code> DOUBLE PRECISION for <code>dlaed7/zlaed7</code>.</p> <p>The subdiagonal element used to create the rank-1 modification.</p>
<i>qstore</i>	<p>REAL for <code>slaed7/claed7</code> DOUBLE PRECISION for <code>dlaed7/zlaed7</code>.</p> <p>Array, dimension (n^2+1). Serves also as output parameter. Stores eigenvectors of submatrices encountered during divide and conquer, packed together. <i>qp</i>tr points to beginning of the submatrices.</p>
<i>qp</i> tr	<p>INTEGER. Array, dimension $(n+2)$. Serves also as output parameter. List of indices pointing to beginning of submatrices stored in <i>qstore</i>. The submatrices are numbered starting at the bottom left of the divide and conquer tree, from left to right and bottom to top.</p>

prmptr, *perm*,
givptr **INTEGER**. Arrays, dimension ($n \lg n$) each.

The array *prmptr*(*) contains a list of pointers which indicate where in *perm* a level's permutation is stored. *prmptr*(i+1) - *prmptr*(i) indicates the size of the permutation and also the size of the full, non-deflated problem.

The array *perm*(*) contains the permutations (from deflation and sorting) to be applied to each eigenblock.

The array *givptr*(*) contains a list of pointers which indicate where in *givcol* a level's Givens rotations are stored. *givptr*(i+1) - *givptr*(i) indicates the number of Givens rotations.

givcol **INTEGER**. Array, dimension (2, $n \lg n$).
Each pair of numbers indicates a pair of columns to take place in a Givens rotation.

givnum **REAL** for *slaed7/claed7*
DOUBLE PRECISION for *dlaed7/zlaed7*.
Array, dimension (2, $n \lg n$).
Each number indicates the *S* value to be used in the corresponding Givens rotation.

iwork **INTEGER**. Workspace array, dimension ($4n$).

rwork **REAL** for *claed7*
DOUBLE PRECISION for *zlaed7*.
Workspace array, dimension ($3n+2q_{siz}*n$). Used in complex flavors only.

Output Parameters

d On exit, contains the eigenvalues of the repaired matrix.

q On exit, *q* contains the eigenvectors of the repaired tridiagonal matrix.

indxq **INTEGER.** Array, dimension (*n*).
Contains the permutation which will reintegrate the subproblems back into sorted order, that is, $d(\text{indxq}(i = 1, n))$ will be in ascending order.

info **INTEGER.**
If *info* = 0, the execution is successful.
If *info* = -*i*, the *i*th parameter had an illegal value.
If *info* = 1, an eigenvalue did not converge.

?laed8

Used by ?stedc. Merges eigenvalues and deflates secular equation. Used when the original matrix is dense.

```
call slaed8(  icompq, k, n, qsiz, d, q, ldq, indxq, rho, cutpnt, z,
             dlamda, q2, ldq2, w, perm, givptr, givcol, givnum, indx,
             info )
call dlaed8(  icompq, k, n, qsiz, d, q, ldq, indxq, rho, cutpnt, z,
             dlamda, q2, ldq2, w, perm, givptr, givcol, givnum, indx,
             info )
call claed8(  k, n, qsiz, q, ldq, d, rho, cutpnt, z, dlamda, q2,
             ldq2, w, indxp, indx, indxq, perm, givptr, givcol, givnum,
             info )
call zlaed8(  k, n, qsiz, q, ldq, d, rho, cutpnt, z, dlamda, q2,
             ldq2, w, indxp, indx, indxq, perm, givptr, givcol, givnum,
             info )
```

Discussion

This routine merges the two sets of eigenvalues together into a single sorted set. Then it tries to deflate the size of the problem. There are two ways in which deflation can occur: when two or more eigenvalues are close together or if there is a tiny element in the Z vector. For each such occurrence the order of the related secular equation problem is reduced by one.

Input Parameters

icompq **INTEGER**. Used with real flavors only.
If *icompq* = 0, compute eigenvalues only.
If *icompq* = 1, compute eigenvectors of original dense symmetric matrix also. On entry, the array *q* must contain the orthogonal matrix used to reduce the original matrix to tridiagonal form.

n **INTEGER**. The dimension of the symmetric tridiagonal matrix ($n \geq 0$).

<i>cutpnt</i>	INTEGER. The location of the last eigenvalue in the leading sub-matrix. $\min(1, n) \leq \textit{cutpnt} \leq n$.
<i>qsiz</i>	INTEGER. The dimension of the orthogonal/unitary matrix used to reduce the full matrix to tridiagonal form; $qsiz \geq n$ (for real flavors, $qsiz \geq n$ if <i>icompq</i> = 1).
<i>d, z</i>	REAL for <i>slaed8/claed8</i> DOUBLE PRECISION for <i>dlaed8/zlaed8</i> . Arrays, dimension at least $\max(1, n)$ each. The array <i>d</i> (*) contains the eigenvalues of the two submatrices to be combined. On entry, <i>z</i> (*) contains the updating vector (the last row of the first sub-eigenvector matrix and the first row of the second sub-eigenvector matrix). The contents of <i>z</i> are destroyed by the updating process.
<i>q</i>	REAL for <i>slaed8</i> DOUBLE PRECISION for <i>dlaed8</i> COMPLEX for <i>claed8</i> COMPLEX*16 for <i>zlaed8</i> . Array <i>q</i> (<i>ldq</i> , *). The second dimension of <i>q</i> must be at least $\max(1, n)$. On entry, <i>q</i> contains the eigenvectors of the partially solved system which has been previously updated in matrix multiplies with other partially solved eigensystems. For real flavors, if <i>icompq</i> = 0, <i>q</i> is not referenced.
<i>ldq</i>	INTEGER. The first dimension of the array <i>q</i> ; $ldq \geq \max(1, n)$.
<i>ldq2</i>	INTEGER. The first dimension of the output array <i>q2</i> ; $ldq2 \geq \max(1, n)$.
<i>indxq</i>	INTEGER. Array, dimension (<i>n</i>). The permutation which separately sorts the two sub-problems in <i>d</i> into ascending order. Note that elements in the second half of this permutation must first have <i>cutpnt</i> added to their values in order to be accurate.

rho REAL for `slaed8/claed8`
 DOUBLE PRECISION for `dlaed8/zlaed8`.
 On entry, the off-diagonal element associated with the rank-1 cut which originally split the two submatrices which are now being recombined.

Output Parameters

k INTEGER. The number of non-deflated eigenvalues, and the order of the related secular equation.

d On exit, contains the trailing ($n-k$) updated eigenvalues (those which were deflated) sorted into increasing order.

q On exit, *q* contains the trailing ($n-k$) updated eigenvectors (those which were deflated) in its last ($n-k$) columns.

rho On exit, *rho* has been modified to the value required by `?laed3`.

dlambda, w REAL for `slaed8/claed8`
 DOUBLE PRECISION for `dlaed8/zlaed8`.
 Arrays, dimension (n) each.
 The array *dlambda*(*) contains a copy of the first k eigenvalues which will be used by `?laed3` to form the secular equation.
 The array *w*(*) will hold the first k values of the final deflation-altered Z-vector and will be passed to `?laed3`.

q2 REAL for `slaed8`
 DOUBLE PRECISION for `dlaed8`
 COMPLEX for `claed8`
 COMPLEX*16 for `zlaed8`.
 Array *q2*(*ldq2*, *). The second dimension of *q2* must be at least $\max(1, n)$.
 Contains a copy of the first k eigenvectors which will be used by `slaed7/dlaed7` in a matrix multiply (`sgemm/dgemm`) to update the new eigenvectors.
 For real flavors, if *icompq* = 0, *q2* is not referenced.

indxp, indx INTEGER. Workspace arrays, dimension (n) each.

The array *indx(*)* will contain the permutation used to place deflated values of *d* at the end of the array. On output, *indx(1:k)* points to the nondeflated *d*-values and *indx(k+1:n)* points to the deflated eigenvalues.

The array *indx(*)* will contain the permutation used to sort the contents of *d* into ascending order.

<i>perm</i>	INTEGER. Array, dimension (<i>n</i>). Contains the permutations (from deflation and sorting) to be applied to each eigenblock.
<i>givptr</i>	INTEGER. Contains the number of Givens rotations which took place in this subproblem.
<i>givcol</i>	INTEGER. Array, dimension (2, <i>n</i>). Each pair of numbers indicates a pair of columns to take place in a Givens rotation.
<i>givnum</i>	REAL for <i>slaed8/claed8</i> DOUBLE PRECISION for <i>dlaed8/zlaed8</i> . Array, dimension (2, <i>n</i>). Each number indicates the <i>S</i> value to be used in the corresponding Givens rotation.
<i>info</i>	INTEGER. If <i>info</i> = 0, the execution is successful. If <i>info</i> = <i>-i</i> , the <i>i</i> th parameter had an illegal value.

?laed9

Used by `sstedc/dstedc`.

Finds the roots of the secular equation and updates the eigenvectors. Used when the original matrix is dense.

```
call slaed9( k, kstart, kstop, n, d, q, ldq, rho,
            dlamda, w, s, lds, info )
call dlaed9( k, kstart, kstop, n, d, q, ldq, rho,
            dlamda, w, s, lds, info )
```

Discussion

This routine finds the roots of the secular equation, as defined by the values in `d`, `Z`, and `rho`, between `kstart` and `kstop`. It makes the appropriate calls to `slaed4/dlaed4` and then stores the new matrix of eigenvectors for use in calculating the next level of `Z` vectors.

Input Parameters

`k` **INTEGER**. The number of terms in the rational function to be solved by `slaed4/dlaed4` ($k \geq 0$).

`kstart`, `kstop` **INTEGER**. The updated eigenvalues `lambda(i)`, $kstart \leq i \leq kstop$ are to be computed. $1 \leq kstart \leq kstop \leq k$.

`n` **INTEGER**. The number of rows and columns in the Q matrix. $n \geq k$ (deflation may result in $n > k$).

`q` **REAL** for `slaed9`
DOUBLE PRECISION for `dlaed9`.
Workspace array, dimension $(ldq, *)$. The second dimension of `q` must be at least $\max(1, n)$.

`ldq` **INTEGER**. The first dimension of the array `q`;
 $ldq \geq \max(1, n)$.

<i>rho</i>	<p>REAL for <code>slaed9</code> DOUBLE PRECISION for <code>dlaed9</code></p> <p>The value of the parameter in the rank one update equation. $rho \geq 0$ required.</p>
<i>dlambda, w</i>	<p>REAL for <code>slaed9</code> DOUBLE PRECISION for <code>dlaed9</code></p> <p>Arrays, dimension (<i>k</i>) each.</p> <p>The first <i>k</i> elements of the array <i>dlambda</i>(*) contain the old roots of the deflated updating problem. These are the poles of the secular equation.</p> <p>The first <i>k</i> elements of the array <i>w</i>(*) contain the components of the deflation-adjusted updating vector.</p>
<i>lds</i>	<p>INTEGER. The first dimension of the output array <i>s</i>; $lds \geq \max(1, k)$.</p>

Output Parameters

<i>d</i>	<p>REAL for <code>slaed9</code> DOUBLE PRECISION for <code>dlaed9</code></p> <p>Array, dimension (<i>n</i>). <i>d</i>(<i>i</i>) contains the updated eigenvalues for $kstart \leq i \leq kstop$.</p>
<i>s</i>	<p>REAL for <code>slaed9</code> DOUBLE PRECISION for <code>dlaed9</code>.</p> <p>Array, dimension (<i>lds</i>, *). The second dimension of <i>s</i> must be at least $\max(1, k)$.</p> <p>Will contain the eigenvectors of the repaired matrix which will be stored for subsequent Z vector calculation and multiplied by the previously accumulated eigenvectors to update the system.</p>
<i>info</i>	<p>INTEGER.</p> <p>If <i>info</i> = 0, the execution is successful. If <i>info</i> = -<i>i</i>, the <i>i</i>th parameter had an illegal value. If <i>info</i> = 1, the eigenvalue did not converge.</p>

?laeda

Used by `?stedc`. Computes the Z vector determining the rank-one modification of the diagonal matrix. Used when the original matrix is dense.

```
call slaeda( n, tlvls, curlvl, curpbm, prmptr, perm, givptr, givcol,
            givnum, q, qptr, z, ztemp, info )
call dlaeda( n, tlvls, curlvl, curpbm, prmptr, perm, givptr, givcol,
            givnum, q, qptr, z, ztemp, info )
```

Discussion

The routine `?laeda` computes the Z vector corresponding to the merge step in the `curlvl`-th step of the merge process with `tlvls` steps for the `curpbm`-th problem.

Input Parameters

`n` **INTEGER**. The dimension of the symmetric tridiagonal matrix ($n \geq 0$).

`tlvls` **INTEGER**. The total number of merging levels in the overall divide and conquer tree.

`curlvl` **INTEGER**. The current level in the overall merge routine, $0 \leq \text{curlvl} \leq \text{tlvls}$.

`curpbm` **INTEGER**. The current problem in the current level in the overall merge routine (counting from upper left to lower right).

`prmptr`, `perm`,
`givptr` **INTEGER**. Arrays, dimension ($n \lg n$) each. The array `prmptr(*)` contains a list of pointers which indicate where in `perm` a level's permutation is stored. `prmptr(i+1) - prmptr(i)` indicates the size of the permutation and also the size of the full, non-deflated problem.

The array *perm*(*) contains the permutations (from deflation and sorting) to be applied to each eigenblock.

The array *givptr*(*) contains a list of pointers which indicate where in *givcol* a level's Givens rotations are stored. *givptr*(i+1) - *givptr*(i) indicates the number of Givens rotations.

<i>givcol</i>	INTEGER. Array, dimension (2, <i>n lgn</i>). Each pair of numbers indicates a pair of columns to take place in a Givens rotation.
<i>givnum</i>	REAL for <i>slaeda</i> DOUBLE PRECISION for <i>dlaeda</i> . Array, dimension (2, <i>n lgn</i>). Each number indicates the <i>S</i> value to be used in the corresponding Givens rotation.
<i>q</i>	REAL for <i>slaeda</i> DOUBLE PRECISION for <i>dlaeda</i> . Array, dimension (<i>n</i> ²). Contains the square eigenblocks from previous levels, the starting positions for blocks are given by <i>qp</i> <i>tr</i> .
<i>qp</i> <i>tr</i>	INTEGER. Array, dimension (<i>n</i> +2). Contains a list of pointers which indicate where in <i>q</i> an eigenblock is stored. $\text{sqrt}(\text{qp}tr(i+1) - \text{qp}tr(i))$ indicates the size of the block.
<i>ztemp</i>	REAL for <i>slaeda</i> DOUBLE PRECISION for <i>dlaeda</i> . Workspace array, dimension (<i>n</i>).

Output Parameters

<i>z</i>	REAL for <i>slaeda</i> DOUBLE PRECISION for <i>dlaeda</i> . Array, dimension (<i>n</i>). Contains the updating vector (the last row of the first sub-eigenvector matrix and the first row of the second sub-eigenvector matrix).
----------	--

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = -*i*, the *i*th parameter had an illegal value.

?laein

Computes a specified right or left eigenvector of an upper Hessenberg matrix by inverse iteration.

```
call slaein( rightv, noinit, n, h, ldh, wr, wi, vr, vi, b, ldb,
            work, eps3, smlnum, bignum, info )
call dlaein( rightv, noinit, n, h, ldh, wr, wi, vr, vi, b, ldb,
            work, eps3, smlnum, bignum, info )
call claein( rightv, noinit, n, h, ldh, w, v, b, ldb,
            rwork, eps3, smlnum, info )
call zlaein( rightv, noinit, n, h, ldh, w, v, b, ldb,
            rwork, eps3, smlnum, info )
```

Discussion

The routine ?laein uses inverse iteration to find a right or left eigenvector corresponding to the eigenvalue (w_r, w_i) of a real upper Hessenberg matrix H (for real flavors `slaein/dlaein`) or to the eigenvalue w of a complex upper Hessenberg matrix H (for complex flavors `claein/zlaein`).

Input Parameters

<code>rightv</code>	LOGICAL. If <code>rightv = .TRUE.</code> , compute right eigenvector; if <code>rightv = .FALSE.</code> , compute left eigenvector.
<code>noinit</code>	LOGICAL. If <code>noinit = .TRUE.</code> , no initial vector is supplied in (<code>vr,vi</code>) or in <code>v</code> (for complex flavors); if <code>noinit = .FALSE.</code> , initial vector is supplied in (<code>vr,vi</code>) or in <code>v</code> (for complex flavors).
<code>n</code>	INTEGER. The order of the matrix H ($n \geq 0$).

h REAL for *slaein*
DOUBLE PRECISION for *dlaein*
COMPLEX for *claein*
COMPLEX*16 for *zlaein*.
Array *h(ldh, *)*. The second dimension of *h* must be at least $\max(1, n)$. Contains the upper Hessenberg matrix *H*.

ldh INTEGER. The first dimension of the array *h*;
 $ldh \geq \max(1, n)$.

wr, wi REAL for *slaein*
DOUBLE PRECISION for *dlaein*.
The real and imaginary parts of the eigenvalue of *H* whose corresponding right or left eigenvector is to be computed (for real flavors of the routine).

w COMPLEX for *claein*
COMPLEX*16 for *zlaein*.
The eigenvalue of *H* whose corresponding right or left eigenvector is to be computed (for complex flavors of the routine).

vr, vi REAL for *slaein*
DOUBLE PRECISION for *dlaein*.
Arrays, dimension (*n*) each. Used for real flavors only. On entry, if *noinit* = *.FALSE.* and *wi* = 0.0, *vr* must contain a real starting vector for inverse iteration using the real eigenvalue *wr*;
if *noinit* = *.FALSE.* and *wi* \neq 0.0, *vr* and *vi* must contain the real and imaginary parts of a complex starting vector for inverse iteration using the complex eigenvalue (*wr,wi*); otherwise *vr* and *vi* need not be set.

v COMPLEX for *claein*
COMPLEX*16 for *zlaein*.
Array, dimension (*n*) . Used for complex flavors only. On entry, if *noinit* = *.FALSE.*, *v* must contain a starting vector for inverse iteration; otherwise *v* need not be set.

<i>b</i>	<p>REAL for <code>slaein</code> DOUBLE PRECISION for <code>dlaein</code> COMPLEX for <code>claein</code> COMPLEX*16 for <code>zlaein</code>.</p> <p>Workspace array <code>b(ldb, *)</code>. The second dimension of <code>b</code> must be at least $\max(1, n)$.</p>
<i>ldb</i>	<p>INTEGER. The first dimension of the array <code>b</code>; $ldb \geq n+1$ for real flavors; $ldb \geq \max(1, n)$ for complex flavors.</p>
<i>work</i>	<p>REAL for <code>slaein</code> DOUBLE PRECISION for <code>dlaein</code>.</p> <p>Workspace array, dimension (<code>n</code>). Used for real flavors only.</p>
<i>rwork</i>	<p>REAL for <code>claein</code> DOUBLE PRECISION for <code>zlaein</code>.</p> <p>Workspace array, dimension (<code>n</code>). Used for complex flavors only.</p>
<i>eps3, smlnum</i>	<p>REAL for <code>slaein/claein</code> DOUBLE PRECISION for <code>dlaein/zlaein</code>.</p> <p><code>eps3</code> is a small machine-dependent value which is used to perturb close eigenvalues, and to replace zero pivots. <code>smlnum</code> is a machine-dependent value close to underflow threshold.</p>
<i>bignum</i>	<p>REAL for <code>slaein</code> DOUBLE PRECISION for <code>dlaein</code>.</p> <p><code>bignum</code> is a machine-dependent value close to overflow threshold. Used for real flavors only.</p>

Output Parameters

<i>vr, vi</i>	<p>On exit, if $w_i = 0.0$ (real eigenvalue), <code>vr</code> contains the computed real eigenvector; if $w_i \neq 0.0$ (complex eigenvalue), <code>vr</code> and <code>vi</code> contain the real and imaginary parts of the computed complex eigenvector. The eigenvector is normalized so that the component of</p>
---------------	--

largest magnitude has magnitude 1; here the magnitude of a complex number (x,y) is taken to be $|x| + |y|$.
 vi is not referenced if $wi = 0.0$.

v On exit, v contains the computed eigenvector, normalized so that the component of largest magnitude has magnitude 1; here the magnitude of a complex number (x,y) is taken to be $|x| + |y|$.

$info$ **INTEGER.**
 If $info = 0$, the execution is successful.
 If $info = 1$, inverse iteration did not converge. For real flavors, vr is set to the last iterate, and so is vi if $wi \neq 0.0$. For complex flavors, v is set to the last iterate.

?laev2

Computes the eigenvalues and eigenvectors of a 2-by-2 symmetric/Hermitian matrix.

```
call slaev2 (a, b, c, rt1, rt2, cs1, sn1)
call dlaev2 (a, b, c, rt1, rt2, cs1, sn1)
call claev2 (a, b, c, rt1, rt2, cs1, sn1)
call zlaev2 (a, b, c, rt1, rt2, cs1, sn1)
```

Discussion

This routine performs the eigendecomposition of a 2-by-2 symmetric matrix

$$\begin{bmatrix} a & b \\ b & c \end{bmatrix} \text{ (for slaev2/dlaev2) or Hermitian matrix } \begin{bmatrix} a & b \\ \text{conjg}(b) & c \end{bmatrix}$$

(for claev2/zlaev2).

On return, $rt1$ is the eigenvalue of larger absolute value, $rt2$ of smaller absolute value, and $(cs1, sn1)$ is the unit right eigenvector for $rt1$, giving the decomposition

$$\begin{bmatrix} cs1 & sn1 \\ -sn1 & cs1 \end{bmatrix} \cdot \begin{bmatrix} a & b \\ b & c \end{bmatrix} \cdot \begin{bmatrix} cs1 & -sn1 \\ sn1 & cs1 \end{bmatrix} = \begin{bmatrix} rt1 & 0 \\ 0 & rt2 \end{bmatrix}$$

(for `slaev2/dlaev2`),

or

$$\begin{bmatrix} cs1 & \text{conjg}(sn1) \\ -sn1 & cs1 \end{bmatrix} \cdot \begin{bmatrix} a & b \\ \text{conjg}(b) & c \end{bmatrix} \cdot \begin{bmatrix} cs1 & -\text{conjg}(sn1) \\ sn1 & cs1 \end{bmatrix} = \begin{bmatrix} rt1 & 0 \\ 0 & rt2 \end{bmatrix}$$

(for `claev2/zlaev2`).

Input Parameters

`a, b, c` `REAL` for `slaev2`
 `DOUBLE PRECISION` for `dlaev2`
 `COMPLEX` for `claev2`
 `COMPLEX*16` for `zlaev2`.
 Elements of the input matrix.

Output Parameters

`rt1, rt2` `REAL` for `slaev2/claev2`
 `DOUBLE PRECISION` for `dlaev2/zlaev2`.
 Eigenvalues of larger and smaller absolute value, respectively.

`cs1` `REAL` for `slaev2/claev2`
 `DOUBLE PRECISION` for `dlaev2/zlaev2`.

`sn1` `REAL` for `slaev2`
 `DOUBLE PRECISION` for `dlaev2`
 `COMPLEX` for `claev2`
 `COMPLEX*16` for `zlaev2`.
 The vector (`cs1, sn1`) is the unit right eigenvector for `rt1`.

Application Notes

rt1 is accurate to a few ulps barring over/underflow. *rt2* may be inaccurate if there is massive cancellation in the determinant $a*c-b*b$; higher precision or correctly rounded or correctly truncated arithmetic would be needed to compute *rt2* accurately in all cases. *cs1* and *sn1* are accurate to a few ulps barring over/underflow. Overflow is possible only if *rt1* is within a factor of 5 of overflow. Underflow is harmless if the input data is 0 or exceeds *underflow_threshold* / *macheps*.

?laexc

Swaps adjacent diagonal blocks of a real upper quasi-triangular matrix in Schur canonical form, by an orthogonal similarity transformation.

```
call slaexc ( wantq, n, t, ldt, q, ldq, j1, n1, n2, work, info )
call dlaexc ( wantq, n, t, ldt, q, ldq, j1, n1, n2, work, info )
```

Discussion

This routine swaps adjacent diagonal blocks T_{11} and T_{22} of order 1 or 2 in an upper quasi-triangular matrix T by an orthogonal similarity transformation.

T must be in Schur canonical form, that is, block upper triangular with 1-by-1 and 2-by-2 diagonal blocks; each 2-by-2 diagonal block has its diagonal elements equal and its off-diagonal elements of opposite sign.

Input Parameters

<code>wantq</code>	LOGICAL. If <code>wantq = .TRUE.</code> , accumulate the transformation in the matrix Q ; If <code>wantq = .FALSE.</code> , do not accumulate the transformation.
<code>n</code>	INTEGER. The order of the matrix T ($n \geq 0$).
<code>t, q</code>	REAL for <code>slaexc</code> DOUBLE PRECISION for <code>dlaexc</code> Arrays: <code>t(ldt,*)</code> contains on entry the upper quasi-triangular matrix T , in Schur canonical form. The second dimension of <code>t</code> must be at least $\max(1, n)$.

$q(ldq, *)$ contains on entry, if $wantq = .TRUE.$, the orthogonal matrix Q . If $wantq = .FALSE.$, q is not referenced.

The second dimension of q must be at least $\max(1, n)$.

<i>ldt</i>	INTEGER . The first dimension of t ; at least $\max(1, n)$.
<i>ldq</i>	INTEGER . The first dimension of q ; If $wantq = .FALSE.$, then $ldq \geq 1$. If $wantq = .TRUE.$, then $ldq \geq \max(1, n)$.
<i>j1</i>	INTEGER . The index of the first row of the first block T_{11} .
<i>n1</i>	INTEGER . The order of the first block T_{11} ($n1 = 0, 1$, or 2).
<i>n2</i>	INTEGER . The order of the second block T_{22} ($n2 = 0, 1$, or 2).
<i>work</i>	REAL for slaexc ; DOUBLE PRECISION for dlaexc . Workspace array, DIMENSION (n).

Output Parameters

<i>t</i>	On exit, the updated matrix T , again in Schur canonical form.
<i>q</i>	On exit, if $wantq = .TRUE.$, the updated matrix Q .
<i>info</i>	INTEGER . If $info = 0$, the execution is successful. If $info = 1$, the transformed matrix T would be too far from Schur form; the blocks are not swapped and T and Q are unchanged.

?lag2

Computes the eigenvalues of a 2-by-2 generalized eigenvalue problem, with scaling as necessary to avoid over-/underflow.

```
call slag2 ( a, lda, b, ldb, safmin, scale1, scale2, wr1, wr2, wi )
call dlag2 ( a, lda, b, ldb, safmin, scale1, scale2, wr1, wr2, wi )
```

Discussion

This routine computes the eigenvalues of a 2 x 2 generalized eigenvalue problem $A - w B$, with scaling as necessary to avoid over-/underflow. The scaling factor, s , results in a modified eigenvalue equation

$$s A - w B,$$

where s is a non-negative scaling factor chosen so that w , $w B$, and $s A$ do not overflow and, if possible, do not underflow, either.

Input Parameters

<code>a, b</code>	<code>REAL</code> for <code>slag2</code> <code>DOUBLE PRECISION</code> for <code>dlag2</code>
	Arrays:
	<code>a(lda,2)</code> contains, on entry, the 2 x 2 matrix A . It is assumed that its 1-norm is less than $1/safmin$. Entries less than $\text{sqrt}(safmin) * \text{norm}(A)$ are subject to being treated as zero.
	<code>b(ldb,2)</code> contains, on entry, the 2 x 2 upper triangular matrix B . It is assumed that the one-norm of B is less than $1/safmin$. The diagonals should be at least $\text{sqrt}(safmin)$ times the largest element of B (in absolute value); if a diagonal is smaller than that, then $\pm \text{sqrt}(safmin)$ will be used instead of that diagonal.
<code>lda</code>	<code>INTEGER</code> . The first dimension of <code>a</code> ; <code>lda</code> \geq 2.

ldb INTEGER. The first dimension of *b*; *ldb* ≥ 2.

safmin REAL for *slag2*;
DOUBLE PRECISION for *dlag2*.
The smallest positive number such that $1/\textit{safmin}$ does not overflow. (This should always be `?lamch('S')` - it is an argument in order to avoid having to call `?lamch` frequently.)

Output Parameters

scale1 REAL for *slag2*;
DOUBLE PRECISION for *dlag2*.
A scaling factor used to avoid over-/underflow in the eigenvalue equation which defines the first eigenvalue. If the eigenvalues are complex, then the eigenvalues are $(\textit{wr1} \pm \textit{wi}i)/\textit{scale1}$ (which may lie outside the exponent range of the machine), *scale1*=*scale2*, and *scale1* will always be positive. If the eigenvalues are real, then the first (real) eigenvalue is $\textit{wr1} / \textit{scale1}$, but this may overflow or underflow, and in fact, *scale1* may be zero or less than the underflow threshold if the exact eigenvalue is sufficiently large.

scale2 REAL for *slag2*;
DOUBLE PRECISION for *dlag2*.
A scaling factor used to avoid over-/underflow in the eigenvalue equation which defines the second eigenvalue. If the eigenvalues are complex, then *scale2*=*scale1*. If the eigenvalues are real, then the second (real) eigenvalue is $\textit{wr2} / \textit{scale2}$, but this may overflow or underflow, and in fact, *scale2* may be zero or less than the underflow threshold if the exact eigenvalue is sufficiently large.

wr1 REAL for *slag2*;
DOUBLE PRECISION for *dlag2*.
If the eigenvalue is real, then *wr1* is *scale1* times the

eigenvalue closest to the (2,2) element of AB^{-1} . If the eigenvalue is complex, then $wr1=wr2$ is *scale1* times the real part of the eigenvalues.

wr2

REAL for *slag2*;

DOUBLE PRECISION for *dlag2*.

If the eigenvalue is real, then *wr2* is *scale2* times the other eigenvalue. If the eigenvalue is complex, then $wr1=wr2$ is *scale1* times the real part of the eigenvalues.

wi

REAL for *slag2*;

DOUBLE PRECISION for *dlag2*.

If the eigenvalue is real, then *wi* is zero. If the eigenvalue is complex, then *wi* is *scale1* times the imaginary part of the eigenvalues. *wi* will always be non-negative.

?lags2

Computes 2-by-2 orthogonal matrices U , V , and Q , and applies them to matrices A and B such that the rows of the transformed A and B are parallel.

```
call slags2 ( upper, a1, a2, a3, b1, b2, b3, csu, snu,
             csv, snv, csq, snq )
call dlags2 ( upper, a1, a2, a3, b1, b2, b3, csu, snu,
             csv, snv, csq, snq )
```

Discussion

This routine computes 2-by-2 orthogonal matrices U , V and Q , such that if `upper = .TRUE.`, then

$$U' * A * Q = U' * \begin{bmatrix} A_1 & A_2 \\ 0 & A_3 \end{bmatrix} * Q = \begin{bmatrix} x & 0 \\ x & x \end{bmatrix}$$

and

$$V' * B * Q = V' * \begin{bmatrix} B_1 & B_2 \\ 0 & B_3 \end{bmatrix} * Q = \begin{bmatrix} x & 0 \\ x & x \end{bmatrix}$$

or if `upper = .FALSE.`, then

$$U' * A * Q = U' * \begin{bmatrix} A_1 & 0 \\ A_2 & A_3 \end{bmatrix} * Q = \begin{bmatrix} x & x \\ 0 & x \end{bmatrix}$$

and

$$V' * B * Q = V' * \begin{bmatrix} B_1 & 0 \\ B_2 & B_3 \end{bmatrix} * Q = \begin{bmatrix} x & x \\ 0 & x \end{bmatrix}$$

The rows of the transformed A and B are parallel, where

$$U = \begin{bmatrix} csu & snu \\ -snu & csu \end{bmatrix}, \quad V = \begin{bmatrix} csv & snv \\ -snv & csv \end{bmatrix}, \quad Q = \begin{bmatrix} csq & snq \\ -snq & csq \end{bmatrix}$$

Here Z' denotes the transpose of Z .

Input Parameters

<i>upper</i>	LOGICAL. If <i>upper</i> = <code>.TRUE.</code> , the input matrices A and B are upper triangular; If <i>upper</i> = <code>.FALSE.</code> , the input matrices A and B are lower triangular.
<i>a1, a2, a3</i>	REAL for <code>slags2</code> DOUBLE PRECISION for <code>dlags2</code> On entry, <i>a1</i> , <i>a2</i> and <i>a3</i> are elements of the input 2-by-2 upper (lower) triangular matrix A .
<i>b1, b2, b3</i>	REAL for <code>slags2</code> DOUBLE PRECISION for <code>dlags2</code> On entry, <i>b1</i> , <i>b2</i> and <i>b3</i> are elements of the input 2-by-2 upper (lower) triangular matrix B .

Output Parameters

<i>csu, snu</i>	REAL for <code>slags2</code> DOUBLE PRECISION for <code>dlags2</code> The desired orthogonal matrix U .
<i>csv, snv</i>	REAL for <code>slags2</code> DOUBLE PRECISION for <code>dlags2</code> The desired orthogonal matrix V .
<i>csq, snq</i>	REAL for <code>slags2</code> DOUBLE PRECISION for <code>dlags2</code> The desired orthogonal matrix Q .

?lagtf

Computes an LU factorization of a matrix $T - \lambda I$, where T is a general tridiagonal matrix, and λ a scalar, using partial pivoting with row interchanges.

```
call slagtf ( n, a, lambda, b, c, tol, d, in, info )
call dlagtf ( n, a, lambda, b, c, tol, d, in, info )
```

Discussion

This routine factorizes the matrix $(T - \textit{lambda} * I)$, where T is an n -by- n tridiagonal matrix and \textit{lambda} is a scalar, as

$$T - \textit{lambda} * I = P L U,$$

where P is a permutation matrix, L is a unit lower tridiagonal matrix with at most one non-zero sub-diagonal elements per column and U is an upper triangular matrix with at most two non-zero super-diagonal elements per column. The factorization is obtained by Gaussian elimination with partial pivoting and implicit row scaling. The parameter \textit{lambda} is included in the routine so that ?lagtf may be used, in conjunction with ?lagts, to obtain eigenvectors of T by inverse iteration..

Input Parameters

n **INTEGER**. The order of the matrix T ($n \geq 0$).

a, b, c **REAL** for **slagtf**
DOUBLE PRECISION for **dlagtf**
Arrays, dimension $a(n)$, $b(n-1)$, $c(n-1)$:
On entry, $a(*)$ must contain the diagonal elements of the matrix T .
On entry, $b(*)$ must contain the $(n-1)$ super-diagonal elements of T .
On entry, $c(*)$ must contain the $(n-1)$ sub-diagonal elements of T .

tol REAL for `slagtf`
 DOUBLE PRECISION for `dlagtf`
 On entry, a relative tolerance used to indicate whether or not the matrix $(T - \textit{lambda} * I)$ is nearly singular. *tol* should normally be chose as approximately the largest relative error in the elements of T . For example, if the elements of T are correct to about 4 significant figures, then *tol* should be set to about $5 * 10^{-4}$. If *tol* is supplied as less than `eps`, where `eps` is the relative machine precision, then the value `eps` is used in place of *tol*.

Output Parameters

a On exit, *a* is overwritten by the n diagonal elements of the upper triangular matrix U of the factorization of T .

b On exit, *b* is overwritten by the $n-1$ super-diagonal elements of the matrix U of the factorization of T .

c On exit, *c* is overwritten by the $n-1$ sub-diagonal elements of the matrix L of the factorization of T .

d REAL for `slagtf`
 DOUBLE PRECISION for `dlagtf`
 Array, dimension $(n-2)$.
 On exit, *d* is overwritten by the $n-2$ second super-diagonal elements of the matrix U of the factorization of T .

in INTEGER.
 Array, dimension (n) .
 On exit, *in* contains details of the permutation matrix P . If an interchange occurred at the k -th step of the elimination, then $in(k) = 1$, otherwise $in(k) = 0$. The element $in(n)$ returns the smallest positive integer j such that

$$\text{abs}(u(j,j)) \leq \text{norm}((T - \textit{lambda} * I)(j)) * \textit{tol},$$

where $\text{norm}(A(j))$ denotes the sum of the absolute values of the j -th row of the matrix A . If no such j exists then $in(n)$ is returned as zero. If $in(n)$ is returned as

positive, then a diagonal element of U is small, indicating that $(T - \textit{lambda} * I)$ is singular or nearly singular.

info

INTEGER.

If *info* = 0, the execution is successful.

If *info* = $-k$, the k th parameter had an illegal value.

?lagtm

Performs a matrix-matrix product of the form $C = \alpha AB + \beta C$, where A is a tridiagonal matrix, B and C are rectangular matrices, and α and β are scalars, which may be 0, 1, or -1.

```
call slagtm( trans, n, nrhs, alpha, dl, d, du, x, ldx, beta, b, ldb)
call dlagtm( trans, n, nrhs, alpha, dl, d, du, x, ldx, beta, b, ldb)
call clagtm( trans, n, nrhs, alpha, dl, d, du, x, ldx, beta, b, ldb)
call zlagtm( trans, n, nrhs, alpha, dl, d, du, x, ldx, beta, b, ldb)
```

Discussion

This routine performs a matrix-vector product of the form :

$$B := \alpha A * X + \beta B$$

where A is a tridiagonal matrix of order n , B and X are n -by- $nrhs$ matrices, and α and β are real scalars, each of which may be 0., 1., or -1.

Input Parameters

<i>trans</i>	CHARACTER*1. Must be 'N' or 'T' or 'C'. Indicates the form of the equations: If <i>trans</i> = 'N', then $B := \alpha A * X + \beta B$ (no transpose); If <i>trans</i> = 'T', then $B := \alpha A^T * X + \beta B$ (transpose); If <i>trans</i> = 'C', then $B := \alpha A^H * X + \beta B$ (conjugate transpose)
<i>n</i>	INTEGER. The order of the matrix A ($n \geq 0$).
<i>nrhs</i>	INTEGER. The number of right-hand sides, i.e., the number of columns in X and B ($nrhs \geq 0$).

alpha, beta REAL for *slagtm/clagtm*
DOUBLE PRECISION for *dlagtm/zlagtm*
The scalars α and β . *alpha* must be 0., 1., or -1.; otherwise, it is assumed to be 0. *beta* must be 0., 1., or -1.; otherwise, it is assumed to be 1.

dl, d, du REAL for *slagtm*
DOUBLE PRECISION for *dlagtm*
COMPLEX for *clagtm*
COMPLEX*16 for *zlagtm*.
Arrays: *dl*(*n* - 1), *d*(*n*), *du*(*n* - 1).
The array *dl* contains the (*n* - 1) sub-diagonal elements of *T*.
The array *d* contains the *n* diagonal elements of *T*.
The array *du* contains the (*n* - 1) super-diagonal elements of *T*.

x, b REAL for *slagtm*
DOUBLE PRECISION for *dlagtm*
COMPLEX for *clagtm*
COMPLEX*16 for *zlagtm*.
Arrays:
x(*ldx*, *) contains the *n*-by-*nrhs* matrix *X*. The second dimension of *x* must be at least max(1, *nrhs*).
b(*ldb*, *) contains the *n*-by-*nrhs* matrix *B*. The second dimension of *b* must be at least max(1, *nrhs*).

ldx INTEGER. The leading dimension of the array *x*;
ldx \geq max(1, *n*).

ldb INTEGER. The leading dimension of the array *b*;
ldb \geq max(1, *n*).

Output Parameters

b Overwritten by the matrix expression
 $B := \mathit{alpha} * A * X + \mathit{beta} * B$

?lagts

Solves the system of equations $(T-\lambda I)x = y$ or $(T-\lambda I)^T x = y$, where T is a general tridiagonal matrix and λ a scalar, using the LU factorization computed by ?lagtf.

```
call slagts ( job, n, a, b, c, d, in, y, tol, info )
call dlagts ( job, n, a, b, c, d, in, y, tol, info )
```

Discussion

This routine may be used to solve for x one of the systems of equations:

$$(T - \text{lambda} * I) * x = y \quad \text{or} \quad (T - \text{lambda} * I)' * x = y,$$

where T is an n -by- n tridiagonal matrix, following the factorization of $(T - \text{lambda} * I)$ as

$$T - \text{lambda} * I = P L U,$$

computed by the routine ?lagtf.

The choice of equation to be solved is controlled by the argument *job*, and in each case there is an option to perturb zero or very small diagonal elements of U , this option being intended for use in applications such as inverse iteration.

Input Parameters

job **INTEGER**. Specifies the job to be performed by ?lagts as follows:

- = 1: The equations $(T - \text{lambda} * I)x = y$ are to be solved, but diagonal elements of U are not to be perturbed.
- = -1: The equations $(T - \text{lambda} * I)x = y$ are to be solved and, if overflow would otherwise occur, the diagonal elements of U are to be perturbed. See argument *tol* below.

$= 2$: The equations $(T - \textit{lambda} * I)' x = y$ are to be solved, but diagonal elements of U are not to be perturbed.

$= -2$: The equations $(T - \textit{lambda} * I)' x = y$ are to be solved and, if overflow would otherwise occur, the diagonal elements of U are to be perturbed. See argument *tol* below.

n **INTEGER**. The order of the matrix T ($n \geq 0$).

a, b, c, d **REAL** for **slagts**
DOUBLE PRECISION for **dlagts**
 Arrays, dimension $a(n)$, $b(n-1)$, $c(n-1)$, $d(n-2)$:
 On entry, $a(*)$ must contain the diagonal elements of U as returned from **?lagtf**.
 On entry, $b(*)$ must contain the first super-diagonal elements of U as returned from **?lagtf**.
 On entry, $c(*)$ must contain the sub-diagonal elements of L as returned from **?lagtf**.
 On entry, $d(*)$ must contain the second super-diagonal elements of U as returned from **?lagtf**.

in **INTEGER**.
 Array, dimension (n).
 On entry, $in(*)$ must contain details of the matrix P as returned from **?lagtf**.

y **REAL** for **slagts**
DOUBLE PRECISION for **dlagts**
 Array, dimension (n). On entry, the right hand side vector y .

tol **REAL** for **slagtf**
DOUBLE PRECISION for **dlagtf**.
 On entry, with $job < 0$, *tol* should be the minimum perturbation to be made to very small diagonal elements of U . *tol* should normally be chosen as about $eps * \text{norm}(U)$, where eps is the relative machine

precision, but if *tol* is supplied as non-positive, then it is reset to $eps * \max(\text{abs}(u(i,j)))$. If *job* > 0 then *tol* is not referenced.

Output Parameters

<i>y</i>	On exit, <i>y</i> is overwritten by the solution vector <i>x</i> .
<i>tol</i>	On exit, <i>tol</i> is changed as described in <i>Input Parameters</i> section above, only if <i>tol</i> is non-positive on entry. Otherwise <i>tol</i> is unchanged.
<i>info</i>	<p>INTEGER.</p> <p>If <i>info</i> = 0, the execution is successful.</p> <p>If <i>info</i> = -<i>i</i>, the <i>i</i>th parameter had an illegal value.</p> <p>If <i>info</i> = <i>i</i> > 0, overflow would occur when computing the <i>i</i>th element of the solution vector <i>x</i>. This can only occur when <i>job</i> is supplied as positive and either means that a diagonal element of <i>U</i> is very small, or that the elements of the right-hand side vector <i>y</i> are very large.</p>

?lagv2

Computes the Generalized Schur factorization of a real 2-by-2 matrix pencil (A,B) where B is upper triangular.

```
call slagv2 ( a, lda, b, ldb, alphas, alpha_i, beta, csl,
             snl, csr, snr )
call dlagv2 ( a, lda, b, ldb, alphas, alpha_i, beta, csl,
             snl, csr, snr )
```

Discussion

This routine computes the Generalized Schur factorization of a real 2-by-2 matrix pencil (A,B) where B is upper triangular. The routine computes orthogonal (rotation) matrices given by *csl*, *snl* and *csr*, *snr* such that:

1) if the pencil (A,B) has two real eigenvalues (include 0/0 or 1/0 types), then

$$\begin{bmatrix} a_{11} & a_{12} \\ 0 & a_{22} \end{bmatrix} = \begin{bmatrix} csl & snl \\ -snl & csl \end{bmatrix} \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} csr & -snr \\ snr & csr \end{bmatrix}$$

$$\begin{bmatrix} b_{11} & b_{12} \\ 0 & b_{22} \end{bmatrix} = \begin{bmatrix} csl & snl \\ -snl & csl \end{bmatrix} \begin{bmatrix} b_{11} & b_{12} \\ 0 & b_{22} \end{bmatrix} \begin{bmatrix} csr & -snr \\ snr & csr \end{bmatrix}$$

2) if the pencil (A,B) has a pair of complex conjugate eigenvalues, then

$$\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} = \begin{bmatrix} csl & snl \\ -snl & csl \end{bmatrix} \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} csr & -snr \\ snr & csr \end{bmatrix}$$

$$\begin{bmatrix} b_{11} & 0 \\ 0 & b_{22} \end{bmatrix} = \begin{bmatrix} csl & snl \\ -snl & csl \end{bmatrix} \begin{bmatrix} b_{11} & b_{12} \\ 0 & b_{22} \end{bmatrix} \begin{bmatrix} csr & -snr \\ snr & csr \end{bmatrix}$$

where $b_{11} \geq b_{22} > 0$.

Input Parameters

- a, b* REAL for `slagv2`
 DOUBLE PRECISION for `dlagv2`
 Arrays:
a(lda, 2) contains the 2-by-2 matrix *A*;
b(ldb, 2) contains the upper triangular 2-by-2 matrix *B*.
- lda* INTEGER. The leading dimension of the array *a*;
lda ≥ 2.
- ldb* INTEGER. The leading dimension of the array *b*;
ldb ≥ 2.

Output Parameters

- a* On exit, *a* is overwritten by the “A-part” of the generalized Schur form.
- b* On exit, *b* is overwritten by the “B-part” of the generalized Schur form.
- alphar, alphai, beta* REAL for `slagv2`
 DOUBLE PRECISION for `dlagv2`.
 Arrays, dimension (2) each.
 (*alphar*(*k*) + *i* * *alphai*(*k*))/*beta*(*k*) are the eigenvalues of the pencil (*A*,*B*), *k*=1,2 and *i* = sqrt(-1).
 Note that *beta*(*k*) may be zero.
- csl, snl* REAL for `slagv2`
 DOUBLE PRECISION for `dlagv2`
 The cosine and sine of the left rotation matrix, respectively.
- csr, snr* REAL for `slagv2`
 DOUBLE PRECISION for `dlagv2`
 The cosine and sine of the right rotation matrix, respectively.

?lahqr

Computes the eigenvalues and Schur factorization of an upper Hessenberg matrix, using the double-shift/single-shift QR algorithm.

```
call slahqr ( wantt, wantz, n, ilo, ihi, h, ldh, wr, wi,
             iloz, ihiz, z, ldz, info )
call dlahqr ( wantt, wantz, n, ilo, ihi, h, ldh, wr, wi,
             iloz, ihiz, z, ldz, info )
call clahqr ( wantt, wantz, n, ilo, ihi, h, ldh, w,
             iloz, ihiz, z, ldz, info )
call zlahqr ( wantt, wantz, n, ilo, ihi, h, ldh, w,
             iloz, ihiz, z, ldz, info )
```

Discussion

This routine is an auxiliary routine called by ?hseqr to update the eigenvalues and Schur decomposition already computed by ?hseqr, by dealing with the Hessenberg submatrix in rows and columns *ilo* to *ihi*.

Input Parameters

wantt LOGICAL.
If *wantt* = .TRUE., the full Schur form *T* is required;
If *wantt* = .FALSE., eigenvalues only are required.

wantz LOGICAL.
If *wantz* = .TRUE., the matrix of Schur vectors *Z* is required;
If *wantz* = .FALSE., Schur vectors are not required.

n INTEGER. The order of the matrix *H* ($n \geq 0$).

ilo, ihi INTEGER.
It is assumed that *H* is already upper quasi-triangular in rows and columns *ihi*+1:*n*, and that $H(i\text{lo}, i\text{lo}-1) = 0$ (unless *ilo* = 1). The routine ?lahqr works primarily

with the Hessenberg submatrix in rows and columns *ilo* to *ihi*, but applies transformations to all of *H* if *wantt* = *.TRUE.*.

Constraints:

$1 \leq ilo \leq \max(1, ihi)$; $ihi \leq n$.

h, z

REAL for *slahqr*

DOUBLE PRECISION for *dlahqr*

COMPLEX for *clahqr*

COMPLEX*16 for *zlahqr*.

Arrays:

*h(ldh, *)* contains the upper Hessenberg matrix *H*.

The second dimension of *h* must be at least $\max(1, n)$.

*z(ldz, *)*

If *wantz* = *.TRUE.*, then, on entry, *z* must contain the current matrix *Z* of transformations accumulated by

?hseqr.

If *wantz* = *.FALSE.*, then *z* is not referenced.

The second dimension of *z* must be at least $\max(1, n)$.

ldh

INTEGER. The first dimension of *h*; at least $\max(1, n)$.

ldz

INTEGER. The first dimension of *z*; at least $\max(1, n)$.

iloz, ihiz

INTEGER. Specify the rows of *Z* to which transformations must be applied if *wantz* = *.TRUE.*.

$1 \leq iloz \leq ilo$; $ihi \leq ihiz \leq n$.

Output Parameters

h

On exit, if *wantt* = *.TRUE.*, *H* is upper quasi-triangular (upper triangular for complex flavors) in rows and columns *ilo:ihi*, with any 2-by-2 diagonal blocks in standard form. If *wantt* = *.FALSE.*, the contents of *H* are unspecified on exit.

wr, wi

REAL for *slahqr*

DOUBLE PRECISION for *dlahqr*

Arrays, DIMENSION at least $\max(1, n)$ each. Used with real flavors only.

The real and imaginary parts, respectively, of the

computed eigenvalues *ilo* to *ihi* are stored in the corresponding elements of *wr* and *wi*. If two eigenvalues are computed as a complex conjugate pair, they are stored in consecutive elements of *wr* and *wi*, say the *i*-th and (*i*+1)th, with *wi*(*i*) > 0 and *wi*(*i*+1) < 0. If *wantt* = *.TRUE.*, the eigenvalues are stored in the same order as on the diagonal of the Schur form returned in *H*, with *wr*(*i*) = *H*(*i*,*i*), and, if *H*(*i*:*i*+1, *i*:*i*+1) is a 2-by-2 diagonal block, $wi(i) = \sqrt{H(i+1,i)*H(i,i+1)}$ and $wi(i+1) = -wi(i)$.

w *COMPLEX* for *clahqr*
*COMPLEX*16* for *zlahqr*.
 Array, *DIMENSION* at least max (1, *n*). Used with complex flavors only.
 The computed eigenvalues *ilo* to *ihi* are stored in the corresponding elements of *w*.
 If *wantt* = *.TRUE.*, the eigenvalues are stored in the same order as on the diagonal of the Schur form returned in *H*, with *w*(*i*) = *H*(*i*,*i*).

z If *wantz* = *.TRUE.*, then, on exit *z* has been updated; transformations are applied only to the submatrix *Z*(*iloz*:*ihiz*, *ilo*:*ihi*).

info *INTEGER*.
 If *info* = 0, the execution is successful.
 If *info* = *i* > 0, *?lahqr* failed to compute all the eigenvalues *ilo* to *ihi* in a total of 30*(*ihi*-*ilo*+1) iterations; elements *i*+1:*ihi* of *wr* and *wi* (for *slahqr*/*dlahqr*) or *w* (for *clahqr*/*zlahqr*) contain those eigenvalues which have been successfully computed.

?lahrd

Reduces the first *nb* columns of a general rectangular matrix *A* so that elements below the *k*-th subdiagonal are zero, and returns auxiliary matrices which are needed to apply the transformation to the unreduced part of *A*.

```
call slahrd ( n, k, nb, a, lda, tau, t, ldt, y, ldy )
call dlahrd ( n, k, nb, a, lda, tau, t, ldt, y, ldy )
call clahrd ( n, k, nb, a, lda, tau, t, ldt, y, ldy )
call zlahrd ( n, k, nb, a, lda, tau, t, ldt, y, ldy )
```

Discussion

The routine reduces the first *nb* columns of a real/complex general *n*-by- $(n-k+1)$ matrix *A* so that elements below the *k*-th subdiagonal are zero. The reduction is performed by an orthogonal/unitary similarity transformation $Q' A Q$. The routine returns the matrices *V* and *T* which determine *Q* as a block reflector $I - V T V'$, and also the matrix $Y = A V T$.

The matrix *Q* is represented as products of *nb* elementary reflectors:
 $Q = H(1) H(2) \dots H(nb)$

Each *H*(*i*) has the form

$$H(i) = I - \tau * v * v'$$

where *tau* is a real/complex scalar, and *v* is a real/complex vector.

This is an auxiliary routine called by `?gehrd`.

Input Parameters

- n* **INTEGER.** The order of the matrix *A* ($n \geq 0$).
- k* **INTEGER.** The offset for the reduction. Elements below the *k*-th subdiagonal in the first *nb* columns are reduced to zero.
- nb* **INTEGER.** The number of columns to be reduced.

a REAL for `slahrd`
DOUBLE PRECISION for `dlahrd`
COMPLEX for `clahrd`
COMPLEX*16 for `zlahrd`.
Array $a(lda, n-k+1)$ contains the n -by- $(n-k+1)$ general matrix A to be reduced.

lda INTEGER. The first dimension of **a**; at least $\max(1, n)$.

ldt INTEGER. The first dimension of the output array **t**; must be at least $\max(1, nb)$.

ldy INTEGER. The first dimension of the output array **y**; must be at least $\max(1, n)$.

Output Parameters

a On exit, the elements on and above the k -th subdiagonal in the first nb columns are overwritten with the corresponding elements of the reduced matrix; the elements below the k -th subdiagonal, with the array **tau**, represent the matrix Q as a product of elementary reflectors. The other columns of **a** are unchanged. See *Application Notes* below.

tau REAL for `slahrd`
DOUBLE PRECISION for `dlahrd`
COMPLEX for `clahrd`
COMPLEX*16 for `zlahrd`.
Array, DIMENSION (nb).
Contains scalar factors of the elementary reflectors.

t, y REAL for `slahrd`
DOUBLE PRECISION for `dlahrd`
COMPLEX for `clahrd`
COMPLEX*16 for `zlahrd`.
Arrays, dimension $t(ldt, nb)$, $y(ldy, nb)$.
The array **t** contains upper triangular matrix T .
The array **y** contains the n -by- nb matrix Y .

Application Notes

For the elementary reflector $H(i)$,

$v(1:i+k-1) = 0$, $v(i+k) = 1$; $v(i+k+1:n)$ is stored on exit in $a(i+k+1:n, i)$ and τ is stored in $\tau a(i)$.

The elements of the vectors v together form the $(n-k+1)$ -by- nb matrix V which is needed, with T and Y , to apply the transformation to the unreduced part of the matrix, using an update of the form:

$$A := (I - VT V') * (A - Y V')$$

The contents of A on exit are illustrated by the following example with $n = 7$, $k = 3$ and $nb = 2$:

$$\begin{bmatrix} a & h & a & a & a \\ a & h & a & a & a \\ a & h & a & a & a \\ h & h & a & a & a \\ v_1 & h & a & a & a \\ v_1 & v_2 & a & a & a \\ v_1 & v_2 & a & a & a \end{bmatrix}$$

where a denotes an element of the original matrix A , h denotes a modified element of the upper Hessenberg matrix H , and v_i denotes an element of the vector defining $H(i)$.

?laic1

Applies one step of incremental condition estimation.

```
call slaic1 ( job, j, x, sest, w, gamma, sestpr, s, c )
call dlaic1 ( job, j, x, sest, w, gamma, sestpr, s, c )
call claic1 ( job, j, x, sest, w, gamma, sestpr, s, c )
call zlaic1 ( job, j, x, sest, w, gamma, sestpr, s, c )
```

Discussion

The routine ?laic1 applies one step of incremental condition estimation in its simplest version.

Let x , $\|x\|_2 = 1$ (where $\|a\|_2$ denotes the 2-norm of a), be an approximate singular vector of an j -by- j lower triangular matrix L , such that

$$\|L^*x\|_2 = \text{sest}$$

Then ?laic1 computes sestpr , s , c such that the vector

$$\hat{x} = \begin{bmatrix} s^*x \\ c \end{bmatrix}$$

is an approximate singular vector of

$$\hat{L} = \begin{bmatrix} L & 0 \\ w' & \text{gamma} \end{bmatrix}$$

in the sense that

$$\|\hat{L} \hat{x}\|_2 = \text{sestpr}.$$

Depending on job , an estimate for the largest or smallest singular value is computed.

Note that $[s \ c]^T$ and sestpr^2 is an eigenpair of the system (for slaic1/claic)

$$\text{diag}(sest * sest, 0) + [alpha \quad gamma] * \begin{bmatrix} alpha \\ gamma \end{bmatrix}$$

where $alpha = x' * w$;

or of the system (for `claic1/zlaic`)

$$\text{diag}(sest * sest, 0) + [alpha \quad gamma] * \begin{bmatrix} \text{conjg}(alpha) \\ \text{conjg}(gamma) \end{bmatrix}$$

where $alpha = \text{conjg}(x)' * w$.

Input Parameters

<i>job</i>	INTEGER. If <i>job</i> =1, an estimate for the largest singular value is computed; If <i>job</i> =2, an estimate for the smallest singular value is computed;
<i>j</i>	INTEGER. Length of <i>x</i> and <i>w</i> .
<i>x, w</i>	REAL for <code>slaic1</code> DOUBLE PRECISION for <code>dlaic1</code> COMPLEX for <code>claic1</code> COMPLEX*16 for <code>zlaic1</code> . Arrays, dimension (<i>j</i>) each. Contain vectors <i>x</i> and <i>w</i> , respectively.
<i>sest</i>	REAL for <code>slaic1/claic1</code> ; DOUBLE PRECISION for <code>dlaic1/zlaic1</code> . Estimated singular value of <i>j</i> -by- <i>j</i> matrix <i>L</i> .
<i>gamma</i>	REAL for <code>slaic1</code> DOUBLE PRECISION for <code>dlaic1</code> COMPLEX for <code>claic1</code> COMPLEX*16 for <code>zlaic1</code> . The diagonal element <i>gamma</i> .

Output Parameters

sestpr REAL for `slaic1/claic1`;
 DOUBLE PRECISION for `dlaic1/zlaic1`.
 Estimated singular value of $(j+1)$ -by- $(j+1)$ matrix
 Lhat.

s, c REAL for `slaic1`
 DOUBLE PRECISION for `dlaic1`
 COMPLEX for `claic1`
 COMPLEX*16 for `zlaic1`.
 Sine and cosine needed in forming *xhat*.

?laln2

Solves a 1-by-1 or 2-by-2 linear system of equations of the specified form.

```
call slaln2( ltrans, na, nw, smin, ca, a, lda, d1, d2,
            b, ldb, wr, wi, x, ldx, scale, xnorm, info )
call dlaln2( ltrans, na, nw, smin, ca, a, lda, d1, d2,
            b, ldb, wr, wi, x, ldx, scale, xnorm, info )
```

Discussion

The routine solves a system of the form

$$(ca A - w D) X = s B \quad \text{or} \quad (ca A' - w D) X = s B$$

with possible scaling (s) and perturbation of A (A' means A -transpose.)

A is an na -by- na real matrix, ca is a real scalar, D is an na -by- na real diagonal matrix, w is a real or complex value, and X and B are na -by-1 matrices: real if w is real, complex if w is complex. The parameter na may be 1 or 2.

If w is complex, X and B are represented as na -by-2 matrices, the first column of each being the real part and the second being the imaginary part.

The routine computes the scaling factor s (≤ 1) so chosen that X can be computed without overflow. X is further scaled if necessary to assure that $\text{norm}(ca A - w D) * \text{norm}(X)$ is less than overflow.

If both singular values of $(ca A - w D)$ are less than $smin$, $smin * I$ (where I stands for identity) will be used instead of $(ca A - w D)$. If only one singular value is less than $smin$, one element of $(ca A - w D)$ will be perturbed enough to make the smallest singular value roughly $smin$. If both singular values are at least $smin$, $(ca A - w D)$ will not be perturbed. In any case, the perturbation will be at most some small multiple of

$\max(\text{*smin*}, \text{ulp} * \text{norm}(\text{ca } A - w D))$.

The singular values are computed by infinity-norm approximations, and thus will only be correct to a factor of 2 or so.



NOTE. All input quantities are assumed to be smaller than overflow by a reasonable factor (see *bignum*).

Input Parameters

trans LOGICAL.
 If *trans* = *.TRUE.*, *A*-transpose will be used.
 If *trans* = *.FALSE.*, *A* will be used (not transposed).

na INTEGER. The size of the matrix *A*. May only be 1 or 2.

nw INTEGER. This parameter must be 1 if *w* is real, and 2 if *w* is complex. May only be 1 or 2.

smin REAL for *slaln2*
 DOUBLE PRECISION for *dlaln2*.
 The desired lower bound on the singular values of *A*. This should be a safe distance away from underflow or overflow, for example, between (*underflow/machine_precision*) and (*machine_precision * overflow*). (See *bignum* and *ulp*).

ca REAL for *slaln2*
 DOUBLE PRECISION for *dlaln2*.
 The coefficient by which *A* is multiplied.

a REAL for *slaln2*
 DOUBLE PRECISION for *dlaln2*.
 Array, DIMENSION (*lda,na*). The *na*-by-*na* matrix *A*.

lda INTEGER. The leading dimension of *a*. Must be at least *na*.

d1, d2 REAL for *slaln2*
 DOUBLE PRECISION for *dlaln2*.
 The (1,1) and (2,2) elements in the diagonal matrix *D*, respectively. *d2* is not used if *nw* = 1.

<i>b</i>	<p>REAL for <code>s1aln2</code> DOUBLE PRECISION for <code>dlaln2</code>. Array, DIMENSION (<i>ldb</i>,<i>nw</i>). The <i>na</i>-by-<i>nw</i> matrix <i>B</i> (right-hand side). If <i>nw</i> =2 (<i>w</i> is complex), column 1 contains the real part of <i>B</i> and column 2 contains the imaginary part.</p>
<i>ldb</i>	<p>INTEGER. The leading dimension of <i>b</i>. Must be at least <i>na</i>.</p>
<i>wr, wi</i>	<p>REAL for <code>s1aln2</code> DOUBLE PRECISION for <code>dlaln2</code>. The real and imaginary part of the scalar <i>w</i>, respectively. <i>wi</i> is not used if <i>nw</i> = 1.</p>
<i>ldx</i>	<p>INTEGER. The leading dimension of the output array <i>x</i>. Must be at least <i>na</i>.</p>

Output Parameters

<i>x</i>	<p>REAL for <code>s1aln2</code> DOUBLE PRECISION for <code>dlaln2</code>. Array, DIMENSION (<i>ldx</i>,<i>nw</i>). The <i>na</i>-by-<i>nw</i> matrix <i>X</i> (unknowns), as computed by the routine. If <i>nw</i> = 2 (<i>w</i> is complex), on exit, column 1 will contain the real part of <i>X</i> and column 2 will contain the imaginary part.</p>
<i>scale</i>	<p>REAL for <code>s1aln2</code> DOUBLE PRECISION for <code>dlaln2</code>. The scale factor that <i>B</i> must be multiplied by to insure that overflow does not occur when computing <i>X</i>. Thus $(ca A - w D) X$ will be <i>scale</i>*<i>B</i>, not <i>B</i> (ignoring perturbations of <i>A</i>.) It will be at most 1.</p>
<i>xnorm</i>	<p>REAL for <code>s1aln2</code> DOUBLE PRECISION for <code>dlaln2</code>. The infinity-norm of <i>X</i>, when <i>X</i> is regarded as an <i>na</i>-by-<i>nw</i> real matrix.</p>
<i>info</i>	<p>INTEGER. An error flag. It will be zero if no error occurs, a negative number if an argument is in error, or a positive</p>

number if $(ca A - w D)$ had to be perturbed.

The possible values are:

If $info = 0$: no error occurred, and $(ca A - w D)$ did not have to be perturbed.

If $info = 1$: $(ca A - w D)$ had to be perturbed to make its smallest (or only) singular value greater than $smin$.



NOTE. *In the interests of speed, this routine does not check the inputs for errors.*

?lals0

Applies back multiplying factors in solving the least squares problem using divide and conquer SVD approach.

Used by ?gelsd.

```
call slals0( icompg, nl, nr, sqre, nrhs, b, ldb, bx, ldbx, perm,
            givptr, givcol, ldgcol, givnum, ldgnum, poles, difl, difr, z,
            k, c, s, work, info )
call dlals0( icompg, nl, nr, sqre, nrhs, b, ldb, bx, ldbx, perm,
            givptr, givcol, ldgcol, givnum, ldgnum, poles, difl, difr, z,
            k, c, s, work, info )
call clals0 ( icompg, nl, nr, sqre, nrhs, b, ldb, bx, ldbx, perm,
            givptr, givcol, ldgcol, givnum, ldgnum, poles, difl, difr, z,
            k, c, s, rwork, info )
call zlals0 ( icompg, nl, nr, sqre, nrhs, b, ldb, bx, ldbx, perm,
            givptr, givcol, ldgcol, givnum, ldgnum, poles, difl, difr, z,
            k, c, s, rwork, info )
```

Discussion

The routine applies back the multiplying factors of either the left or right singular vector matrix of a diagonal matrix appended by a row to the right hand side matrix B in solving the least squares problem using the divide-and-conquer SVD approach.

For the left singular vector matrix, three types of orthogonal matrices are involved:

(1L) Givens rotations: the number of such rotations is *givptr*; the pairs of columns/rows they were applied to are stored in *givcol*; and the *c*- and *s*-values of these rotations are stored in *givnum*.

(2L) Permutation. The $(nl+1)$ -st row of B is to be moved to the first row, and for $j=2:n$, *perm*(j)-th row of B is to be moved to the j -th row.

(3L) The left singular vector matrix of the remaining matrix.

For the right singular vector matrix, four types of orthogonal matrices are involved:

(1R) The right singular vector matrix of the remaining matrix.

(2R) If *sqre* = 1, one extra Givens rotation to generate the right null space.

(3R) The inverse transformation of (2L).

(4R) The inverse transformation of (1L).

Input Parameters

<i>icompq</i>	INTEGER. Specifies whether singular vectors are to be computed in factored form: If <i>icompq</i> = 0: Left singular vector matrix. If <i>icompq</i> = 1: Right singular vector matrix.
<i>nl</i>	INTEGER. The row dimension of the upper block. $nl \geq 1$.
<i>nr</i>	INTEGER. The row dimension of the lower block. $nr \geq 1$.
<i>sqre</i>	INTEGER. If <i>sqre</i> = 0: the lower block is an <i>nr</i> -by- <i>nr</i> square matrix. If <i>sqre</i> = 1: the lower block is an <i>nr</i> -by- $(nr+1)$

rectangular matrix. The bidiagonal matrix has row dimension $n = nl + nr + 1$, and column dimension $m = n + sqre$.

nrhs **INTEGER**. The number of columns of *b* and *bx*. Must be at least 1.

b **REAL** for *slals0*
DOUBLE PRECISION for *dlals0*
COMPLEX for *clals0*
COMPLEX*16 for *zlals0*.
Array, **DIMENSION** (*ldb*, *nrhs*). Contains the right hand sides of the least squares problem in rows 1 through *m*.

ldb **INTEGER**. The leading dimension of *b*. Must be at least $\max(1, \max(m, n))$.

bx **REAL** for *slals0*
DOUBLE PRECISION for *dlals0*
COMPLEX for *clals0*
COMPLEX*16 for *zlals0*.
Workspace array, **DIMENSION** (*ldb*, *nrhs*).

ldb **INTEGER**. The leading dimension of *bx*.

perm **INTEGER**.
Array, **DIMENSION** (*n*). The permutations (from deflation and sorting) applied to the two blocks.

givptr **INTEGER**. The number of Givens rotations which took place in this subproblem.

givcol **INTEGER**.
Array, **DIMENSION** (*ldgcol*, 2). Each pair of numbers indicates a pair of rows/columns involved in a Givens rotation.

ldgcol **INTEGER**. The leading dimension of *givcol*, must be at least *n*.

<i>givnum</i>	<p>REAL for <code>slals0 / clals0</code> DOUBLE PRECISION for <code>dlals0 / zlals0</code> Array, DIMENSION (<i>ldgnum</i>, 2). Each number indicates the <i>c</i> or <i>s</i> value used in the corresponding Givens rotation.</p>
<i>ldgnum</i>	<p>INTEGER. The leading dimension of arrays <i>diffr</i>, <i>poles</i> and <i>givnum</i>, must be at least <i>k</i>.</p>
<i>poles</i>	<p>REAL for <code>slals0 / clals0</code> DOUBLE PRECISION for <code>dlals0 / zlals0</code> Array, DIMENSION (<i>ldgnum</i>, 2). On entry, <i>poles</i>(1:<i>k</i>, 1) contains the new singular values obtained from solving the secular equation, and <i>poles</i>(1:<i>k</i>, 2) is an array containing the poles in the secular equation.</p>
<i>difl</i>	<p>REAL for <code>slals0 / clals0</code> DOUBLE PRECISION for <code>dlals0 / zlals0</code> Array, DIMENSION (<i>k</i>). On entry, <i>difl</i>(<i>i</i>) is the distance between <i>i</i>-th updated (undeflated) singular value and the <i>i</i>-th (undeflated) old singular value.</p>
<i>diffr</i>	<p>REAL for <code>slals0 / clals0</code> DOUBLE PRECISION for <code>dlals0 / zlals0</code> Array, DIMENSION (<i>ldgnum</i>, 2). On entry, <i>diffr</i>(<i>i</i>, 1) contains the distances between <i>i</i>-th updated (undeflated) singular value and the <i>i+1</i>-th (undeflated) old singular value. And <i>diffr</i>(<i>i</i>, 2) is the normalizing factor for the <i>i</i>-th right singular vector.</p>
<i>z</i>	<p>REAL for <code>slals0 / clals0</code> DOUBLE PRECISION for <code>dlals0 / zlals0</code> Array, DIMENSION (<i>k</i>). Contains the components of the deflation-adjusted updating row vector.</p>
<i>k</i>	<p>INTEGER. Contains the dimension of the non-deflated matrix. This is the order of the related secular equation. $1 \leq k \leq n$.</p>

<i>c</i>	REAL for <code>slals0/clals0</code> DOUBLE PRECISION for <code>dlals0/zlals0</code> Contains garbage if <code>sqre</code> = 0 and the <i>c</i> value of a Givens rotation related to the right null space if <code>sqre</code> = 1.
<i>s</i>	REAL for <code>slals0/clals0</code> DOUBLE PRECISION for <code>dlals0/zlals0</code> Contains garbage if <code>sqre</code> = 0 and the <i>s</i> value of a Givens rotation related to the right null space if <code>sqre</code> = 1.
<i>work</i>	REAL for <code>slals0</code> DOUBLE PRECISION for <code>dlals0</code> Workspace array, DIMENSION (<i>k</i>). Used with real flavors only.
<i>rwork</i>	REAL for <code>clals0</code> DOUBLE PRECISION for <code>zlals0</code> Workspace array, DIMENSION ($k*(1+nrhs) + 2*nrhs$). Used with complex flavors only.

Output Parameters

<i>b</i>	On exit, contains the solution <i>X</i> in rows 1 through <i>n</i> .
<i>info</i>	INTEGER. If <i>info</i> = 0: successful exit. If <i>info</i> = <i>-i</i> < 0, the <i>i</i> -th argument had an illegal value.

?lalsa

Computes the SVD of the coefficient matrix in compact form. Used by

?gelsd.

```

call slalsa (  icompq, smlsiz, n, nrhs, b, ldb, bx, ldbx,
              u, ldu, vt, k, difl, difr, z, poles, givptr,
              givcol, ldgcol, perm, givnum, c, s, work,
              iwork, info )
call dlalsa (  icompq, smlsiz, n, nrhs, b, ldb, bx, ldbx,
              u, ldu, vt, k, difl, difr, z, poles, givptr,
              givcol, ldgcol, perm, givnum, c, s, work,
              iwork, info )
call clalsa (  icompq, smlsiz, n, nrhs, b, ldb, bx, ldbx,
              u, ldu, vt, k, difl, difr, z, poles, givptr,
              givcol, ldgcol, perm, givnum, c, s, rwork,
              iwork, info )
call zlalsa (  icompq, smlsiz, n, nrhs, b, ldb, bx, ldbx,
              u, ldu, vt, k, difl, difr, z, poles, givptr,
              givcol, ldgcol, perm, givnum, c, s, rwork,
              iwork, info )

```

Discussion

The routine is an intermediate step in solving the least squares problem by computing the SVD of the coefficient matrix in compact form. The singular vectors are computed as products of simple orthogonal matrices.

If `icompq = 0`, `?lalsa` applies the inverse of the left singular vector matrix of an upper bidiagonal matrix to the right hand side; and if `icompq = 1`, the routine applies the right singular vector matrix to the right hand side. The singular vector matrices were generated in the compact form by `?lalsa`.

Input Parameters

<i>icompq</i>	INTEGER . Specifies whether the left or the right singular vector matrix is involved. If <i>icompq</i> = 0: left singular vector matrix is used If <i>icompq</i> = 1: right singular vector matrix is used.
<i>smlsiz</i>	INTEGER . The maximum size of the subproblems at the bottom of the computation tree.
<i>n</i>	INTEGER . The row and column dimensions of the upper bidiagonal matrix.
<i>nrhs</i>	INTEGER . The number of columns of <i>b</i> and <i>bx</i> . Must be at least 1.
<i>b</i>	REAL for <i>slalsa</i> DOUBLE PRECISION for <i>dlalsa</i> COMPLEX for <i>clalsa</i> COMPLEX*16 for <i>zlalsa</i> Array, DIMENSION (<i>ldb</i> , <i>nrhs</i>). Contains the right hand sides of the least squares problem in rows 1 through <i>m</i> .
<i>ldb</i>	INTEGER . The leading dimension of <i>b</i> in the calling subprogram. Must be at least $\max(1, \max(m, n))$.
<i>ldb_x</i>	INTEGER . The leading dimension of the output array <i>bx</i> .
<i>u</i>	REAL for <i>slalsa/clalsa</i> DOUBLE PRECISION for <i>dlalsa/zlalsa</i> Array, DIMENSION (<i>ldu</i> , <i>smlsiz</i>). On entry, <i>u</i> contains the left singular vector matrices of all subproblems at the bottom level.
<i>ldu</i>	INTEGER , <i>ldu</i> $\geq n$. The leading dimension of arrays <i>u</i> , <i>vt</i> , <i>difl</i> , <i>difr</i> , <i>poles</i> , <i>givnum</i> , and <i>z</i> .
<i>vt</i>	REAL for <i>slalsa/clalsa</i> DOUBLE PRECISION for <i>dlalsa/zlalsa</i> Array, DIMENSION (<i>ldu</i> , <i>smlsiz</i> + 1). On entry, contains the right singular vector matrices of all subproblems at the bottom level.
<i>k</i>	INTEGER array, DIMENSION (<i>n</i>).

<i>difl</i>	<p>REAL for <i>slalsa/clalsa</i> DOUBLE PRECISION for <i>dlalsa/zlalsa</i> Array, DIMENSION (<i>ldu, nlvl</i>), where <i>nlvl</i> = $\text{int}(\log_2(n/(smlsiz+1))) + 1$.</p>
<i>difr</i>	<p>REAL for <i>slalsa/clalsa</i> DOUBLE PRECISION for <i>dlalsa/zlalsa</i> Array, DIMENSION (<i>ldu, 2*nlvl</i>). On entry, <i>difl</i>(*, <i>i</i>) and <i>difr</i>(*, <i>2i-1</i>) record distances between singular values on the <i>i</i>-th level and singular values on the (<i>i-1</i>)-th level, and <i>difr</i>(*, <i>2i</i>) record the normalizing factors of the right singular vectors matrices of subproblems on <i>i</i>-th level.</p>
<i>z</i>	<p>REAL for <i>slalsa/clalsa</i> DOUBLE PRECISION for <i>dlalsa/zlalsa</i> Array, DIMENSION (<i>ldu, nlvl</i>). On entry, <i>z</i>(1, <i>i</i>) contains the components of the deflation- adjusted updating the row vector for subproblems on the <i>i</i>-th level.</p>
<i>poles</i>	<p>REAL for <i>slalsa/clalsa</i> DOUBLE PRECISION for <i>dlalsa/zlalsa</i> Array, DIMENSION (<i>ldu, 2*nlvl</i>). On entry, <i>poles</i>(*, <i>2i-1: 2i</i>) contains the new and old singular values involved in the secular equations on the <i>i</i>-th level.</p>
<i>givptr</i>	<p>INTEGER. Array, DIMENSION (<i>n</i>). On entry, <i>givptr</i>(<i>i</i>) records the number of Givens rotations performed on the <i>i</i>-th problem on the computation tree.</p>
<i>givcol</i>	<p>INTEGER. Array, DIMENSION (<i>ldgcol, 2*nlvl</i>). On entry, for each <i>i</i>, <i>givcol</i>(*, <i>2i-1: 2i</i>) records the locations of Givens rotations performed on the <i>i</i>-th level on the computation tree.</p>

<i>ldgcol</i>	INTEGER, $ldgcol \geq n$. The leading dimension of arrays <i>givcol</i> and <i>perm</i> .
<i>perm</i>	INTEGER. Array, DIMENSION (<i>ldgcol</i> , <i>nlvl</i>). On entry, <i>perm</i> (<i>i</i> , <i>i</i>) records permutations done on the <i>i</i> -th level of the computation tree.
<i>givnum</i>	REAL for <i>slalsa/clalsa</i> DOUBLE PRECISION for <i>dlalsa/zlalsa</i> Array, DIMENSION (<i>ldu</i> , $2 * nlvl$). On entry, <i>givnum</i> (<i>i</i> , $2i-1 : 2i$) records the <i>c</i> and <i>s</i> values of Givens rotations performed on the <i>i</i> -th level on the computation tree.
<i>c</i>	REAL for <i>slalsa/clalsa</i> DOUBLE PRECISION for <i>dlalsa/zlalsa</i> Array, DIMENSION (<i>n</i>). On entry, if the <i>i</i> -th subproblem is not square, <i>c</i> (<i>i</i>) contains the <i>c</i> value of a Givens rotation related to the right null space of the <i>i</i> -th subproblem.
<i>s</i>	REAL for <i>slalsa/clalsa</i> DOUBLE PRECISION for <i>dlalsa/zlalsa</i> Array, DIMENSION (<i>n</i>). On entry, if the <i>i</i> -th subproblem is not square, <i>s</i> (<i>i</i>) contains the <i>s</i> -value of a Givens rotation related to the right null space of the <i>i</i> -th subproblem.
<i>work</i>	REAL for <i>slalsa</i> DOUBLE PRECISION for <i>dlalsa</i> Workspace array, DIMENSION at least (<i>n</i>). Used with real flavors only.
<i>rwork</i>	REAL for <i>clalsa</i> DOUBLE PRECISION for <i>zlalsa</i> Workspace array, DIMENSION at least $\max(n, (smlsz+1) * nrhs * 3)$. Used with complex flavors only.
<i>iwork</i>	INTEGER. Workspace array, DIMENSION at least ($3n$).

Output Parameters

<i>b</i>	On exit, contains the solution X in rows 1 through n .
<i>bx</i>	REAL for <code>slalsa</code> DOUBLE PRECISION for <code>dlalsa</code> COMPLEX for <code>clalsa</code> COMPLEX*16 for <code>zlalsa</code> Array, DIMENSION (<i>ldb</i> , <i>nrhs</i>). On exit, the result of applying the left or right singular vector matrix to <i>b</i> .
<i>info</i>	INTEGER. If <i>info</i> = 0: successful exit If <i>info</i> = - <i>i</i> < 0, the <i>i</i> -th argument had an illegal value.

?lalsd

Uses the singular value decomposition of A to solve the least squares problem.

```
call slalsd ( uplo, smlsiz, n, nrhs, d, e, b, ldb,
             rcond, rank, work, iwork, info )
call dlalsd ( uplo, smlsiz, n, nrhs, d, e, b, ldb,
             rcond, rank, work, iwork, info )
call clalsd ( uplo, smlsiz, n, nrhs, d, e, b, ldb,
             rcond, rank, work, rwork, iwork, info )
call zlalsd ( uplo, smlsiz, n, nrhs, d, e, b, ldb,
             rcond, rank, work, rwork, iwork, info )
```

Discussion

The routine uses the singular value decomposition of A to solve the least squares problem of finding X to minimize the Euclidean norm of each column of $AX-B$, where A is n -by- n upper bidiagonal, and X and B are n -by- $nrhs$. The solution X overwrites B .

The singular values of A smaller than $rcond$ times the largest singular value are treated as zero in solving the least squares problem; in this case a minimum norm solution is returned. The actual singular values are returned in d in ascending order.

This code makes very mild assumptions about floating point arithmetic. It will work on machines with a guard digit in add/subtract, or on those binary machines without guard digits which subtract like the Cray XMP, Cray YMP, Cray C 90, or Cray 2.

It could conceivably fail on hexadecimal or decimal machines without guard digits, but we know of none.

Input Parameters

uplo CHARACTER*1.
 If *uplo* = 'U', d and e define an upper bidiagonal matrix.
 If *uplo* = 'L', d and e define a lower bidiagonal matrix.

smlsiz INTEGER. The maximum size of the subproblems at the bottom of the computation tree.

n INTEGER. The dimension of the bidiagonal matrix.
 $n \geq 0$.

nrhs INTEGER. The number of columns of B . Must be at least 1.

d REAL for `slalsd/clalsd`
 DOUBLE PRECISION for `dlalsd/zlalsd`
 Array, DIMENSION (n). On entry, d contains the main diagonal of the bidiagonal matrix.

e REAL for `slalsd/clalsd`
 DOUBLE PRECISION for `dlalsd/zlalsd`
 Array, DIMENSION ($n-1$). Contains the super-diagonal entries of the bidiagonal matrix. On exit, e is destroyed.

b REAL for `slalsd`
 DOUBLE PRECISION for `dlalsd`
 COMPLEX for `clalsd`
 COMPLEX*16 for `zlalsd`

Array, **DIMENSION** (*ldb,nrhs*). On input, *b* contains the right hand sides of the least squares problem. On output, *b* contains the solution *X*.

ldb **INTEGER**. The leading dimension of *b* in the calling subprogram. Must be at least $\max(1,n)$.

rcond **REAL** for **slalsd/clalsd**
DOUBLE PRECISION for **dlalsd/zlalsd**
The singular values of *A* less than or equal to *rcond* times the largest singular value are treated as zero in solving the least squares problem.
If *rcond* is negative, machine precision is used instead. For example, if $\text{diag}(S)*X=B$ were the least squares problem, where $\text{diag}(S)$ is a diagonal matrix of singular values, the solution would be $X(i) = B(i) / S(i)$ if $S(i)$ is greater than $rcond * \max(S)$, and $X(i) = 0$ if $S(i)$ is less than or equal to $rcond * \max(S)$.

rank **INTEGER**. The number of singular values of *A* greater than *rcond* times the largest singular value.

work **REAL** for **slalsd**
DOUBLE PRECISION for **dlalsd**
COMPLEX for **clalsd**
COMPLEX*16 for **zlalsd**
Workspace array.
DIMENSION for real flavors at least $(9n+2n*smlsiz+8n*nlvl+n*nrhs+(smlsiz+1)^2)$, where
 $nlvl = \max(0, \text{int}(\log_2(n / (smlsiz+1))) + 1)$.
DIMENSION for complex flavors at least $(n*nrhs)$.

rwork **REAL** for **clalsd**
DOUBLE PRECISION for **zlalsd**
Workspace array, used with complex flavors only.
DIMENSION at least $(9n + 2n*smlsiz + 8n*nlvl + 3*smlsiz*nrhs + (smlsiz+1)^2)$, where
 $nlvl = \max(0, \text{int}(\log_2(\min(m,n)/(smlsiz+1))) + 1)$.

iwork **INTEGER**.
 Workspace array, **DIMENSION** at least $(3n*nlvl + 11n)$.

Output Parameters

d On exit, if *info* = 0, *d* contains singular values of the bidiagonal matrix.

b On exit, *b* contains the solution *X*.

info **INTEGER**.
 If *info* = 0: successful exit.
 If *info* = -*i* < 0, the *i*-th argument had an illegal value.
 If *info* > 0: The algorithm failed to compute a singular value while working on the submatrix lying in rows and columns *info*/(*n*+1) through mod(*info*,*n*+1).

?lamch

Determines machine parameters for floating-point arithmetic.

```
val = slamch ( cmach )
val = dlamch ( cmach )
```

Discussion

The function ?lamch determines single precision and double precision machine parameters.

Input Parameters

cmach **CHARACTER*1**. Specifies the value to be returned by ?lamch:
 = 'E' or 'e', *val* = *eps*
 = 'S' or 's', *val* = *sfmin*
 = 'B' or 'b', *val* = *base*
 = 'P' or 'p', *val* = *eps*base*

```

= 'N' or 'n', val = t
= 'R' or 'r', val = rnd
= 'M' or 'm', val = emin
= 'U' or 'u', val = rmin
= 'L' or 'l', val = emax
= 'O' or 'o', val = rmax
where
eps = relative machine precision;
sfmin = safe minimum, such that 1/sfmin does not
overflow;
base = base of the machine;
prec = eps*base;
t = number of (base) digits in the mantissa;
rnd = 1.0 when rounding occurs in addition, 0.0
otherwise;
emin = minimum exponent before (gradual) underflow;
rmin = underflow_threshold - base**(emin-1);
emax = largest exponent before overflow;
rmax = overflow_threshold - (base**emax)*(1-eps).

```

Output Parameters

```

val          REAL for slamch
             DOUBLE PRECISION for dlamch
             Value returned by the function.

```

?lamc1

Called from ?lamc2.
Determines machine parameters given
by beta, t, rnd, ieee1.

```

call slamc1 ( beta, t, rnd, ieee1 )
call dlamc1 ( beta, t, rnd, ieee1 )

```

Discussion

The routine `?lamc1` determines machine parameters given by `beta`, `t`, `rnd`, `ieee1`.

Output Parameters

`beta` **INTEGER.** The base of the machine.

`t` **INTEGER.** The number of (`beta`) digits in the mantissa.

`rnd` **LOGICAL.**
Specifies whether proper rounding (`rnd = .TRUE.`) or chopping (`rnd = .FALSE.`) occurs in addition. This may not be a reliable guide to the way in which the machine performs its arithmetic.

`ieee1` **LOGICAL.**
Specifies whether rounding appears to be done in the `ieee` 'round to nearest' style.

?lamc2

Used by ?lamch.

Determines machine parameters specified in its arguments list.

```
call slamc2 ( beta, t, rnd, eps, emin, rmin, emax, rmax )
call dlamc2 ( beta, t, rnd, eps, emin, rmin, emax, rmax )
```

Discussion

The routine `?lamc2` determines machine parameters specified in its arguments list.

Output Parameters

`beta` **INTEGER.** The base of the machine.

`t` **INTEGER.** The number of (`beta`) digits in the mantissa.

<i>rnd</i>	LOGICAL. Specifies whether proper rounding (<i>rnd</i> = <code>.TRUE.</code>) or chopping (<i>rnd</i> = <code>.FALSE.</code>) occurs in addition. This may not be a reliable guide to the way in which the machine performs its arithmetic.
<i>eps</i>	REAL for <code>slamc2</code> DOUBLE PRECISION for <code>dlamc2</code> The smallest positive number such that $fl(1.0 - eps) < 1.0$, where <i>fl</i> denotes the computed value.
<i>emin</i>	INTEGER. The minimum exponent before (gradual) underflow occurs.
<i>rmin</i>	REAL for <code>slamc2</code> DOUBLE PRECISION for <code>dlamc2</code> The smallest normalized number for the machine, given by $base^{emin-1}$, where <i>base</i> is the floating point value of <i>beta</i> .
<i>emax</i>	INTEGER. The maximum exponent before overflow occurs.
<i>rmax</i>	REAL for <code>slamc2</code> DOUBLE PRECISION for <code>dlamc2</code> The largest positive number for the machine, given by $base^{emax}(1 - eps)$, where <i>base</i> is the floating point value of <i>beta</i> .

?lamc3

Called from ?lamc1-?lamc5. Intended to force *a* and *b* to be stored prior to doing the addition of *a* and *b*.

```
val = slamc3 (a, b)
val = dlamc3 (a, b)
```


Discussion

The routine is intended to force *a* and *b* to be stored prior to doing the addition of *a* and *b*, for use in situations where optimizers might hold one of these in a register.

Input Parameters

a,b REAL for `s1amc3`
 DOUBLE PRECISION for `d1amc3`
 The values *a* and *b*.

Output Parameters

val REAL for `s1amc3`
 DOUBLE PRECISION for `d1amc3`
 The result of adding values *a* and *b*.

?lamc4

This is a service routine for ?lamc2.

```
call s1amc4 (emin, start, base)
call d1amc4 (emin, start, base)
```

Discussion

This is a service routine for ?lamc2.

Input Parameters

start REAL for `s1amc4`
 DOUBLE PRECISION for `d1amc4`
 The starting point for determining *emin*.

base INTEGER. The base of the machine.

Output Parameters

emin **INTEGER**. The minimum exponent before (gradual) underflow, computed by setting *a* = *start* and dividing by *base* until the previous *a* can not be recovered.

?lamc5

Called from ?lamc2.

Attempts to compute the largest machine floating-point number, without overflow.

```
call slamc5 ( beta, p, emin, ieee, emax, rmax )
call dlamc5 ( beta, p, emin, ieee, emax, rmax )
```

Discussion

The routine ?lamc5 attempts to compute *rmax*, the largest machine floating-point number, without overflow. It assumes that $emax + \text{abs}(emin)$ sum approximately to a power of 2. It will fail on machines where this assumption does not hold, for example, the Cyber 205 (*emin* = -28625, *emax* = 28718). It will also fail if the value supplied for *emin* is too large (that is, too close to zero), probably with overflow.

Input Parameters

beta **INTEGER**. The base of floating-point arithmetic.

p **INTEGER**. The number of base *beta* digits in the mantissa of a floating-point value.

emin **INTEGER**. The minimum exponent before (gradual) underflow.

ieee **LOGICAL**. A logical flag specifying whether or not the arithmetic system is thought to comply with the IEEE standard.

Output Parameters.

emax **INTEGER**. The largest exponent before overflow.

rmax **REAL** for **slamc5**
DOUBLE PRECISION for **dlamc5**
The largest machine floating-point number.

?lamrg

Creates a permutation list to merge the entries of two independently sorted sets into a single set sorted in ascending order.

```
call slamrg ( n1, n2, a, strd1, strd2, index )
call dlamrg ( n1, n2, a, strd1, strd2, index )
```

Discussion

The routine creates a permutation list which will merge the elements of *a* (which is composed of two independently sorted sets) into a single set which is sorted in ascending order.

Input Parameters

n1, n2 **INTEGER**.
These arguments contain the respective lengths of the two sorted lists to be merged.

a **REAL** for **slamrg**
DOUBLE PRECISION for **dlamrg**.
Array, **DIMENSION** (*n1+n2*).
The first *n1* elements of *a* contain a list of numbers which are sorted in either ascending or descending order. Likewise for the final *n2* elements.

strd1, strd2 INTEGER.

These are the strides to be taken through the array *a*. Allowable strides are 1 and -1. They indicate whether a subset of *a* is sorted in ascending (*strdx* = 1) or descending (*strdx* = -1) order.

Output Parameters

index INTEGER.

Array, DIMENSION (*n1+n2*).
On exit, this array will contain a permutation such that if $b(i) = a(index(i))$ for $i=1, n1+n2$, then *b* will be sorted in ascending order.

?langb

Returns the value of the 1-norm, Frobenius norm, infinity-norm, or the largest absolute value of any element of general band matrix.

```

val = slangb ( norm, n, kl, ku, ab, ldab, work )
val = dlangb ( norm, n, kl, ku, ab, ldab, work )
val = clangb ( norm, n, kl, ku, ab, ldab, work )
val = zlangb ( norm, n, kl, ku, ab, ldab, work )

```

Discussion

The function returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of an *n*-by-*n* band matrix *A*, with *kl* sub-diagonals and *ku* super-diagonals.

The value *val* returned by the function is:

```

val = max(abs(Aij)), if norm = 'M' or 'm'
      = norm1(A),    if norm = '1' or 'O' or 'o'
      = normI(A),   if norm = 'I' or 'i'

```

= normF(A), if *norm* = 'F', 'f', 'E' or 'e'

where norm1 denotes the 1-norm of a matrix (maximum column sum), normI denotes the infinity norm of a matrix (maximum row sum) and normF denotes the Frobenius norm of a matrix (square root of sum of squares). Note that $\max(\text{abs}(A_{ij}))$ is not a matrix norm.

Input Parameters

norm CHARACTER*1. Specifies the value to be returned by the routine as described above.

n INTEGER. The order of the matrix A.
 $n \geq 0$. When $n = 0$, *langb* is set to zero.

kl INTEGER. The number of sub-diagonals of the matrix A.
 $kl \geq 0$.

ku INTEGER. The number of super-diagonals of the matrix A.
 $ku \geq 0$.

ab REAL for slangb
DOUBLE PRECISION for dlangb
COMPLEX for clangb
COMPLEX*16 for zlangb
Array, DIMENSION (*ldab*,*n*). The band matrix A, stored in rows 1 to $kl+ku+1$. The *j*-th column of A is stored in the *j*-th column of the array *ab* as follows:
 $ab(ku+1+i-j, j) = a(i, j)$
for $\max(1, j-ku) \leq i \leq \min(n, j+kl)$.

ldab INTEGER. The leading dimension of the array *ab*.
 $ldab \geq kl+ku+1$.

work REAL for slangb/clangb
DOUBLE PRECISION for dlangb/zlangb
Workspace array, DIMENSION (*lwork*), where $lwork \geq n$ when *norm* = 'I'; otherwise, *work* is not referenced.

Output Parameters

val REAL for `slangb/clangb`
 DOUBLE PRECISION for `dlangb/zlangb`
 Value returned by the function.

?lange

Returns the value of the 1-norm, Frobenius norm, infinity-norm, or the largest absolute value of any element of a general rectangular matrix.

```
val = slange ( norm, m, n, a, lda, work )
val = dlange ( norm, m, n, a, lda, work )
val = clange ( norm, m, n, a, lda, work )
val = zlange ( norm, m, n, a, lda, work )
```

Discussion

The function `?lange` returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a real/complex matrix *A*.

The value *val* returned by the function is:

```
val = max(abs(Aij)), if norm = 'M' or 'm'
      = norm1(A),    if norm = '1' or 'O' or 'o'
      = normI(A),   if norm = 'I' or 'i'
      = normF(A),   if norm = 'F', 'f', 'E' or 'e'
```

where `norm1` denotes the 1-norm of a matrix (maximum column sum), `normI` denotes the infinity norm of a matrix (maximum row sum) and `normF` denotes the Frobenius norm of a matrix (square root of sum of squares). Note that `max(abs(Aij))` is not a matrix norm.

Input Parameters

<i>norm</i>	CHARACTER*1. Specifies the value to be returned in ?lange as described above.
<i>m</i>	INTEGER. The number of rows of the matrix A. $m \geq 0$. When $m = 0$, ?lange is set to zero.
<i>n</i>	INTEGER. The number of columns of the matrix A. $n \geq 0$. When $n = 0$, ?lange is set to zero.
<i>a</i>	REAL for slange DOUBLE PRECISION for dlange COMPLEX for clange COMPLEX*16 for zlange Array, DIMENSION (<i>lda</i> , <i>n</i>). The <i>m</i> -by- <i>n</i> matrix A.
<i>lda</i>	INTEGER. The leading dimension of the array <i>a</i> . $lda \geq \max(m,1)$.
<i>work</i>	REAL for slange and clange. DOUBLE PRECISION for dlange and zlange. Workspace array, DIMENSION (<i>lwork</i>), where $lwork \geq m$ when <i>norm</i> = 'I'; otherwise, <i>work</i> is not referenced.

Output Parameters

<i>val</i>	REAL for slange/clange DOUBLE PRECISION for dlange/zlange Value returned by the function.
------------	---

?langt

Returns the value of the 1-norm, Frobenius norm, infinity-norm, or the largest absolute value of any element of a general tridiagonal matrix.

```
val = slangt ( norm, n, dl, d, du )
```

```

val = dlangt ( norm, n, dl, d, du )
val = clangt ( norm, n, dl, d, du )
val = zlangt ( norm, n, dl, d, du )

```

Discussion

The routine returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a real/complex tridiagonal matrix *A*.

The value *val* returned by the function is:

```

val = max(abs(Aij)), if norm = 'M' or 'm'
      = norm1(A),    if norm = '1' or 'O' or 'o'
      = normI(A),   if norm = 'I' or 'i'
      = normF(A),   if norm = 'F', 'f', 'E' or 'e'

```

where `norm1` denotes the 1-norm of a matrix (maximum column sum), `normI` denotes the infinity norm of a matrix (maximum row sum) and `normF` denotes the Frobenius norm of a matrix (square root of sum of squares). Note that `max(abs(Aij))` is not a matrix norm.

Input Parameters

<i>norm</i>	CHARACTER*1. Specifies the value to be returned in <i>?langt</i> as described above.
<i>n</i>	INTEGER. The order of the matrix <i>A</i> . $n \geq 0$. When $n = 0$, <i>?langt</i> is set to zero.
<i>dl, d, du</i>	REAL for <i>slangt</i> DOUBLE PRECISION for <i>dlangt</i> COMPLEX for <i>clangt</i> COMPLEX*16 for <i>zlangt</i> Arrays: <i>dl</i> ($n-1$), <i>d</i> (n), <i>du</i> ($n-1$). The array <i>dl</i> contains the ($n-1$) sub-diagonal elements of <i>A</i> . The array <i>d</i> contains the diagonal elements of <i>A</i> . The array <i>du</i> contains the ($n-1$) super-diagonal elements of <i>A</i> .

Output Parameters

`val` REAL for `slangt/clangt`
 DOUBLE PRECISION for `dlangt/zlangt`
 Value returned by the function.

?lanhs

Returns the value of the 1-norm, Frobenius norm, infinity-norm, or the largest absolute value of any element of an upper Hessenberg matrix.

```
val = slanhs ( norm, n, a, lda, work )
val = dlanhs ( norm, n, a, lda, work )
val = clanhs ( norm, n, a, lda, work )
val = zlanhs ( norm, n, a, lda, work )
```

Discussion

The function `?lanhs` returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a Hessenberg matrix A .

The value `val` returned by the function is:

```
val = max(abs(Aij)), if norm = 'M' or 'm'
      = norm1(A),    if norm = '1' or 'O' or 'o'
      = normI(A),   if norm = 'I' or 'i'
      = normF(A),   if norm = 'F', 'f', 'E' or 'e'
```

where `norm1` denotes the 1-norm of a matrix (maximum column sum), `normI` denotes the infinity norm of a matrix (maximum row sum) and `normF` denotes the Frobenius norm of a matrix (square root of sum of squares). Note that `max(abs(Aij))` is not a matrix norm.

Input Parameters

<i>norm</i>	CHARACTER*1. Specifies the value to be returned in ?lanhs as described above.
<i>n</i>	INTEGER. The order of the matrix A. $n \geq 0$. When $n = 0$, ?lanhs is set to zero.
<i>a</i>	REAL for slanhs DOUBLE PRECISION for dlanhs COMPLEX for clanhs COMPLEX*16 for zlanhs Array, DIMENSION (<i>lda</i> , <i>n</i>). The <i>n</i> -by- <i>n</i> upper Hessenberg matrix A; the part of A below the first sub-diagonal is not referenced.
<i>lda</i>	INTEGER. The leading dimension of the array <i>a</i> . $lda \geq \max(n,1)$.
<i>work</i>	REAL for slanhs and clanhs. DOUBLE PRECISION for dlange and zlange. Workspace array, DIMENSION (<i>lwork</i>), where $lwork \geq n$ when <i>norm</i> = 'I'; otherwise, <i>work</i> is not referenced.

Output Parameters

<i>val</i>	REAL for slanhs/clanhs DOUBLE PRECISION for dlanhs/zlanhs Value returned by the function.
------------	---

?lansb

Returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a symmetric band matrix.

```
val = slansb ( norm, uplo, n, k, ab, ldab, work )
val = dlansb ( norm, uplo, n, k, ab, ldab, work )
```

```

val = clansb ( norm, uplo, n, k, ab, ldab, work )
val = zlansb ( norm, uplo, n, k, ab, ldab, work )

```

Discussion

The function `?lansb` returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of an n -by- n real/complex symmetric band matrix A , with k super-diagonals.

The value `val` returned by the function is:

```

val = max(abs(Aij)), if norm = 'M' or 'm'
      = norm1(A),    if norm = '1' or 'O' or 'o'
      = normI(A),   if norm = 'I' or 'i'
      = normF(A),   if norm = 'F', 'f', 'E' or 'e'

```

where `norm1` denotes the 1-norm of a matrix (maximum column sum), `normI` denotes the infinity norm of a matrix (maximum row sum) and `normF` denotes the Frobenius norm of a matrix (square root of sum of squares). Note that `max(abs(Aij))` is not a matrix norm.

Input Parameters

`norm` CHARACTER*1. Specifies the value to be returned in `?lansb` as described above.

`uplo` CHARACTER*1. Specifies whether the upper or lower triangular part of the band matrix A is supplied.
If `uplo` = 'U': upper triangular part is supplied;
If `uplo` = 'L': lower triangular part is supplied.

`n` INTEGER. The order of the matrix A . $n \geq 0$.
When $n = 0$, `?lansb` is set to zero.

`k` INTEGER. The number of super-diagonals or sub-diagonals of the band matrix A . $k \geq 0$.

`ab` REAL for `slansb`
DOUBLE PRECISION for `dlansb`
COMPLEX for `clansb`
COMPLEX*16 for `zlansb`
Array, DIMENSION (`ldab`, n). The upper or lower triangle of the symmetric band matrix A , stored in the

first $k+1$ rows of ab . The j -th column of A is stored in the j -th column of the array ab as follows:

if $uplo = 'U'$, $ab(k+1+i-j, j) = a(i, j)$

for $\max(1, j-k) \leq i \leq j$;

if $uplo = 'L'$, $ab(1+i-j, j) = a(i, j)$ for $j \leq \min(n, j+k)$.

ldab INTEGER. The leading dimension of the array ab .
 $ldab \geq k+1$.

work REAL for `slansb` and `clansb`.
 DOUBLE PRECISION for `dlansb` and `zlansb`.
 Workspace array, DIMENSION ($lwork$), where
 $lwork \geq n$ when $norm = 'I'$ or $'1'$ or $'O'$; otherwise, $work$ is not referenced.

Output Parameters

val REAL for `slansb/clansb`
 DOUBLE PRECISION for `dlansb/zlansb`
 Value returned by the function.

?lanhb

Returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a Hermitian band matrix.

`val = clanhb (norm, uplo, n, k, ab, ldab, work)`

`val = zlanhb (norm, uplo, n, k, ab, ldab, work)`

Discussion

The routine returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of an n -by- n Hermitian band matrix A , with k super-diagonals.

The value `val` returned by the function is:

$val = \max(\text{abs}(A_{ij}))$, if $norm = 'M'$ or $'m'$
 $= \text{norm1}(A)$, if $norm = '1'$ or $'O'$ or $'o'$
 $= \text{normI}(A)$, if $norm = 'I'$ or $'i'$
 $= \text{normF}(A)$, if $norm = 'F', 'f', 'E'$ or $'e'$

where norm1 denotes the 1-norm of a matrix (maximum column sum), normI denotes the infinity norm of a matrix (maximum row sum) and normF denotes the Frobenius norm of a matrix (square root of sum of squares). Note that $\max(\text{abs}(A_{ij}))$ is not a matrix norm.

Input Parameters

norm CHARACTER*1. Specifies the value to be returned in *?lanhb* as described above.

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the band matrix *A* is supplied.
 If *uplo* = 'U': upper triangular part is supplied;
 If *uplo* = 'L': lower triangular part is supplied.

n INTEGER. The order of the matrix *A*. $n \geq 0$. When $n = 0$, *?lanhb* is set to zero.

k INTEGER. The number of super-diagonals or sub-diagonals of the band matrix *A*. $k \geq 0$.

ab COMPLEX for *clanhb*.
 COMPLEX*16 for *zlanhb*.
 Array, DIMENSION (*ldab*,*n*). The upper or lower triangle of the Hermitian band matrix *A*, stored in the first $k+1$ rows of *ab*. The *j*-th column of *A* is stored in the *j*-th column of the array *ab* as follows:
 if *uplo* = 'U', $ab(k+1+i-j, j) = a(i, j)$
 for $\max(1, j-k) \leq i \leq j$;
 if *uplo* = 'L', $ab(1+i-j, j) = a(i, j)$ for $j \leq \min(n, j+k)$.
 Note that the imaginary parts of the diagonal elements need not be set and are assumed to be zero.

ldab INTEGER. The leading dimension of the array *ab*.
 $ldab \geq k+1$.

work REAL for `clanhb`.
 DOUBLE PRECISION for `zlanhb`.
 Workspace array, DIMENSION (*lwork*), where
lwork ≥ *n* when *norm* = 'I' or 'l' or 'O'; otherwise, *work*
 is not referenced.

Output Parameters

val REAL for `slanhb/clanhb`
 DOUBLE PRECISION for `dlanhb/zlanhb`
 Value returned by the function.

?lansp

Returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a symmetric matrix supplied in packed form.

```
val = slansp ( norm, uplo, n, ap, work )
val = dlansp ( norm, uplo, n, ap, work )
val = clansp ( norm, uplo, n, ap, work )
val = zlansp ( norm, uplo, n, ap, work )
```

Discussion

The function `?lansp` returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a real/complex symmetric matrix *A*, supplied in packed form.

The value *val* returned by the function is:

$$\begin{aligned}
 \text{val} &= \max(\text{abs}(A_{ij})), & \text{if } \text{norm} = \text{'M'} \text{ or 'm'} \\
 &= \text{norm1}(A), & \text{if } \text{norm} = \text{'1'} \text{ or 'O'} \text{ or 'o'} \\
 &= \text{normI}(A), & \text{if } \text{norm} = \text{'I'} \text{ or 'i'}
 \end{aligned}$$

= $\text{normF}(A)$, if *norm* = 'F', 'f', 'E' or 'e'

where norm1 denotes the 1-norm of a matrix (maximum column sum), normI denotes the infinity norm of a matrix (maximum row sum) and normF denotes the Frobenius norm of a matrix (square root of sum of squares). Note that $\max(\text{abs}(A_{ij}))$ is not a matrix norm.

Input Parameters

norm CHARACTER*1. Specifies the value to be returned in *?lansp* as described above.

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the symmetric matrix *A* is supplied.
If *uplo* = 'U': Upper triangular part of *A* is supplied
If *uplo* = 'L': Lower triangular part of *A* is supplied.

n INTEGER. The order of the matrix *A*. $n \geq 0$. When $n = 0$, *?lansp* is set to zero.

ap REAL for *slansp*
DOUBLE PRECISION for *dlansp*
COMPLEX for *clansp*
COMPLEX*16 for *zlansp*
Array, DIMENSION $(n(n+1)/2)$. The upper or lower triangle of the symmetric matrix *A*, packed columnwise in a linear array. The *j*-th column of *A* is stored in the array *ap* as follows:
if *uplo* = 'U', $ap(i + (j-1)j/2) = A(i, j)$ for $1 \leq i \leq j$;
if *uplo* = 'L', $ap(i + (j-1)(2n-j)/2) = A(i, j)$ for $j \leq i \leq n$.

work REAL for *slansp* and *clansp*.
DOUBLE PRECISION for *dlansp* and *zlansp*.
Workspace array, DIMENSION (*lwork*), where $lwork \geq n$ when *norm* = 'I' or 'l' or 'O'; otherwise, *work* is not referenced.

Output Parameters

val REAL for *slansp/clansp*
DOUBLE PRECISION for *dlansp/zlansp*
Value returned by the function.

?lanhp

Returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a complex Hermitian matrix supplied in packed form.

```
val = clanhp ( norm, uplo, n, ap, work )
val = zlanhp ( norm, uplo, n, ap, work )
```

Discussion

The function `?lanhp` returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a complex Hermitian matrix A , supplied in packed form.

The value `val` returned by the function is:

```
val = max(abs(Aij)), if norm = 'M' or 'm'
      = norm1(A),    if norm = '1' or 'O' or 'o'
      = normI(A),   if norm = 'I' or 'i'
      = normF(A),   if norm = 'F', 'f', 'E' or 'e'
```

where `norm1` denotes the 1-norm of a matrix (maximum column sum), `normI` denotes the infinity norm of a matrix (maximum row sum) and `normF` denotes the Frobenius norm of a matrix (square root of sum of squares). Note that `max(abs(Aij))` is not a matrix norm.

Input Parameters

`norm` CHARACTER*1. Specifies the value to be returned in `?lanhp` as described above.

`uplo` CHARACTER*1. Specifies whether the upper or lower triangular part of the Hermitian matrix A is supplied. If `uplo = 'U'`: Upper triangular part of A is supplied. If `uplo = 'L'`: Lower triangular part of A is supplied.

n **INTEGER**. The order of the matrix *A*.
 $n \geq 0$. When $n = 0$, `?lanhp` is set to zero.

ap **COMPLEX** for `clanhp`.
COMPLEX*16 for `zlanhp`.
 Array, **DIMENSION** ($n(n+1)/2$). The upper or lower triangle of the Hermitian matrix *A*, packed columnwise in a linear array. The *j*-th column of *A* is stored in the array *ap* as follows:
 if *uplo* = 'U', $ap(i + (j-1)j/2) = A(i, j)$ for $1 \leq i \leq j$;
 if *uplo* = 'L', $ap(i + (j-1)(2n-j)/2) = A(i, j)$ for $j \leq i \leq n$.

work **REAL** for `clanhp`.
DOUBLE PRECISION for `zlanhp`.
 Workspace array, **DIMENSION** (*lwork*), where
 $lwork \geq n$ when *norm* = 'I' or '1' or 'O'; otherwise, *work* is not referenced.

Output Parameters

val **REAL** for `clanhp`.
DOUBLE PRECISION for `zlanhp`.
 Value returned by the function.

?lanst/?lanht

Returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a real symmetric or complex Hermitian tridiagonal matrix.

```
val = slanst ( norm, n, d, e )
val = dlanst ( norm, n, d, e )
val = clanht ( norm, n, d, e )
val = zlanht ( norm, n, d, e )
```

Discussion

The functions `?lanst/?lanht` return the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a real symmetric or a complex Hermitian tridiagonal matrix A .

The value `val` returned by the function is:

$$\begin{aligned}
 \text{val} &= \max(\text{abs}(A_{ij})), & \text{if } \text{norm} = \text{'M'} \text{ or 'm'} \\
 &= \text{norm1}(A), & \text{if } \text{norm} = \text{'1'} \text{ or 'O'} \text{ or 'o'} \\
 &= \text{normI}(A), & \text{if } \text{norm} = \text{'I'} \text{ or 'i'} \\
 &= \text{normF}(A), & \text{if } \text{norm} = \text{'F'}, \text{'f'}, \text{'E'} \text{ or 'e'}
 \end{aligned}$$

where `norm1` denotes the 1-norm of a matrix (maximum column sum), `normI` denotes the infinity norm of a matrix (maximum row sum) and `normF` denotes the Frobenius norm of a matrix (square root of sum of squares). Note that $\max(\text{abs}(A_{ij}))$ is not a matrix norm.

Input Parameters

<code>norm</code>	CHARACTER*1. Specifies the value to be returned in <code>?lanst/?lanht</code> as described above.
<code>n</code>	INTEGER. The order of the matrix A . $n \geq 0$. When $n = 0$, <code>?lanst/?lanht</code> is set to zero.
<code>d</code>	REAL for <code>slanst/clanht</code> DOUBLE PRECISION for <code>dlanst/zlanht</code> Array, DIMENSION (n). The diagonal elements of A .
<code>e</code>	REAL for <code>slanst</code> DOUBLE PRECISION for <code>dlanst</code> COMPLEX for <code>clanht</code> COMPLEX*16 for <code>zlanht</code> Array, DIMENSION ($n-1$). The ($n-1$) sub-diagonal or super-diagonal elements of A .

Output Parameters

<code>val</code>	REAL for <code>slanst/clanht</code> DOUBLE PRECISION for <code>dlanst/zlanht</code> Value returned by the function.
------------------	---

?lansy

Returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a real/complex symmetric matrix.

```
val = slansy ( norm, uplo, n, a, lda, work )
val = dlansy ( norm, uplo, n, a, lda, work )
val = clansy ( norm, uplo, n, a, lda, work )
val = zlansy ( norm, uplo, n, a, lda, work )
```

Discussion

The function `?lansy` returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a real/complex symmetric matrix A .

The value `val` returned by the function is:

```
val = max(abs(Aij)), if norm = 'M' or 'm'
      = norm1(A),    if norm = '1' or 'O' or 'o'
      = normI(A),   if norm = 'I' or 'i'
      = normF(A),   if norm = 'F', 'f', 'E' or 'e'
```

where `norm1` denotes the 1-norm of a matrix (maximum column sum), `normI` denotes the infinity norm of a matrix (maximum row sum) and `normF` denotes the Frobenius norm of a matrix (square root of sum of squares). Note that `max(abs(Aij))` is not a matrix norm.

Input Parameters

`norm` CHARACTER*1. Specifies the value to be returned in `?lansy` as described above.

<i>uplo</i>	<p>CHARACTER*1. Specifies whether the upper or lower triangular part of the symmetric matrix <i>A</i> is to be referenced.</p> <p>= 'U': Upper triangular part of <i>A</i> is referenced.</p> <p>= 'L': Lower triangular part of <i>A</i> is referenced</p>
<i>n</i>	<p>INTEGER. The order of the matrix <i>A</i>. $n \geq 0$. When $n = 0$, <i>?lansy</i> is set to zero.</p>
<i>a</i>	<p>REAL for <i>slansy</i></p> <p>DOUBLE PRECISION for <i>dlansy</i></p> <p>COMPLEX for <i>clansy</i></p> <p>COMPLEX*16 for <i>zlansy</i></p> <p>Array, DIMENSION (<i>lda</i>,<i>n</i>). The symmetric matrix <i>A</i>. If <i>uplo</i> = 'U', the leading <i>n</i>-by-<i>n</i> upper triangular part of <i>a</i> contains the upper triangular part of the matrix <i>A</i>, and the strictly lower triangular part of <i>a</i> is not referenced. If <i>uplo</i> = 'L', the leading <i>n</i>-by-<i>n</i> lower triangular part of <i>a</i> contains the lower triangular part of the matrix <i>A</i>, and the strictly upper triangular part of <i>a</i> is not referenced.</p>
<i>lda</i>	<p>INTEGER. The leading dimension of the array <i>a</i>.</p> <p>$lda \geq \max(n,1)$.</p>
<i>work</i>	<p>REAL for <i>slansy</i> and <i>clansy</i>.</p> <p>DOUBLE PRECISION for <i>dlansy</i> and <i>zlansy</i>.</p> <p>Workspace array, DIMENSION (<i>lwork</i>), where $lwork \geq n$ when <i>norm</i> = 'I' or '1' or 'O'; otherwise, <i>work</i> is not referenced.</p>

Output Parameters

<i>val</i>	<p>REAL for <i>slansy/clansy</i></p> <p>DOUBLE PRECISION for <i>dlansy/zlansy</i></p> <p>Value returned by the function.</p>
------------	--

?lanhe

Returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a complex Hermitian matrix.

```
val = clanhe ( norm, uplo, n, a, lda, work )  
val = zlanhe ( norm, uplo, n, a, lda, work )
```

Discussion

The function ?lanhe returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a complex Hermitian matrix A .

The value *val* returned by the function is:

```
val = max(abs(Aij)), if norm = 'M' or 'm'  
      = norm1(A),    if norm = '1' or 'O' or 'o'  
      = normI(A),   if norm = 'I' or 'i'  
      = normF(A),   if norm = 'F', 'f', 'E' or 'e'
```

where `norm1` denotes the 1-norm of a matrix (maximum column sum), `normI` denotes the infinity norm of a matrix (maximum row sum) and `normF` denotes the Frobenius norm of a matrix (square root of sum of squares). Note that `max(abs(Aij))` is not a matrix norm.

Input Parameters

norm CHARACTER*1. Specifies the value to be returned in ?lanhe as described above.

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the Hermitian matrix A is to be referenced.
= 'U': Upper triangular part of A is referenced.
= 'L': Lower triangular part of A is referenced

n **INTEGER**. The order of the matrix *A*.
 $n \geq 0$. When $n = 0$, `?lanhe` is set to zero.

a **COMPLEX** for `clanhe`.
COMPLEX*16 for `zlanhe`.
 Array, **DIMENSION** (*lda*,*n*). The Hermitian matrix *A*.
 If *uplo* = 'U', the leading *n*-by-*n* upper triangular part of *a* contains the upper triangular part of the matrix *A*, and the strictly lower triangular part of *a* is not referenced.
 If *uplo* = 'L', the leading *n*-by-*n* lower triangular part of *a* contains the lower triangular part of the matrix *A*, and the strictly upper triangular part of *a* is not referenced.

lda **INTEGER**. The leading dimension of the array *a*.
 $lda \geq \max(n,1)$.

work **REAL** for `clanhe`.
DOUBLE PRECISION for `zlanhe`.
 Workspace array, **DIMENSION** (*lwork*), where
 $lwork \geq n$ when *norm* = 'I' or '1' or 'O'; otherwise, *work* is not referenced.

Output Parameters

val **REAL** for `clanhe`.
DOUBLE PRECISION for `zlanhe`.
 Value returned by the function.

?lantb

Returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a triangular band matrix.

```
val = slantb ( norm, uplo, diag, n, k, ab, ldab, work )
val = dlantb ( norm, uplo, diag, n, k, ab, ldab, work )
```

```

val = clantb ( norm, uplo, diag, n, k, ab, ldab, work )
val = zlantb ( norm, uplo, diag, n, k, ab, ldab, work )

```

Discussion

The function `?lantb` returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of an n -by- n triangular band matrix A , with $(k + 1)$ diagonals.

The value `val` returned by the function is:

```

val = max(abs(Aij)), if norm = 'M' or 'm'
      = norm1(A),    if norm = '1' or 'O' or 'o'
      = normI(A),   if norm = 'I' or 'i'
      = normF(A),   if norm = 'F', 'f', 'E' or 'e'

```

where `norm1` denotes the 1-norm of a matrix (maximum column sum), `normI` denotes the infinity norm of a matrix (maximum row sum) and `normF` denotes the Frobenius norm of a matrix (square root of sum of squares). Note that `max(abs(Aij))` is not a matrix norm.

Input Parameters

`norm` CHARACTER*1. Specifies the value to be returned in `?lantb` as described above.

`uplo` CHARACTER*1. Specifies whether the matrix A is upper or lower triangular.
 = 'U': Upper triangular
 = 'L': Lower triangular.

`diag` CHARACTER*1. Specifies whether or not the matrix A is unit triangular.
 = 'N': Non-unit triangular
 = 'U': Unit triangular.

`n` INTEGER. The order of the matrix A .
 $n \geq 0$. When $n = 0$, `?lantb` is set to zero.

`k` INTEGER. The number of super-diagonals of the matrix A if `uplo = 'U'`, or the number of sub-diagonals of the matrix A if `uplo = 'L'`. $k \geq 0$.

ab REAL for `slantb`
DOUBLE PRECISION for `dlantb`
COMPLEX for `clantb`
COMPLEX*16 for `zlantb`
Array, DIMENSION (*ldab*,*n*). The upper or lower triangular band matrix *A*, stored in the first *k*+1 rows of *ab*. The *j*-th column of *A* is stored in the *j*-th column of the array *ab* as follows:
if *uplo* = 'U', $ab(k+1+i-j, j) = a(i, j)$ for $\max(1, j-k) \leq i \leq j$;
if *uplo* = 'L', $ab(1+i-j, j) = a(i, j)$ for $j \leq i \leq \min(n, j+k)$.
Note that when *diag* = 'U', the elements of the array *ab* corresponding to the diagonal elements of the matrix *A* are not referenced, but are assumed to be one.

ldab INTEGER. The leading dimension of the array *ab*.
 $ldab \geq k+1$.

work REAL for `slantb` and `clantb`.
DOUBLE PRECISION for `dlantb` and `zlantb`.
Workspace array, DIMENSION (*lwork*), where $lwork \geq n$ when *norm* = 'I'; otherwise, *work* is not referenced.

Output Parameters

val REAL for `slantb/clantb`.
DOUBLE PRECISION for `dlantb/zlantb`.
Value returned by the function.

?lantp

Returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a triangular matrix supplied in packed form.

```
val = slantp ( norm, uplo, diag, n, ap, work )  
val = dlantp ( norm, uplo, diag, n, ap, work )  
val = clantp ( norm, uplo, diag, n, ap, work )  
val = zlantp ( norm, uplo, diag, n, ap, work )
```

Discussion

The function `?lantp` returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a triangular matrix A , supplied in packed form.

The value `val` returned by the function is:

```
val = max(abs(Aij)), if norm = 'M' or 'm'  
      = norm1(A),    if norm = '1' or 'O' or 'o'  
      = normI(A),   if norm = 'I' or 'i'  
      = normF(A),   if norm = 'F', 'f', 'E' or 'e'
```

where `norm1` denotes the 1-norm of a matrix (maximum column sum), `normI` denotes the infinity norm of a matrix (maximum row sum) and `normF` denotes the Frobenius norm of a matrix (square root of sum of squares). Note that `max(abs(Aij))` is not a matrix norm.

Input Parameters

`norm` CHARACTER*1. Specifies the value to be returned in `?lantp` as described above.

<i>uplo</i>	<p>CHARACTER*1. Specifies whether the matrix <i>A</i> is upper or lower triangular.</p> <p>= 'U': Upper triangular</p> <p>= 'L': Lower triangular.</p>
<i>diag</i>	<p>CHARACTER*1. Specifies whether or not the matrix <i>A</i> is unit triangular.</p> <p>= 'N': Non-unit triangular</p> <p>= 'U': Unit triangular.</p>
<i>n</i>	<p>INTEGER. The order of the matrix <i>A</i>.</p> <p>$n \geq 0$. When $n = 0$, <i>?lantp</i> is set to zero.</p>
<i>ap</i>	<p>REAL for <i>slantp</i></p> <p>DOUBLE PRECISION for <i>dlantp</i></p> <p>COMPLEX for <i>clantp</i></p> <p>COMPLEX*16 for <i>zlantp</i></p> <p>Array, DIMENSION ($n(n+1)/2$). The upper or lower triangular matrix <i>A</i>, packed columnwise in a linear array. The <i>j</i>-th column of <i>A</i> is stored in the array <i>ap</i> as follows:</p> <p>if <i>uplo</i> = 'U', $AP(i + (j-1)j/2) = a(i,j)$ for $1 \leq i \leq j$;</p> <p>if <i>uplo</i> = 'L', $ap(i + (j-1)(2n-j)/2) = a(i,j)$ for $j \leq i \leq n$.</p> <p>Note that when <i>diag</i> = 'U', the elements of the array <i>ap</i> corresponding to the diagonal elements of the matrix <i>A</i> are not referenced, but are assumed to be one.</p>
<i>work</i>	<p>REAL for <i>slantp</i> and <i>clantp</i>.</p> <p>DOUBLE PRECISION for <i>dlantp</i> and <i>zlantp</i>.</p> <p>Workspace array, DIMENSION (<i>lwork</i>), where $lwork \geq n$ when <i>norm</i> = 'I'; otherwise, <i>work</i> is not referenced.</p>

Output Parameters

<i>val</i>	<p>REAL for <i>slantp/clantp</i>.</p> <p>DOUBLE PRECISION for <i>dlantp/zlantp</i>.</p> <p>Value returned by the function.</p>
------------	--

?lantr

Returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a trapezoidal or triangular matrix.

```
val = slantr ( norm, uplo, diag, m, n, a, lda, work )
val = dlantr ( norm, uplo, diag, m, n, a, lda, work )
val = clantr ( norm, uplo, diag, m, n, a, lda, work )
val = zlantr ( norm, uplo, diag, m, n, a, lda, work )
```

Discussion

The function `?lantr` returns the value of the 1-norm, or the Frobenius norm, or the infinity norm, or the element of largest absolute value of a trapezoidal or triangular matrix A .

The value `val` returned by the function is:

```
val = max(abs(Aij)), if norm = 'M' or 'm'
      = norm1(A),    if norm = '1' or 'O' or 'o'
      = normI(A),   if norm = 'I' or 'i'
      = normF(A),   if norm = 'F', 'f', 'E' or 'e'
```

where `norm1` denotes the 1-norm of a matrix (maximum column sum), `normI` denotes the infinity norm of a matrix (maximum row sum) and `normF` denotes the Frobenius norm of a matrix (square root of sum of squares). Note that `max(abs(Aij))` is not a matrix norm.

Input Parameters

`norm` CHARACTER*1. Specifies the value to be returned in `?lantr` as described above.

<i>uplo</i>	<p>CHARACTER*1. Specifies whether the matrix <i>A</i> is upper or lower trapezoidal.</p> <p>= 'U': Upper trapezoidal</p> <p>= 'L': Lower trapezoidal.</p> <p>Note that <i>A</i> is triangular instead of trapezoidal if $m = n$.</p>
<i>diag</i>	<p>CHARACTER*1. Specifies whether or not the matrix <i>A</i> has unit diagonal.</p> <p>= 'N': Non-unit diagonal</p> <p>= 'U': Unit diagonal.</p>
<i>m</i>	<p>INTEGER. The number of rows of the matrix <i>A</i>.</p> <p>$m \geq 0$, and if <i>uplo</i> = 'U', $m \leq n$. When $m = 0$, ?lantr is set to zero.</p>
<i>n</i>	<p>INTEGER. The number of columns of the matrix <i>A</i>.</p> <p>$n \geq 0$, and if <i>uplo</i> = 'L', $n \leq m$. When $n = 0$, ?lantr is set to zero.</p>
<i>a</i>	<p>REAL for slantr</p> <p>DOUBLE PRECISION for dlantr</p> <p>COMPLEX for clantr</p> <p>COMPLEX*16 for zlantr</p> <p>Array, DIMENSION (<i>lda</i>,<i>n</i>).</p> <p>The trapezoidal matrix <i>A</i> (<i>A</i> is triangular if $m = n$).</p> <p>If <i>uplo</i> = 'U', the leading <i>m</i>-by-<i>n</i> upper trapezoidal part of the array <i>a</i> contains the upper trapezoidal matrix, and the strictly lower triangular part of <i>a</i> is not referenced.</p> <p>If <i>uplo</i> = 'L', the leading <i>m</i>-by-<i>n</i> lower trapezoidal part of the array <i>a</i> contains the lower trapezoidal matrix, and the strictly upper triangular part of <i>a</i> is not referenced.</p> <p>Note that when <i>diag</i> = 'U', the diagonal elements of <i>a</i> are not referenced and are assumed to be one.</p>
<i>lda</i>	<p>INTEGER. The leading dimension of the array <i>a</i>.</p> <p>$lda \geq \max(m,1)$.</p>

work REAL for *slantr/clantrp*.
 DOUBLE PRECISION for *dlantr/zlantr*.
 Workspace array, DIMENSION (*lwork*), where
lwork ≥ *m* when *norm* = 'I'; otherwise, *work* is not
 referenced.

Output Parameters

val REAL for *slantr/clantrp*.
 DOUBLE PRECISION for *dlantr/zlantr*.
 Value returned by the function.

?lanv2

*Computes the Schur factorization of a
 real 2-by-2 nonsymmetric matrix in
 standard form.*

```
call slanv2 ( a, b, c, d, rt1r, rt1i, rt2r, rt2i, cs, sn )
call dlanv2 ( a, b, c, d, rt1r, rt1i, rt2r, rt2i, cs, sn )
```

Discussion

The routine computes the Schur factorization of a real 2-by-2 nonsymmetric matrix in standard form:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix} = \begin{bmatrix} cs & -sn \\ sn & cs \end{bmatrix} \begin{bmatrix} aa & bb \\ cc & dd \end{bmatrix} \begin{bmatrix} cs & sn \\ -sn & cs \end{bmatrix}$$

where either

1. $cc = 0$ so that aa and dd are real eigenvalues of the matrix, or
2. $aa = dd$ and $bb*cc < 0$, so that $aa \pm \text{sqrt}(bb*cc)$ are complex conjugate eigenvalues.

The routine was adjusted to reduce the risk of cancellation errors, when computing real eigenvalues, and to ensure, if possible, that $\text{abs}(rt1r) \geq \text{abs}(rt2r)$.

Input Parameters

a, b, c, d REAL for `slanv2`
 DOUBLE PRECISION for `dlanv2`.
 On entry, elements of the input matrix.

Output Parameters

a, b, c, d On exit, overwritten by the elements of the standardized Schur form.

rt1r, rt1i,
rt2r, rt2i, REAL for `slanv2`
 DOUBLE PRECISION for `dlanv2`.
 The real and imaginary parts of the eigenvalues. If the eigenvalues are a complex conjugate pair, $rt1i > 0$.

cs, sn REAL for `slanv2`
 DOUBLE PRECISION for `dlanv2`.
 Parameters of the rotation matrix.

?lapll

Measures the linear dependence of two vectors.

```
call slapll ( n, x, incx, y, incy, ssmi )
call dlapll ( n, x, incx, y, incy, ssmi )
call clapll ( n, x, incx, y, incy, ssmi )
call zlapll ( n, x, incx, y, incy, ssmi )
```

Discussion

Given two column vectors x and y of length n , let

$A = (x \ y)$ be the n -by-2 matrix.

The routine `?lap11` first computes the QR factorization of A as $A = QR$ and then computes the SVD of the 2-by-2 upper triangular matrix R . The smaller singular value of R is returned in `ssmin`, which is used as the measurement of the linear dependency of the vectors x and y .

Input Parameters

`n` `INTEGER`. The length of the vectors x and y .

`x` `REAL` for `slap11`
 `DOUBLE PRECISION` for `dlap11`
 `COMPLEX` for `clap11`
 `COMPLEX*16` for `zlap11`
 Array, `DIMENSION (1+(n-1)incx)`.
 On entry, `x` contains the n -vector x .

`y` `REAL` for `slap11`
 `DOUBLE PRECISION` for `dlap11`
 `COMPLEX` for `clap11`
 `COMPLEX*16` for `zlap11`
 Array, `DIMENSION (1+(n-1)incy)`. On entry, `y`
 contains the n -vector y .

`incx` `INTEGER`. The increment between successive elements
 of x ; `incx` > 0.

`incy` `INTEGER`. The increment between successive elements
 of y ; `incy` > 0.

Output Parameters

`x` On exit, `x` is overwritten.

`y` On exit, `y` is overwritten.

`ssmin` `REAL` for `slap11/clap11`
 `DOUBLE PRECISION` for `dlap11/zlap11`
 The smallest singular value of the n -by-2 matrix
 $A = (x \ y)$.

?lapmt

Performs a forward or backward permutation of the columns of a matrix.

```
call slapmt ( forwrđ, m, n, x, ldx, k )
call dlapmt ( forwrđ, m, n, x, ldx, k )
call clapmt ( forwrđ, m, n, x, ldx, k )
call zlapmt ( forwrđ, m, n, x, ldx, k )
```

Discussion

The routine ?lapmt rearranges the columns of the m -by- n matrix X as specified by the permutation $k(1),k(2),\dots,k(n)$ of the integers $1,\dots,n$.

If $forwrđ = .TRUE.$, forward permutation:

$X(*,k(j))$ is moved to $X(*,j)$ for $j = 1,2,\dots,n$.

If $forwrđ = .FALSE.$, backward permutation:

$X(*,j)$ is moved to $X(*,k(j))$ for $j = 1,2,\dots,n$.

Input Parameters

<i>forwrđ</i>	LOGICAL. If $forwrđ = .TRUE.$, forward permutation If $forwrđ = .FALSE.$, backward permutation
<i>m</i>	INTEGER. The number of rows of the matrix X . $m \geq 0$.
<i>n</i>	INTEGER. The number of columns of the matrix X . $n \geq 0$.
<i>x</i>	REAL for slapmt DOUBLE PRECISION for dlapmt COMPLEX for clapmt COMPLEX*16 for zlapmt Array, DIMENSION (ldx,n). On entry, the m -by- n matrix X .

ldx INTEGER. The leading dimension of the array *x*,
ldx $\geq \max(1,m)$.

k INTEGER.
Array, DIMENSION (*n*). On entry, *k* contains the
permutation vector.

Output Parameters

x On exit, *x* contains the permuted matrix *X*.

?lapy2

Returns $\sqrt{x^2+y^2}$.

```
val = slapy2 ( x, y )  
val = dlapy2 ( x, y )
```

Discussion

The function ?lapy2 returns $\sqrt{x^2+y^2}$, avoiding unnecessary overflow or harmful underflow.

Input Parameters

x, y REAL for slapy2
 DOUBLE PRECISION for dlapy2
Specify the input values *x* and *y*.

Output Parameters

val REAL for slapy2
 DOUBLE PRECISION for dlapy2.
Value returned by the function.

?lapy3

Returns $\text{sqrt}(x^2+y^2+z^2)$.

```
val = slapy3 ( x, y, z )
val = dlapy3 ( x, y, z )
```

Discussion

The function ?lapy3 returns $\text{sqrt}(x^2+y^2+z^2)$, avoiding unnecessary overflow or harmful underflow.

Input Parameters

x, y, z REAL for slapy3
 DOUBLE PRECISION for dlapy3
 Specify the input values *x*, *y* and *z*.

Output Parameters

val REAL for slapy3
 DOUBLE PRECISION for dlapy3.
 Value returned by the function.

?laqgb

Scales a general band matrix, using row and column scaling factors computed by ?gbequ.

```
call slaqgb ( m, n, kl, ku, ab, ldab, r, c, rowcnd,
             colcnd, amax, equed )
call dlaqgb ( m, n, kl, ku, ab, ldab, r, c, rowcnd,
             colcnd, amax, equed )
```

```

call claqqb ( m, n, kl, ku, ab, ldab, r, c, rowcnd,
              colcnd, amax, equed )
call zlaqqb ( m, n, kl, ku, ab, ldab, r, c, rowcnd,
              colcnd, amax, equed )

```

Discussion

The routine equilibrates a general m -by- n band matrix A with kl subdiagonals and ku superdiagonals using the row and column scaling factors in the vectors r and c .

Input Parameters

m **INTEGER.** The number of rows of the matrix A .
 $m \geq 0$.

n **INTEGER.** The number of columns of the matrix A .
 $n \geq 0$.

kl **INTEGER.** The number of subdiagonals within the band of A . $kl \geq 0$.

ku **INTEGER.** The number of superdiagonals within the band of A . $ku \geq 0$.

ab **REAL** for slaqqb
DOUBLE PRECISION for dlaqqb
COMPLEX for claqqb
COMPLEX*16 for zlaqqb
Array, **DIMENSION** ($ldab, n$). On entry, the matrix A in band storage, in rows 1 to $kl+ku+1$. The j -th column of A is stored in the j -th column of the array ab as follows:
 $ab(ku+1+i-j, j) = A(i, j)$ for
 $\max(1, j-ku) \leq i \leq \min(m, j+kl)$.

ldab **INTEGER.** The leading dimension of the array ab .
 $lda \geq kl+ku+1$.

amax **REAL** for slaqqb/claqqb
DOUBLE PRECISION for dlaqqb/zlaqqb
Absolute value of largest matrix entry.

Output Parameters

<i>ab</i>	On exit, the equilibrated matrix, in the same storage format as <i>A</i> . See <i>equed</i> for the form of the equilibrated matrix.
<i>r, c</i>	REAL for <i>slaqgb/claqgb</i> DOUBLE PRECISION for <i>dlaqgb/zlaqgb</i> Arrays <i>r</i> (<i>m</i>), <i>c</i> (<i>n</i>). Contain the row and column scale factors for <i>A</i> , respectively.
<i>rowcnd</i>	REAL for <i>slaqgb/claqgb</i> DOUBLE PRECISION for <i>dlaqgb/zlaqgb</i> Ratio of the smallest <i>r</i> (<i>i</i>) to the largest <i>r</i> (<i>i</i>).
<i>colcnd</i>	REAL for <i>slaqgb/claqgb</i> DOUBLE PRECISION for <i>dlaqgb/zlaqgb</i> Ratio of the smallest <i>c</i> (<i>i</i>) to the largest <i>c</i> (<i>i</i>).
<i>equed</i>	CHARACTER*1. Specifies the form of equilibration that was done. If <i>equed</i> = 'N': No equilibration If <i>equed</i> = 'R': Row equilibration, that is, <i>A</i> has been premultiplied by <i>diag</i> (<i>r</i>). If <i>equed</i> = 'C': Column equilibration, that is, <i>A</i> has been postmultiplied by <i>diag</i> (<i>c</i>). If <i>equed</i> = 'B': Both row and column equilibration, that is, <i>A</i> has been replaced by <i>diag</i> (<i>r</i>)* <i>A</i> * <i>diag</i> (<i>c</i>).

Application Notes

The routine uses internal parameters *thresh*, *large*, and *small*, which have the following meaning. *thresh* is a threshold value used to decide if row or column scaling should be done based on the ratio of the row or column scaling factors. If *rowcnd* < *thresh*, row scaling is done, and if *colcnd* < *thresh*, column scaling is done. *large* and *small* are threshold values used to decide if row scaling should be done based on the absolute size of the largest matrix element. If *amax* > *large* or *amax* < *small*, row scaling is done.

?laqge

Scales a general rectangular matrix,
using row and column scaling factors
computed by ?geequ.

```
call slaqge ( m, n, a, lda, r, c, rowcnd, colcnd, amax, equed )
call dlaqge ( m, n, a, lda, r, c, rowcnd, colcnd, amax, equed )
call claqge ( m, n, a, lda, r, c, rowcnd, colcnd, amax, equed )
call zlaqge ( m, n, a, lda, r, c, rowcnd, colcnd, amax, equed )
```

Discussion

The routine equilibrates a general m -by- n matrix A using the row and scaling factors in the vectors r and c .

Input Parameters

m **INTEGER.** The number of rows of the matrix A .
 $m \geq 0$.

n **INTEGER.** The number of columns of the matrix A .
 $n \geq 0$.

a **REAL** for slaqge
DOUBLE PRECISION for dlaqge
COMPLEX for claqge
COMPLEX*16 for zlaqge
Array, **DIMENSION** (lda,n). On entry, the m -by- n matrix A .

lda **INTEGER.** The leading dimension of the array A .
 $lda \geq \max(m,1)$.

r **REAL** for slangge/claqge
DOUBLE PRECISION for dlaqge/zlaqge
Array, **DIMENSION** (m). The row scale factors for A .

<i>c</i>	REAL for <code>slanqge/claqge</code> DOUBLE PRECISION for <code>dlaqge/zlaqge</code> Array, DIMENSION (<i>n</i>). The column scale factors for A.
<i>rowcnd</i>	REAL for <code>slanqge/claqge</code> DOUBLE PRECISION for <code>dlaqge/zlaqge</code> Ratio of the smallest <i>r</i> (<i>i</i>) to the largest <i>r</i> (<i>i</i>).
<i>colcnd</i>	REAL for <code>slanqge/claqge</code> DOUBLE PRECISION for <code>dlaqge/zlaqge</code> Ratio of the smallest <i>c</i> (<i>i</i>) to the largest <i>c</i> (<i>i</i>).
<i>amax</i>	REAL for <code>slanqge/claqge</code> DOUBLE PRECISION for <code>dlaqge/zlaqge</code> Absolute value of largest matrix entry.

Output Parameters

<i>a</i>	On exit, the equilibrated matrix. See <i>equed</i> for the form of the equilibrated matrix.
<i>equed</i>	CHARACTER*1. Specifies the form of equilibration that was done. If <i>equed</i> = 'N': No equilibration If <i>equed</i> = 'R': Row equilibration, that is, A has been premultiplied by $\text{diag}(r)$. If <i>equed</i> = 'C': Column equilibration, that is, A has been postmultiplied by $\text{diag}(c)$. If <i>equed</i> = 'B': Both row and column equilibration, that is, A has been replaced by $\text{diag}(r)*A*\text{diag}(c)$.

Application Notes

The routine uses internal parameters *thresh*, *large*, and *small*, which have the following meaning. *thresh* is a threshold value used to decide if row or column scaling should be done based on the ratio of the row or column scaling factors. If *rowcnd* < *thresh*, row scaling is done, and if *colcnd* < *thresh*, column scaling is done. *large* and *small* are threshold values used to decide if row scaling should be done based on the absolute size of the largest matrix element. If *amax* > *large* or *amax* < *small*, row scaling is done.

?laqp2

Computes a *QR* factorization with column pivoting of the matrix block.

```
call slaqp2 ( m, n, offset, a, lda, jpvt, tau, vn1, vn2, work )
call dlaqp2 ( m, n, offset, a, lda, jpvt, tau, vn1, vn2, work )
call claqp2 ( m, n, offset, a, lda, jpvt, tau, vn1, vn2, work )
call zlaqp2 ( m, n, offset, a, lda, jpvt, tau, vn1, vn2, work )
```

Discussion

The routine computes a *QR* factorization with column pivoting of the block $A(\text{offset}+1:m, 1:n)$. The block $A(1:\text{offset}, 1:n)$ is accordingly pivoted, but not factorized.

Input Parameters

m **INTEGER**. The number of rows of the matrix *A*.
m ≥ 0.

n **INTEGER**. The number of columns of the matrix *A*.
n ≥ 0.

offset **INTEGER**. The number of rows of the matrix *A* that must be pivoted but no factorized. *offset* ≥ 0.

a **REAL** for `slaqp2`
DOUBLE PRECISION for `dlaqp2`
COMPLEX for `claqp2`
COMPLEX*16 for `zlaqp2`
Array, **DIMENSION** (*lda*,*n*). On entry, the *m*-by-*n* matrix *A*.

lda **INTEGER**. The leading dimension of the array *A*. *lda* ≥ max(1,*m*).

<i>jpvt</i>	<p>INTEGER .</p> <p>Array, DIMENSION (<i>n</i>). On entry, if <i>jpvt</i>(<i>i</i>) ≠ 0, the <i>i</i>-th column of <i>A</i> is permuted to the front of <i>A*P</i> (a leading column); if <i>jpvt</i>(<i>i</i>) = 0, the <i>i</i>-th column of <i>A</i> is a free column.</p>
<i>vn1, vn2</i>	<p>REAL for <i>slaqp2/claqp2</i></p> <p>DOUBLE PRECISION for <i>dlaqp2/zlaqp2</i></p> <p>Arrays, DIMENSION (<i>n</i>) each. Contain the vectors with the partial and exact column norms, respectively.</p>
<i>work</i>	<p>REAL for <i>slaqp2</i></p> <p>DOUBLE PRECISION for <i>dlaqp2</i></p> <p>COMPLEX for <i>claqp2</i></p> <p>COMPLEX*16 for <i>zlaqp2</i></p> <p>Workspace array, DIMENSION (<i>n</i>).</p>

Output Parameters

<i>a</i>	<p>On exit, the upper triangle of block <i>A</i>(<i>offset</i>+1:<i>m</i>,1:<i>n</i>) is the triangular factor obtained; the elements in block <i>A</i>(<i>offset</i>+1:<i>m</i>,1:<i>n</i>) below the diagonal, together with the array <i>tau</i>, represent the orthogonal matrix <i>Q</i> as a product of elementary reflectors. Block <i>A</i>(1:<i>offset</i>,1:<i>n</i>) has been accordingly pivoted, but not factorized.</p>
<i>jpvt</i>	<p>On exit, if <i>jpvt</i>(<i>i</i>) = <i>k</i>, then the <i>i</i>-th column of <i>A*P</i> was the <i>k</i>-th column of <i>A</i>.</p>
<i>tau</i>	<p>REAL for <i>slaqp2</i></p> <p>DOUBLE PRECISION for <i>dlaqp2</i></p> <p>COMPLEX for <i>claqp2</i></p> <p>COMPLEX*16 for <i>zlaqp2</i></p> <p>Array, DIMENSION (min(<i>m</i>,<i>n</i>)). The scalar factors of the elementary reflectors.</p>
<i>vn1, vn2</i>	<p>Contain the vectors with the partial and exact column norms, respectively.</p>

?laqps

Computes a step of *QR* factorization with column pivoting of a real *m*-by-*n* matrix *A* by using BLAS level 3.

```
call slaqps ( m, n, offset, nb, kb, a, lda, jpvt, tau,
             vn1, vn2, auxv, f, ldf )
call dlaqps ( m, n, offset, nb, kb, a, lda, jpvt, tau,
             vn1, vn2, auxv, f, ldf )
call claqps ( m, n, offset, nb, kb, a, lda, jpvt, tau,
             vn1, vn2, auxv, f, ldf )
call zlaqps ( m, n, offset, nb, kb, a, lda, jpvt, tau,
             vn1, vn2, auxv, f, ldf )
```

Discussion

This routine computes a step of *QR* factorization with column pivoting of a real *m*-by-*n* matrix *A* by using BLAS level 3. The routine tries to factorize *nb* columns from *A* starting from the row *offset*+1, and updates all of the matrix with BLAS level 3 routine `?gemm`.

In some cases, due to catastrophic cancellations, `?laqps` cannot factorize *nb* columns. Hence, the actual number of factorized columns is returned in *kb*.

Block $A(1:offset, 1:n)$ is accordingly pivoted, but not factorized.

Input Parameters

<i>m</i>	INTEGER. The number of rows of the matrix <i>A</i> . $m \geq 0$.
<i>n</i>	INTEGER. The number of columns of the matrix <i>A</i> . $n \geq 0$.
<i>offset</i>	INTEGER. The number of rows of <i>A</i> that have been factorized in previous steps.
<i>nb</i>	INTEGER. The number of columns to factorize.

<i>a</i>	<p>REAL for <code>slaqps</code> DOUBLE PRECISION for <code>dlaqps</code> COMPLEX for <code>claqps</code> COMPLEX*16 for <code>zlaqps</code> Array, DIMENSION (<i>lda</i>,<i>n</i>). On entry, the <i>m</i>-by-<i>n</i> matrix <i>A</i>.</p>
<i>lda</i>	<p>INTEGER. The leading dimension of the array <i>a</i>. <i>lda</i> ≥ max(1,<i>m</i>).</p>
<i>jpvt</i>	<p>INTEGER. Array, DIMENSION (<i>n</i>). If <i>jpvt</i>(<i>i</i>) = <i>k</i> then column <i>k</i> of the full matrix <i>A</i> has been permuted into position <i>i</i> in <i>AP</i>.</p>
<i>vn1, vn2</i>	<p>REAL for <code>slaqps/claqps</code> DOUBLE PRECISION for <code>dlaqps/zlaqps</code> Arrays, DIMENSION (<i>n</i>) each. Contain the vectors with the partial and exact column norms, respectively.</p>
<i>auxv</i>	<p>REAL for <code>slaqps</code> DOUBLE PRECISION for <code>dlaqps</code> COMPLEX for <code>claqps</code> COMPLEX*16 for <code>zlaqps</code> Array, DIMENSION (<i>nb</i>). Auxiliary vector.</p>
<i>f</i>	<p>REAL for <code>slaqps</code> DOUBLE PRECISION for <code>dlaqps</code> COMPLEX for <code>claqps</code> COMPLEX*16 for <code>zlaqps</code> Array, DIMENSION (<i>ldf</i>,<i>nb</i>). Matrix $F' = L*Y' *A$.</p>
<i>ldf</i>	<p>INTEGER. The leading dimension of the array <i>f</i>. <i>ldf</i> ≥ max(1,<i>n</i>).</p>

Output Parameters

<i>kb</i>	<p>INTEGER. The number of columns actually factorized.</p>
<i>a</i>	<p>On exit, block <i>A</i>(<i>offset</i>+1:<i>m</i>,1:<i>kb</i>) is the triangular factor obtained and block <i>A</i>(1:<i>offset</i>,1:<i>n</i>) has been accordingly pivoted, but no factorized. The rest of the matrix, block <i>A</i>(<i>offset</i>+1:<i>m</i>,<i>kb</i>+1:<i>n</i>) has been updated.</p>

<i>jpvt</i>	INTEGER array, DIMENSION (<i>n</i>). If <i>jpvt</i> (<i>i</i>) = <i>k</i> then column <i>k</i> of the full matrix <i>A</i> has been permuted into position <i>i</i> in <i>AP</i> .
<i>tau</i>	REAL for <i>slaqps</i> DOUBLE PRECISION for <i>dlaqps</i> COMPLEX for <i>claqps</i> COMPLEX*16 for <i>zlaqps</i> Array, DIMENSION (<i>kb</i>). The scalar factors of the elementary reflectors.
<i>vn1, vn2</i>	The vectors with the partial and exact column norms, respectively.
<i>auxv</i>	Auxiliary vector.
<i>f</i>	Matrix $F' = L*Y' *A$.

?laqsb

Scales a symmetric/Hermitian band matrix, using scaling factors computed by ?pbequ.

```
call slaqsb ( uplo, n, kd, ab, ldab, s, scond, amax, equed )
call dlaqsb ( uplo, n, kd, ab, ldab, s, scond, amax, equed )
call claqsb ( uplo, n, kd, ab, ldab, s, scond, amax, equed )
call zlaqsb ( uplo, n, kd, ab, ldab, s, scond, amax, equed )
```

Discussion

The routine equilibrates a symmetric band matrix *A* using the scaling factors in the vector *s*.

Input Parameters

<i>uplo</i>	CHARACTER*1 . Specifies whether the upper or lower triangular part of the symmetric matrix <i>A</i> is stored. If <i>uplo</i> = 'U': upper triangular. If <i>uplo</i> = 'L': lower triangular.
<i>n</i>	INTEGER . The order of the matrix <i>A</i> . $n \geq 0$.
<i>kd</i>	INTEGER . The number of super-diagonals of the matrix <i>A</i> if <i>uplo</i> = 'U', or the number of sub-diagonals if <i>uplo</i> = 'L'. $kd \geq 0$.
<i>ab</i>	REAL for slaqs DOUBLE PRECISION for dlaqs COMPLEX for claqs COMPLEX*16 for zlaqs Array, DIMENSION (<i>ldab</i> , <i>n</i>). On entry, the upper or lower triangle of the symmetric band matrix <i>A</i> , stored in the first <i>kd</i> +1 rows of the array. The <i>j</i> -th column of <i>A</i> is stored in the <i>j</i> -th column of the array <i>ab</i> as follows: if <i>uplo</i> = 'U', $ab(kd+1+i-j, j) = A(i, j)$ for $\max(1, j-kd) \leq i \leq j$; if <i>uplo</i> = 'L', $ab(1+i-j, j) = A(i, j)$ for $j \leq i \leq \min(n, j+kd)$.
<i>ldab</i>	INTEGER . The leading dimension of the array <i>ab</i> . $ldab \geq kd+1$.
<i>scond</i>	REAL for slaqs / claqs DOUBLE PRECISION for dlaqs / zlaqs Ratio of the smallest $s(i)$ to the largest $s(i)$.
<i>amax</i>	REAL for slaqs / claqs DOUBLE PRECISION for dlaqs / zlaqs Absolute value of largest matrix entry.

Output Parameters

<i>ab</i>	On exit, if <i>info</i> = 0, the triangular factor <i>U</i> or <i>L</i> from the Cholesky factorization $A = U' U$ or $A = L L'$ of the band matrix <i>A</i> , in the same storage format as <i>A</i> .
-----------	---

s REAL for `slaqsb/claqsb`
 DOUBLE PRECISION for `dlaqsb/zlaqsb`
 Array, DIMENSION (*n*). The scale factors for *A*.

equed CHARACTER*1.
 Specifies whether or not equilibration was done.
 If *equed* = 'N': No equilibration.
 If *equed* = 'Y': Equilibration was done, that is, *A* has
 been replaced by $\text{diag}(s) * A * \text{diag}(s)$.

Application Notes

The routine uses internal parameters *thresh*, *large*, and *small*, which have the following meaning. *thresh* is a threshold value used to decide if scaling should be based on the ratio of the scaling factors. If *scond* < *thresh*, scaling is done. *large* and *small* are threshold values used to decide if scaling should be done based on the absolute size of the largest matrix element. If *amax* > *large* or *amax* < *small*, scaling is done.

?laqsp

Scales a symmetric/Hermitian matrix in packed storage, using scaling factors computed by ?ppequ.

```
call slaqsp ( uplo, n, ap, s, scnd, amax, equed )
call dlaqsp ( uplo, n, ap, s, scnd, amax, equed )
call claqsp ( uplo, n, ap, s, scnd, amax, equed )
call zlaqsp ( uplo, n, ap, s, scnd, amax, equed )
```

Discussion

The routine `?laqsp` equilibrates a symmetric matrix *A* using the scaling factors in the vector *s*.

Internal Parameters

<i>uplo</i>	<p>CHARACTER*1. Specifies whether the upper or lower triangular part of the symmetric matrix <i>A</i> is stored.</p> <p>If <i>uplo</i> = 'U': upper triangular.</p> <p>If <i>uplo</i> = 'L': lower triangular.</p>
<i>n</i>	<p>INTEGER. The order of the matrix <i>A</i>.</p> <p>$n \geq 0$.</p>
<i>ap</i>	<p>REAL for <i>slaqsp</i></p> <p>DOUBLE PRECISION for <i>dlaqsp</i></p> <p>COMPLEX for <i>claqsp</i></p> <p>COMPLEX*16 for <i>zlaqsp</i></p> <p>Array, DIMENSION $(n(n+1)/2)$. On entry, the upper or lower triangle of the symmetric matrix <i>A</i>, packed columnwise in a linear array. The <i>j</i>-th column of <i>A</i> is stored in the array <i>ap</i> as follows:</p> <p>if <i>uplo</i> = 'U', $ap(i + (j-1)j/2) = A(i, j)$ for $1 \leq i \leq j$;</p> <p>if <i>uplo</i> = 'L', $ap(i + (j-1)(2n-j)/2) = A(i, j)$ for $j \leq i \leq n$.</p>
<i>s</i>	<p>REAL for <i>slaqsp/claqsp</i></p> <p>DOUBLE PRECISION for <i>dlaqsp/zlaqsp</i></p> <p>Array, DIMENSION (n). The scale factors for <i>A</i>.</p>
<i>scond</i>	<p>REAL for <i>slaqsp/claqsp</i></p> <p>DOUBLE PRECISION for <i>dlaqsp/zlaqsp</i></p> <p>Ratio of the smallest $s(i)$ to the largest $s(i)$.</p>
<i>amax</i>	<p>REAL for <i>slaqsp/claqsp</i></p> <p>DOUBLE PRECISION for <i>dlaqsp/zlaqsp</i></p> <p>Absolute value of largest matrix entry.</p>

Output Parameters

<i>ap</i>	<p>On exit, the equilibrated matrix: $\text{diag}(s) * A * \text{diag}(s)$, in the same storage format as <i>A</i>.</p>
<i>equed</i>	<p>CHARACTER*1.</p> <p>Specifies whether or not equilibration was done.</p> <p>If <i>equed</i> = 'N': No equilibration.</p> <p>If <i>equed</i> = 'Y': Equilibration was done, that is, <i>A</i> has been replaced by $\text{diag}(s) * A * \text{diag}(s)$.</p>

Application Notes

The routine uses internal parameters *thresh*, *large*, and *small*, which have the following meaning. *thresh* is a threshold value used to decide if scaling should be based on the ratio of the scaling factors. If *scond* < *thresh*, scaling is done. *large* and *small* are threshold values used to decide if scaling should be done based on the absolute size of the largest matrix element. If *amax* > *large* or *amax* < *small*, scaling is done.

?laqsy

*Scales a symmetric/Hermitian matrix,
using scaling factors computed by
?poequ.*

```
call slaqsy ( uplo, n, a, lda, s, scnd, amax, equed )
call dlaqsy ( uplo, n, a, lda, s, scnd, amax, equed )
call claqsy ( uplo, n, a, lda, s, scnd, amax, equed )
call zlaqsy ( uplo, n, a, lda, s, scnd, amax, equed )
```

Discussion

The routine equilibrates a symmetric matrix *A* using the scaling factors in the vector *s*.

Input Parameters

uplo CHARACTER*1. Specifies whether the upper or lower triangular part of the symmetric matrix *A* is stored.
If *uplo* = 'U': upper triangular.
If *uplo* = 'L': lower triangular.

n INTEGER. The order of the matrix *A*.
 $n \geq 0$.

a REAL for *slaqsy*
DOUBLE PRECISION for *dlaqsy*
COMPLEX for *claqsy*

COMPLEX*16 for `zlaqsy`

Array, DIMENSION (lda, n). On entry, the symmetric matrix A . If `uplo` = 'U', the leading n -by- n upper triangular part of a contains the upper triangular part of the matrix A , and the strictly lower triangular part of a is not referenced. If `uplo` = 'L', the leading n -by- n lower triangular part of a contains the lower triangular part of the matrix A , and the strictly upper triangular part of a is not referenced.

`lda` INTEGER. The leading dimension of the array a .
 $lda \geq \max(n, 1)$.

`s` REAL for `slaqsy/claqsy`
 DOUBLE PRECISION for `dlaqsy/zlaqsy`
 Array, DIMENSION (n). The scale factors for A .

`scond` REAL for `slaqsy/claqsy`
 DOUBLE PRECISION for `dlaqsy/zlaqsy`
 Ratio of the smallest $s(i)$ to the largest $s(i)$.

`amax` REAL for `slaqsy/claqsy`
 DOUBLE PRECISION for `dlaqsy/zlaqsy`
 Absolute value of largest matrix entry.

Output Parameters

`a` On exit, if `equed` = 'Y', the equilibrated matrix:
 $\text{diag}(s) * A * \text{diag}(s)$.

`equed` CHARACTER*1.
 Specifies whether or not equilibration was done.
 If `equed` = 'N': No equilibration.
 If `equed` = 'Y': Equilibration was done, i.e., A has been replaced by $\text{diag}(s) * A * \text{diag}(s)$.

Application Notes

The routine uses internal parameters `thresh`, `large`, and `small`, which have the following meaning. `thresh` is a threshold value used to decide if scaling should be based on the ratio of the scaling factors. If `scond` <

thresh, scaling is done. *large* and *small* are threshold values used to decide if scaling should be done based on the absolute size of the largest matrix element. If *amax* > *large* or *amax* < *small*, scaling is done.

?laqtr

Solves a real quasi-triangular system of equations, or a complex quasi-triangular system of special form, in real arithmetic.

```
call slaqtr ( ltran, lreal, n, t, ldt, b, w, scale, x,
             work, info )
call dlaqtr ( ltran, lreal, n, t, ldt, b, w, scale, x,
             work, info )
```

Discussion

The routine ?laqtr solves the real quasi-triangular system

$$\text{op}(T) * p = \text{scale} * c, \quad \text{if } \textit{lreal} = \textit{.TRUE.}$$

or the complex quasi-triangular systems

$$\text{op}(T + \mathbf{i}B) * (p + \mathbf{i}q) = \text{scale} * (c + \mathbf{i}d), \quad \text{if } \textit{lreal} = \textit{.FALSE.}$$

in real arithmetic, where T is upper quasi-triangular.

If *lreal* = *.FALSE.*, then the first diagonal block of T must be 1-by-1, B is the specially structured matrix

$$B = \begin{bmatrix} b_1 & b_2 & \dots & b_n \\ & w & & \\ & & w & \\ & & & \dots \\ & & & & w \end{bmatrix}$$

$\text{op}(A) = A$ or A' , A' denotes the conjugate transpose of matrix A .

On input,

$$x = \begin{bmatrix} c \\ d \end{bmatrix}, \text{ on output } x = \begin{bmatrix} p \\ q \end{bmatrix}$$

This routine is designed for the condition number estimation in routine `?trsna`.

Input Parameters

<code>ltran</code>	LOGICAL. On entry, <code>ltran</code> specifies the option of conjugate transpose: = <code>.FALSE.</code> , $\text{op}(T + iB) = T + iB$, = <code>.TRUE.</code> , $\text{op}(T + iB) = (T + iB)'$.
<code>lreal</code>	LOGICAL. On entry, <code>lreal</code> specifies the input matrix structure: = <code>.FALSE.</code> , the input is complex = <code>.TRUE.</code> , the input is real.
<code>n</code>	INTEGER. On entry, <code>n</code> specifies the order of $T + iB$. $n \geq 0$.
<code>t</code>	REAL for <code>slaqtr</code> DOUBLE PRECISION for <code>dlaqtr</code> Array, dimension (ldt, n) . On entry, <code>t</code> contains a matrix in Schur canonical form. If <code>lreal = .FALSE.</code> , then the first diagonal block of <code>t</code> must be 1-by-1.
<code>ldt</code>	INTEGER. The leading dimension of the matrix T . $ldt \geq \max(1, n)$.
<code>b</code>	REAL for <code>slaqtr</code> DOUBLE PRECISION for <code>dlaqtr</code> Array, dimension (n) . On entry, <code>b</code> contains the elements to form the matrix B as described above. If <code>lreal = .TRUE.</code> , <code>b</code> is not referenced.
<code>w</code>	REAL for <code>slaqtr</code> DOUBLE PRECISION for <code>dlaqtr</code> On entry, <code>w</code> is the diagonal element of the matrix B . If <code>lreal = .TRUE.</code> , <code>w</code> is not referenced.

x REAL for `slaqtr`
DOUBLE PRECISION for `dlaqtr`
Array, dimension $(2n)$. On entry, *x* contains the right hand side of the system.

work REAL for `slaqtr`
DOUBLE PRECISION for `dlaqtr`
Workspace array, dimension (n) .

Output Parameters

scale REAL for `slaqtr`
DOUBLE PRECISION for `dlaqtr`
On exit, *scale* is the scale factor.

x On exit, *x* is overwritten by the solution.

info INTEGER.
If *info* = 0: successful exit.
If *info* = 1: the some diagonal 1-by-1 block has been perturbed by a small number `smin` to keep nonsingularity.
If *info* = 2: the some diagonal 2-by-2 block has been perturbed by a small number in `?1a1n2` to keep nonsingularity.



NOTE. *In the interests of speed, this routine does not check the inputs for errors.*

?lar1v

Computes the (scaled) r -th column of the inverse of the submatrix in rows $b1$ through bn of the tridiagonal matrix $LDL^T - \sigma I$.

```
call slar1v ( n, b1, bn, sigma, d, l, ld, lld, gersch, z,
             ztz, mingma, r, isuppz, work )
call dlar1v ( n, b1, bn, sigma, d, l, ld, lld, gersch, z,
             ztz, mingma, r, isuppz, work )
call clar1v ( n, b1, bn, sigma, d, l, ld, lld, gersch, z,
             ztz, mingma, r, isuppz, work )
call zlar1v ( n, b1, bn, sigma, d, l, ld, lld, gersch, z,
             ztz, mingma, r, isuppz, work )
```

Discussion

The routine `?lar1v` computes the (scaled) r -th column of the inverse of the submatrix in rows $b1$ through bn of the tridiagonal matrix $LDL^T - \sigma I$.

The following steps accomplish this computation :

1. Stationary qd transform, $LDL^T - \sigma I = L(+)\ D(+)\ L(+)^T$
2. Progressive qd transform, $LDL^T - \sigma I = U(-)\ D(-)\ U(-)^T$,
3. Computation of the diagonal elements of the inverse of $LDL^T - \sigma I$ by combining the above transforms, and choosing r as the index where the diagonal of the inverse is (one of the) largest in magnitude.
4. Computation of the (scaled) r -th column of the inverse using the twisted factorization obtained by combining the top part of the stationary and the bottom part of the progressive transform.

Input Parameters

n **INTEGER**. The order of the matrix LDL^T .

$b1$ **INTEGER**. First index of the submatrix of LDL^T .

<i>bn</i>	INTEGER. Last index of the submatrix of LDL^T .
<i>sigma</i>	REAL for <code>slarlv/clarlv</code> DOUBLE PRECISION for <code>dclarlv/zclarlv</code> The shift. Initially, when $r = 0$, <i>sigma</i> should be a good approximation to an eigenvalue of LDL^T .
<i>l</i>	REAL for <code>slarlv/clarlv</code> DOUBLE PRECISION for <code>dclarlv/zclarlv</code> Array, DIMENSION ($n-1$). The ($n-1$) subdiagonal elements of the unit bidiagonal matrix L , in elements 1 to $n-1$.
<i>d</i>	REAL for <code>slarlv/clarlv</code> DOUBLE PRECISION for <code>dclarlv/zclarlv</code> Array, DIMENSION (n). The n diagonal elements of the diagonal matrix D .
<i>ld</i>	REAL for <code>slarlv/clarlv</code> DOUBLE PRECISION for <code>dclarlv/zclarlv</code> Array, DIMENSION ($n-1$). The $n-1$ elements $L_i * D_i$.
<i>lld</i>	REAL for <code>slarlv/clarlv</code> DOUBLE PRECISION for <code>dclarlv/zclarlv</code> Array, DIMENSION ($n-1$). The $n-1$ elements $L_i * L_i * D_i$.
<i>gersch</i>	REAL for <code>slarlv/clarlv</code> DOUBLE PRECISION for <code>dclarlv/zclarlv</code> Array, DIMENSION ($2n$). The n Gerschgorin intervals. These are used to restrict the initial search for r , when r is input as 0.
<i>r</i>	INTEGER. Initially r should be input to be 0 and is then output as the index where the diagonal element of the inverse is largest in magnitude. In later iterations, this same value of r should be input.
<i>work</i>	REAL for <code>slarlv/clarlv</code> DOUBLE PRECISION for <code>dclarlv/zclarlv</code> Workspace array, DIMENSION ($4n$).

Output Parameters

<i>z</i>	REAL for <code>slarlv</code> DOUBLE PRECISION for <code>dlarlv</code> COMPLEX for <code>clarlv</code> COMPLEX*16 for <code>zlarlv</code> Array, DIMENSION (<i>n</i>). The (scaled) <i>r</i> -th column of the inverse. <i>z</i> (<i>r</i>) is returned to be 1.
<i>ztz</i>	REAL for <code>slarlv/clarlv</code> DOUBLE PRECISION for <code>dlarlv/zlarlv</code> The square of the norm of <i>z</i> .
<i>mingma</i>	REAL for <code>slarlv/clarlv</code> DOUBLE PRECISION for <code>dlarlv/zlarlv</code> The reciprocal of the largest (in magnitude) diagonal element of the inverse of $LDL^T - \sigma * I$.
<i>r</i>	On output, <i>r</i> is the index where the diagonal element of the inverse is largest in magnitude.
<i>isuppz</i>	INTEGER. Array, DIMENSION (2). The support of the vector in <i>z</i> , that is, the vector <i>z</i> is nonzero only in elements <i>isuppz</i> (1) through <i>isuppz</i> (2).

?lar2v

Applies a vector of plane rotations with real cosines and real/complex sines from both sides to a sequence of 2-by-2 symmetric/Hermitian matrices.

```
call slar2v ( n, x, y, z, incx, c, s, incc )
call dlar2v ( n, x, y, z, incx, c, s, incc )
call clar2v ( n, x, y, z, incx, c, s, incc )
call zlar2v ( n, x, y, z, incx, c, s, incc )
```

Discussion

The routine `?lar2v` applies a vector of real/complex plane rotations with real cosines from both sides to a sequence of 2-by-2 real symmetric or complex Hermitian matrices, defined by the elements of the vectors x , y and z . For $i = 1, 2, \dots, n$

$$\begin{bmatrix} x_i & z_i \\ \text{conjg}(z_i) & y_i \end{bmatrix} = \begin{bmatrix} c(i) & \text{conjg}(s(i)) \\ -s(i) & c(i) \end{bmatrix} \begin{bmatrix} x_i & z_i \\ \text{conjg}(z_i) & y_i \end{bmatrix} \begin{bmatrix} c(i) & -\text{conjg}(s(i)) \\ s(i) & c(i) \end{bmatrix}$$

Input Parameters

- n*** **INTEGER**. The number of plane rotations to be applied.
- x, y, z*** **REAL** for `slar2v`
DOUBLE PRECISION for `dlar2v`
COMPLEX for `clar2v`
COMPLEX*16 for `zlar2v`
 Arrays, **DIMENSION** $(1+(n-1)*incx)$ each. Contain the vectors x , y and z , respectively. For all flavors of `?lar2v`, elements of x and y are assumed to be real.
- incx*** **INTEGER**. The increment between elements of x , y , and z . $incx > 0$.
- c*** **REAL** for `slar2v/clar2v`
DOUBLE PRECISION for `dlar2v/zlar2v`
 Array, **DIMENSION** $(1+(n-1)*incc)$. The cosines of the plane rotations.
- s*** **REAL** for `slar2v`
DOUBLE PRECISION for `dlar2v`
COMPLEX for `clar2v`
COMPLEX*16 for `zlar2v`
 Array, **DIMENSION** $(1+(n-1)*incc)$. The sines of the plane rotations.
- incc*** **INTEGER**. The increment between elements of c and s . $incc > 0$.

Output Parameters

x, y, z Vectors *x*, *y* and *z*, containing the results of transform.

?larf

Applies an elementary reflector to a general rectangular matrix.

```
call slarf ( side, m, n, v, incv, tau, c, ldc, work )
call dlarf ( side, m, n, v, incv, tau, c, ldc, work )
call clarf ( side, m, n, v, incv, tau, c, ldc, work )
call zlarf ( side, m, n, v, incv, tau, c, ldc, work )
```

Discussion

The routine applies a real/complex elementary reflector H to a real/complex m -by- n matrix C , from either the left or the right. H is represented in the form

$$H = I - \tau * v * v',$$

where τ is a real/complex scalar and v is a real/complex vector.

If $\tau = 0$, then H is taken to be the unit matrix.

For `clarf/zlarf`, to apply H' (the conjugate transpose of H), supply `conjg(tau)` instead of `tau`.

Input Parameters

side CHARACTER*1.
 If *side* = 'L': form $H * C$
 If *side* = 'R': form $C * H$.

m INTEGER. The number of rows of the matrix C .

n INTEGER. The number of columns of the matrix C .

v REAL for `slarf`
 DOUBLE PRECISION for `dlarf`
 COMPLEX for `clarf`

COMPLEX*16 for `zlarf`
 Array, DIMENSION
 $(1 + (m-1)*\text{abs}(\text{incv}))$ if `side = 'L'` or
 $(1 + (n-1)*\text{abs}(\text{incv}))$ if `side = 'R'`.
 The vector v in the representation of H . v is not used if
 $\text{tau} = 0$.

`incv` INTEGER. The increment between elements of v .
 $\text{incv} \neq 0$.

`tau` REAL for `slarf`
 DOUBLE PRECISION for `dlarf`
 COMPLEX for `clarf`
 COMPLEX*16 for `zlarf`
 The value `tau` in the representation of H .

`c` REAL for `slarf`
 DOUBLE PRECISION for `dlarf`
 COMPLEX for `clarf`
 COMPLEX*16 for `zlarf`
 Array, DIMENSION (ldc, n) .
 On entry, the m -by- n matrix C .

`ldc` INTEGER. The leading dimension of the array `c`.
 $\text{ldc} \geq \max(1, m)$.

`work` REAL for `slarf`
 DOUBLE PRECISION for `dlarf`
 COMPLEX for `clarf`
 COMPLEX*16 for `zlarf`
 Workspace array, DIMENSION
 (n) if `side = 'L'` or
 (m) if `side = 'R'`.

Output Parameters

`c` On exit, `c` is overwritten by the matrix $H*C$ if `side = 'L'`, or $C*H$ if `side = 'R'`.

?larfb

Applies a block reflector or its transpose/conjugate-transpose to a general rectangular matrix.

```
call slarfb ( side, trans, direct, storev, m, n, k, v,
             ldv, t, ldt, c, ldc, work, ldwork )
call dlarfb ( side, trans, direct, storev, m, n, k, v,
             ldv, t, ldt, c, ldc, work, ldwork )
call clarfb ( side, trans, direct, storev, m, n, k, v,
             ldv, t, ldt, c, ldc, work, ldwork )
call zlarfb ( side, trans, direct, storev, m, n, k, v,
             ldv, t, ldt, c, ldc, work, ldwork )
```

Discussion

The routine ?larfb applies a complex block reflector H or its transpose H' to a complex m -by- n matrix C from either left or right.

Input Parameters

side CHARACTER*1.
 If *side* = 'L': apply H or H' from the left
 If *side* = 'R': apply H or H' from the right

trans CHARACTER*1.
 If *trans* = 'N': apply H (No transpose)
 If *trans* = 'C': apply H' (Conjugate transpose)

direct CHARACTER*1. Indicates how H is formed from a product of elementary reflectors
 If *direct* = 'F': $H = H(1) H(2) \dots H(k)$ (forward)
 If *direct* = 'B': $H = H(k) \dots H(2) H(1)$ (backward)

storev CHARACTER*1. Indicates how the vectors which define the elementary reflectors are stored:
 If *storev* = 'C': Column-wise
 If *storev* = 'R': Row-wise

<i>m</i>	INTEGER. The number of rows of the matrix <i>C</i> .
<i>n</i>	INTEGER. The number of columns of the matrix <i>C</i> .
<i>k</i>	INTEGER. The order of the matrix <i>T</i> (equal to the number of elementary reflectors whose product defines the block reflector).
<i>v</i>	REAL for slarfb DOUBLE PRECISION for dlarfb COMPLEX for clarfb COMPLEX*16 for zlarfb Array, DIMENSION (<i>ldv</i> , <i>k</i>) if <i>storev</i> = 'C' (<i>ldv</i> , <i>m</i>) if <i>storev</i> = 'R' and <i>side</i> = 'L' (<i>ldv</i> , <i>n</i>) if <i>storev</i> = 'R' and <i>side</i> = 'R' The matrix <i>V</i> .
<i>ldv</i>	INTEGER. The leading dimension of the array <i>v</i> . If <i>storev</i> = 'C' and <i>side</i> = 'L', $ldv \geq \max(1, m)$; if <i>storev</i> = 'C' and <i>side</i> = 'R', $ldv \geq \max(1, n)$; if <i>storev</i> = 'R', $ldv \geq k$.
<i>t</i>	REAL for slarfb DOUBLE PRECISION for dlarfb COMPLEX for clarfb COMPLEX*16 for zlarfb Array, DIMENSION (<i>ldt</i> , <i>k</i>). Contains the triangular <i>k</i> -by- <i>k</i> matrix <i>T</i> in the representation of the block reflector.
<i>ldt</i>	INTEGER. The leading dimension of the array <i>t</i> . $ldt \geq k$.
<i>c</i>	REAL for slarfb DOUBLE PRECISION for dlarfb COMPLEX for clarfb COMPLEX*16 for zlarfb Array, DIMENSION (<i>ldc</i> , <i>n</i>). On entry, the <i>m</i> -by- <i>n</i> matrix <i>C</i> .

ldc **INTEGER**. The leading dimension of the array *c*.
ldc ≥ max(1,*m*).

work **REAL** for **slarfb**
DOUBLE PRECISION for **dlarfb**
COMPLEX for **clarfb**
COMPLEX*16 for **zlarfb**
Workspace array, **DIMENSION** (*ldwork*, *k*).

ldwork **INTEGER**. The leading dimension of the array *work*.
If *side* = 'L', *ldwork* ≥ max(1, *n*);
if *side* = 'R', *ldwork* ≥ max(1, *m*).

Output parameters

c On exit, *c* is overwritten by $H * C$ or $H' * C$ or $C * H$ or $C * H'$.

?larfg

*Generates an elementary reflector
(Householder matrix).*

```
call slarfg ( n, alpha, x, incx, tau )
call dlarfg ( n, alpha, x, incx, tau )
call clarfg ( n, alpha, x, incx, tau )
call zlarfg ( n, alpha, x, incx, tau )
```

Discussion

The routine **?larfg** generates a real/complex elementary reflector H of order n , such that

$$H' * \begin{bmatrix} alpha \\ x \end{bmatrix} = \begin{bmatrix} beta \\ 0 \end{bmatrix}, \quad H' * H = I,$$

where α and β are scalars (with β real for all flavors), and x is an $(n-1)$ -element real/complex vector. H is represented in the form

$$H = I - \tau * \begin{bmatrix} 1 \\ v \end{bmatrix} * \begin{bmatrix} 1 & v' \end{bmatrix}$$

where τ is a real/complex scalar and v is a real/complex $(n-1)$ -element vector. Note that for `clarfg/zlarfg`, H is not Hermitian.

If the elements of x are all zero (and, for complex flavors, α is real), then $\tau = 0$ and H is taken to be the unit matrix.

Otherwise, $1 \leq \tau \leq 1$ (for real flavors), or
 $1 \leq \text{Re}(\tau) \leq 1$ and $|\text{Im}(\tau)| \leq 1$ (for complex flavors).

Input Parameters

n **INTEGER**. The order of the elementary reflector.

α **REAL** for `sclarfg`
DOUBLE PRECISION for `dclarfg`
COMPLEX for `clarfg`
COMPLEX*16 for `zlarfg`
On entry, the value α .

x **REAL** for `sclarfg`
DOUBLE PRECISION for `dclarfg`
COMPLEX for `clarfg`
COMPLEX*16 for `zlarfg`
Array, **DIMENSION** $(1+(n-2)*\text{abs}(\text{incx}))$.
On entry, the vector x .

incx **INTEGER**.
The increment between elements of x . $\text{incx} > 0$.

Output Parameters

α On exit, it is overwritten with the value β .

x On exit, it is overwritten with the vector v .

`tau` REAL for `slarfg`
 DOUBLE PRECISION for `dlarfg`
 COMPLEX for `clarfg`
 COMPLEX*16 for `zlarfg`
 The value `tau`.

?larft

Forms the triangular factor T of a block reflector $H = I - VTV^H$.

```
call slarft ( direct, storev, n, k, v, ldv, tau, t, ldt )
call dlarft ( direct, storev, n, k, v, ldv, tau, t, ldt )
call clarft ( direct, storev, n, k, v, ldv, tau, t, ldt )
call zlarft ( direct, storev, n, k, v, ldv, tau, t, ldt )
```

Discussion

The routine `?larft` forms the triangular factor T of a real/complex block reflector H of order n , which is defined as a product of k elementary reflectors.

If `direct` = 'F', $H = H(1) H(2) \dots H(k)$ and T is upper triangular;

If `direct` = 'B', $H = H(k) \dots H(2) H(1)$ and T is lower triangular.

If `storev` = 'C', the vector which defines the elementary reflector $H(i)$ is stored in the i -th column of the array `v`, and $H = I - V^*T^*V'$.

If `storev` = 'R', the vector which defines the elementary reflector $H(i)$ is stored in the i -th row of the array `v`, and $H = I - V' *T*V$.

Input Parameters

`direct` CHARACTER*1. Specifies the order in which the elementary reflectors are multiplied to form the block reflector:
 = 'F': $H = H(1) H(2) \dots H(k)$ (forward)
 = 'B': $H = H(k) \dots H(2) H(1)$ (backward)

storev CHARACTER*1. Specifies how the vectors which define the elementary reflectors are stored (see also *Application Notes* below):
 = 'C': column-wise
 = 'R': row-wise.

n INTEGER. The order of the block reflector H . $n \geq 0$.

k INTEGER. The order of the triangular factor T (equal to the number of elementary reflectors). $k \geq 1$.

v REAL for *slarft*
 DOUBLE PRECISION for *dlarft*
 COMPLEX for *clarft*
 COMPLEX*16 for *zlarft*
 Array, DIMENSION
 (*ldv*, *k*) if *storev* = 'C' or
 (*ldv*, *n*) if *storev* = 'R'.
 The matrix V .

ldv INTEGER. The leading dimension of the array *v*.
 If *storev* = 'C', $ldv \geq \max(1, n)$;
 if *storev* = 'R', $ldv \geq k$.

tau REAL for *slarft*
 DOUBLE PRECISION for *dlarft*
 COMPLEX for *clarft*
 COMPLEX*16 for *zlarft*
 Array, DIMENSION (*k*). *tau*(*i*) must contain the scalar factor of the elementary reflector $H(i)$.

ldt INTEGER. The leading dimension of the output array *t*.
 $ldt \geq k$.

Output Parameters

t REAL for *slarft*
 DOUBLE PRECISION for *dlarft*
 COMPLEX for *clarft*
 COMPLEX*16 for *zlarft*
 Array, DIMENSION (*ldt*, *k*). The *k*-by-*k* triangular factor

T of the block reflector. If $direct = 'F'$, T is upper triangular; if $direct = 'B'$, T is lower triangular. The rest of the array is not used.

v The matrix V .

Application Notes

The shape of the matrix V and the storage of the vectors which define the $H(i)$ is best illustrated by the following example with $n = 5$ and $k = 3$. The elements equal to 1 are not stored; the corresponding array elements are modified but restored on exit. The rest of the array is not used.

$direct = 'F'$ and $storev = 'C'$: $direct = 'F'$ and $storev = 'R'$:

$$\begin{bmatrix} 1 \\ v_1 & 1 \\ v_1 & v_2 & 1 \\ v_1 & v_2 & v_3 \\ v_1 & v_2 & v_3 \end{bmatrix}$$

$$\begin{bmatrix} 1 & v_1 & v_1 & v_1 & v_1 \\ & 1 & v_2 & v_2 & v_2 \\ & & 1 & v_3 & v_3 \end{bmatrix}$$

$direct = 'B'$ and $storev = 'C'$: $direct = 'B'$ and $storev = 'R'$:

$$\begin{bmatrix} v_1 & v_2 & v_3 \\ v_1 & v_2 & v_3 \\ 1 & v_2 & v_3 \\ & 1 & v_3 \\ & & 1 \end{bmatrix}$$

$$\begin{bmatrix} v_1 & v_1 & 1 \\ v_2 & v_2 & v_2 & 1 \\ v_3 & v_3 & v_3 & v_3 & 1 \end{bmatrix}$$

?larfx

Applies an elementary reflector to a general rectangular matrix, with loop unrolling when the reflector has order ≤ 10 .

```
call slarfx ( side, m, n, v, tau, c, ldc, work )
call dlarfx ( side, m, n, v, tau, c, ldc, work )
call clarfx ( side, m, n, v, tau, c, ldc, work )
call zlarfx ( side, m, n, v, tau, c, ldc, work )
```

Discussion

The routine `?larfx` applies a real/complex elementary reflector H to a real/complex m -by- n matrix C , from either the left or the right.

H is represented in the form

$H = I - \tau * v * v'$, where τ is a real/complex scalar and v is a real/complex vector.

If $\tau = 0$, then H is taken to be the unit matrix

Input Parameters

side CHARACTER*1.
If **side** = 'L': form $H*C$
If **side** = 'R': form $C*H$.

m INTEGER. The number of rows of the matrix C .

n INTEGER. The number of columns of the matrix C .

v REAL for `slarfx`
DOUBLE PRECISION for `dlarfx`
COMPLEX for `clarfx`
COMPLEX*16 for `zlarfx`
Array, DIMENSION
(m) if **side** = 'L' or
(n) if **side** = 'R'.
The vector v in the representation of H .

<code>tau</code>	<p>REAL for <code>slarfx</code> DOUBLE PRECISION for <code>dlarfx</code> COMPLEX for <code>clarfx</code> COMPLEX*16 for <code>zlarfx</code> The value <code>tau</code> in the representation of H.</p>
<code>c</code>	<p>REAL for <code>slarfx</code> DOUBLE PRECISION for <code>dlarfx</code> COMPLEX for <code>clarfx</code> COMPLEX*16 for <code>zlarfx</code> Array, DIMENSION (<code>ldc</code>,<code>n</code>). On entry, the m-by-n matrix C.</p>
<code>ldc</code>	<p>INTEGER. The leading dimension of the array <code>c</code>. $lda \geq (1,m)$.</p>
<code>work</code>	<p>REAL for <code>slarfx</code> DOUBLE PRECISION for <code>dlarfx</code> COMPLEX for <code>clarfx</code> COMPLEX*16 for <code>zlarfx</code> Workspace array, DIMENSION (n) if <code>side = 'L'</code> or (m) if <code>side = 'R'</code>. <code>work</code> is not referenced if H has order < 11.</p>

Output Parameters

<code>c</code>	<p>On exit, <code>C</code> is overwritten by the matrix H^*C if <code>side = 'L'</code>, or C^*H if <code>side = 'R'</code>.</p>
----------------	--

?largv

Generates a vector of plane rotations
with real cosines and real/complex
sines.

```
call slargv ( n, x, incx, y, incy, c, incc )
```

```
call dlargv ( n, x, incx, y, incy, c, incc )
call clargv ( n, x, incx, y, incy, c, incc )
call zlargv ( n, x, incx, y, incy, c, incc )
```

Discussion

The routine generates a vector of real/complex plane rotations with real cosines, determined by elements of the real/complex vectors x and y .

For `slargv/dlargv`:

$$\begin{bmatrix} c(i) & s(i) \\ -s(i) & c(i) \end{bmatrix} \begin{bmatrix} x_i \\ y_i \end{bmatrix} = \begin{bmatrix} a_i \\ 0 \end{bmatrix}, \text{ for } i = 1, 2, \dots, n$$

For `clargv/zlargv`:

$$\begin{bmatrix} c(i) & s(i) \\ -\text{conjg}(s(i)) & c(i) \end{bmatrix} \begin{bmatrix} x_i \\ y_i \end{bmatrix} = \begin{bmatrix} r_i \\ 0 \end{bmatrix}, \text{ for } i = 1, 2, \dots, n$$

where $c(i)^2 + \text{abs}(s(i))^2 = 1$ and the following conventions are used (these are the same as in `clartg/zlartg` but differ from the BLAS Level 1 routine `crotg/zrotg`):

If $y_i = 0$, then $c(i) = 1$ and $s(i) = 0$;

If $x_i = 0$, then $c(i) = 0$ and $s(i)$ is chosen so that r_i is real.

Input Parameters

n **INTEGER.** The number of plane rotations to be generated.

x, y **REAL** for `slargv`
DOUBLE PRECISION for `dlargv`
COMPLEX for `clargv`
COMPLEX*16 for `zlargv`
 Arrays, **DIMENSION** $(1+(n-1)*incx)$ and $(1+(n-1)*incy)$, respectively.
 On entry, the vectors x and y .

$incx$ **INTEGER.** The increment between elements of x .
 $incx > 0$.

incy **INTEGER**. The increment between elements of *y*.
incy > 0.

incc **INTEGER**. The increment between elements of the
output array *c*. *incc* > 0.

Output Parameters

x On exit, *x*(*i*) is overwritten by *a*_{*i*} (for real flavors), or by
*r*_{*i*} (for complex flavors), for *i* = 1,...,*n*.

y On exit, the sines *s*(*i*) of the plane rotations.

c **REAL** for **slargv/clargv**
DOUBLE PRECISION for **dlargv/zlargv**
Array, **DIMENSION** (1+(*n*-1)**incc*). The cosines of the
plane rotations.

?larnv

Returns a vector of random numbers
from a uniform or normal distribution.

```
call slarnv ( idist, iseed, n, x )
call dlarnv ( idist, iseed, n, x )
call clarnv ( idist, iseed, n, x )
call zlarnv ( idist, iseed, n, x )
```

Discussion

The routine **?larnv** returns a vector of *n* random real/complex numbers from a uniform or normal distribution.

This routine calls the auxiliary routine **?laruv** to generate random real numbers from a uniform (0,1) distribution, in batches of up to 128 using vectorisable code. The Box-Muller method is used to transform numbers from a uniform to a normal distribution.

Input Parameters

- idist* **INTEGER**. Specifies the distribution of the random numbers:
for **slarnv** and **dlanrv**:
= 1: uniform (0,1)
= 2: uniform (-1,1)
= 3: normal (0,1).
for **clarnv** and **zlanrv**:
= 1: real and imaginary parts each uniform (0,1)
= 2: real and imaginary parts each uniform (-1,1)
= 3: real and imaginary parts each normal (0,1)
= 4: uniformly distributed on the disc $\text{abs}(z) < 1$
= 5: uniformly distributed on the circle $\text{abs}(z) = 1$
- iseed* **INTEGER**.
Array, **DIMENSION** (4).
On entry, the seed of the random number generator; the array elements must be between 0 and 4095, and *iseed*(4) must be odd.
- n* **INTEGER**. The number of random numbers to be generated.

Output Parameters

- x* **REAL** for **slarnv**
DOUBLE PRECISION for **dlanrv**
COMPLEX for **clarnv**
COMPLEX*16 for **zlanrv**
Array, **DIMENSION** (*n*). The generated random numbers.
- iseed* On exit, the seed is updated.

?larrb

Provides limited bisection to locate eigenvalues for more accuracy.

```
call slarrb ( n, d, l, ld, lld, ifirst, ilast, sigma,
             reltol, w, wgap, werr, work, iwork, info )
call dlarrb ( n, d, l, ld, lld, ifirst, ilast, sigma,
             reltol, w, wgap, werr, work, iwork, info )
```

Discussion

Given the relatively robust representation(RRR) LDL^T , the routine does “limited” bisection to locate the eigenvalues of LDL^T , $w(ifirst)$ through $w(ilast)$, to more accuracy. Intervals [*left*, *right*] are maintained by storing their mid-points and semi-widths in the arrays *w* and *werr* respectively.

Input Parameters

<i>n</i>	INTEGER. The order of the matrix.
<i>d</i>	REAL for <code>slarrb</code> DOUBLE PRECISION for <code>dlarrb</code> Array, DIMENSION (<i>n</i>). The <i>n</i> diagonal elements of the diagonal matrix <i>D</i> .
<i>l</i>	REAL for <code>slarrb</code> DOUBLE PRECISION for <code>dlarrb</code> Array, DIMENSION (<i>n</i> -1). The <i>n</i> -1 subdiagonal elements of the unit bidiagonal matrix <i>L</i> .
<i>ld</i>	REAL for <code>slarrb</code> DOUBLE PRECISION for <code>dlarrb</code> Array, DIMENSION (<i>n</i> -1). The <i>n</i> -1 elements $L_i * D_i$.
<i>lld</i>	REAL for <code>slarrb</code> DOUBLE PRECISION for <code>dlarrb</code> Array, DIMENSION (<i>n</i> -1). The <i>n</i> -1 elements $L_i * L_i * D_i$.
<i>ifirst</i>	INTEGER. The index of the first eigenvalue in the cluster.

<i>ilast</i>	INTEGER. The index of the last eigenvalue in the cluster.
<i>sigma</i>	REAL for <i>slarrb</i> DOUBLE PRECISION for <i>dlarrb</i> The shift used to form LDL^T (see <i>?larrf</i>).
<i>reltol</i>	REAL for <i>slarrb</i> DOUBLE PRECISION for <i>dlarrb</i> The relative tolerance.
<i>w</i>	REAL for <i>slarrb</i> DOUBLE PRECISION for <i>dlarrb</i> Array, DIMENSION (<i>n</i>). On input, <i>w(ifirst)</i> through <i>w(ilast)</i> are estimates of the corresponding eigenvalues of LDL^T .
<i>wgap</i>	REAL for <i>slarrb</i> DOUBLE PRECISION for <i>dlarrb</i> Array, DIMENSION (<i>n</i>). The gaps between the eigenvalues of LDL^T .
<i>werr</i>	REAL for <i>slarrb</i> DOUBLE PRECISION for <i>dlarrb</i> Array, DIMENSION (<i>n</i>). On input, <i>werr(ifirst)</i> through <i>werr(ilast)</i> are the errors in the estimates <i>w(ifirst)</i> through <i>w(ilast)</i> .
<i>work</i>	REAL for <i>slarrb</i> DOUBLE PRECISION for <i>dlarrb</i> Workspace array. Note that this parameter is never used in the routine.
<i>iwork</i>	INTEGER. Workspace array, DIMENSION ($2n$).

Output Parameters

<i>w</i>	On output these estimates of the eigenvalues are “refined”.
<i>wgap</i>	Very small gaps are changed on output.
<i>werr</i>	On output, “refined” errors in the estimates <i>w(ifirst)</i> through <i>w(ilast)</i> .

info

INTEGER.

Error flag. Note that this parameter is never set in the routine.

?larre

Given the tridiagonal matrix T , sets small off-diagonal elements to zero and for each unreduced block T_i , finds base representations and eigenvalues.

```
call slarre ( n, d, e, tol, nsplit, isplit, m, w, woff,
             gersch, work, info )
call dlarre ( n, d, e, tol, nsplit, isplit, m, w, woff,
             gersch, work, info )
```

Discussion

Given the tridiagonal matrix T , the routine sets "small" off-diagonal elements to zero, and for each unreduced block T_i , it finds

- the numbers σ_i
- the base $T_i - \sigma_i I = L_i D_i L_i^T$ representations and
- eigenvalues of each $L_i D_i L_i^T$.

The representations and eigenvalues found are then used by `?stegr` to compute the eigenvectors of a symmetric tridiagonal matrix. Currently, the base representations are limited to being positive or negative definite, and the eigenvalues of the definite matrices are found by the *dqds* algorithm (subroutine `?lasq2`). As an added benefit, `?larre` also outputs the n Gerschgorin intervals for each $L_i D_i L_i^T$.

Input Parameters

n INTEGER. The order of the matrix.

<i>d</i>	REAL for <i>slarre</i> DOUBLE PRECISION for <i>dlarre</i> Array, DIMENSION (<i>n</i>). On entry, the <i>n</i> diagonal elements of the tridiagonal matrix <i>T</i> .
<i>e</i>	REAL for <i>slarre</i> DOUBLE PRECISION for <i>dlarre</i> Array, DIMENSION (<i>n</i>). On entry, the (<i>n</i> -1) subdiagonal elements of the tridiagonal matrix <i>T</i> ; <i>e</i> (<i>n</i>) need not be set.
<i>tol</i>	REAL for <i>slarre</i> DOUBLE PRECISION for <i>dlarre</i> The threshold for splitting. If on input $ e(i) < tol$, then the matrix <i>T</i> is split into smaller blocks.
<i>nsplit</i>	INTEGER. The number of blocks <i>T</i> splits into. $1 \leq nsplit \leq n$.
<i>work</i>	REAL for <i>slarre</i> DOUBLE PRECISION for <i>dlarre</i> Workspace array, DIMENSION (4* <i>n</i>).

Output Parameters

<i>d</i>	On exit, the <i>n</i> diagonal elements of the diagonal matrices D_i .
<i>e</i>	On exit, the subdiagonal elements of the unit bidiagonal matrices L_i .
<i>isplit</i>	INTEGER. Array, DIMENSION (2 <i>n</i>). The splitting points, at which <i>T</i> breaks up into submatrices. The first submatrix consists of rows/columns 1 to <i>isplit</i> (1), the second of rows/columns <i>isplit</i> (1)+1 through <i>isplit</i> (2), etc., and the <i>nsplit</i> -th consists of rows/columns <i>isplit</i> (<i>nsplit</i> -1)+1 through <i>isplit</i> (<i>nsplit</i>)= <i>n</i> .
<i>m</i>	INTEGER. The total number of eigenvalues (of all the $L_i D_i L_i^T$) found.

<code>w</code>	REAL for <code>slarre</code> DOUBLE PRECISION for <code>dlarre</code> Array, DIMENSION (n). The first m elements contain the eigenvalues. The eigenvalues of each of the blocks, $L_i D_i L_i^T$, are sorted in ascending order.
<code>woff</code>	REAL for <code>slarre</code> DOUBLE PRECISION for <code>dlarre</code> Array, DIMENSION (n). The <code>nsplit</code> base points σ_j .
<code>gersch</code>	REAL for <code>slarre</code> DOUBLE PRECISION for <code>dlarre</code> Array, DIMENSION ($2n$). The n Gerschgorin intervals.
<code>info</code>	INTEGER. Output error code from <code>?lasq2</code> .

?larrf

Finds a new relatively robust representation such that at least one of the eigenvalues is relatively isolated.

```
call slarrf ( n, d, l, ld, lld, ifirst, ilast, w, dplus,
             lplus, work, iwork, info )
call dlarrf ( n, d, l, ld, lld, ifirst, ilast, w, dplus,
             lplus, work, iwork, info )
```

Discussion

Given the initial representation LDL^T and its cluster of close eigenvalues (in a relative measure), $w(\text{ifirst})$, $w(\text{ifirst}+1)$, ... $w(\text{ilast})$, the routine `?larrf` finds a new relatively robust representation

$$LDL^T - \sigma_j I = L(+)D(+)L(+)^T$$

such that at least one of the eigenvalues of $L(+)D(+)L(+)^T$ is relatively isolated.

Input Parameters

<i>n</i>	INTEGER. The order of the matrix.
<i>d</i>	REAL for <code>slarrf</code> DOUBLE PRECISION for <code>dlarrf</code> Array, DIMENSION (<i>n</i>). The <i>n</i> diagonal elements of the diagonal matrix <i>D</i> .
<i>l</i>	REAL for <code>slarrf</code> DOUBLE PRECISION for <code>dlarrf</code> Array, DIMENSION (<i>n</i> -1). The (<i>n</i> -1) subdiagonal elements of the unit bidiagonal matrix <i>L</i> .
<i>ld</i>	REAL for <code>slarrf</code> DOUBLE PRECISION for <code>dlarrf</code> Array, DIMENSION (<i>n</i> -1). The <i>n</i> -1 elements $L_i * D_i$.
<i>lld</i>	REAL for <code>slarrf</code> DOUBLE PRECISION for <code>dlarrf</code> Array, DIMENSION (<i>n</i> -1). The <i>n</i> -1 elements $L_i * L_i * D_i$.
<i>ifirst</i>	INTEGER. The index of the first eigenvalue in the cluster.
<i>ilast</i>	INTEGER. The index of the last eigenvalue in the cluster.
<i>w</i>	REAL for <code>slarrf</code> DOUBLE PRECISION for <code>dlarrf</code> Array, DIMENSION (<i>n</i>). On input, the eigenvalues of LDL^T in ascending order. $w(ifirst)$ through $w(ilast)$ form the cluster of relatively close eigenvalues.
<i>sigma</i>	REAL for <code>slarrf</code> DOUBLE PRECISION for <code>dlarrf</code> The shift used to form $L(+)D(+)L(+)^T$.
<i>work</i>	REAL for <code>slarrf</code> DOUBLE PRECISION for <code>dlarrf</code> Workspace array.

Output Parameters

<i>w</i>	On output, <i>w</i> (<i>ifirst</i>) through <i>w</i> (<i>ilast</i>) are estimates of the corresponding eigenvalues of $L(+)D(+)L(+)^T$.
<i>dplus</i>	REAL for <code>slarrf</code> DOUBLE PRECISION for <code>dlarrf</code> Array, DIMENSION (<i>n</i>). The <i>n</i> diagonal elements of the diagonal matrix $D(+)$.
<i>lplus</i>	REAL for <code>slarrf</code> DOUBLE PRECISION for <code>dlarrf</code> Array, DIMENSION (<i>n</i>). The first (<i>n</i> -1) elements of <i>lplus</i> contain the subdiagonal elements of the unit bidiagonal matrix $L(+)$. <i>lplus</i> (<i>n</i>) is set to <i>sigma</i> .

?larrv

Computes the eigenvectors of the tridiagonal matrix $T = L D L^T$ given L , D and the eigenvalues of $L D L^T$.

```
call slarrv ( n, d, l, isplit, m, w, iblock, gersch,
             tol, z, ldz, isuppz, work, iwork, info )
call dlarrv ( n, d, l, isplit, m, w, iblock, gersch,
             tol, z, ldz, isuppz, work, iwork, info )
call clarrv ( n, d, l, isplit, m, w, iblock, gersch,
             tol, z, ldz, isuppz, work, iwork, info )
call zlarrv ( n, d, l, isplit, m, w, iblock, gersch,
             tol, z, ldz, isuppz, work, iwork, info )
```

Discussion

The routine `?larrv` computes the eigenvectors of the tridiagonal matrix $T = L D L^T$ given L , D and the eigenvalues of $L D L^T$. The input eigenvalues should have high relative accuracy with respect to the entries of L and D . The desired accuracy of the output can be specified by the input parameter *tol*.

Input Parameters

- n* **INTEGER**. The order of the matrix. $n \geq 0$.
- d* **REAL** for **slarrv/clarrv**
DOUBLE PRECISION for **dlarrv/zlarrv**
 Array, **DIMENSION** (*n*). On entry, the *n* diagonal elements of the diagonal matrix *D*.
- l* **REAL** for **slarrv/clarrv**
DOUBLE PRECISION for **dlarrv/zlarrv**
 Array, **DIMENSION** (*n*-1). On entry, the (*n*-1) subdiagonal elements of the unit bidiagonal matrix *L* are contained in elements 1 to *n*-1 of *l*. *l*(*n*) need not be set.
- isplit* **INTEGER**.
 Array, **DIMENSION** (*n*). The splitting points, at which *T* breaks up into submatrices. The first submatrix consists of rows/columns 1 to *isplit*(1), the second of rows/columns *isplit*(1)+1 through *isplit*(2), etc.
- tol* **REAL** for **slarrv/clarrv**
DOUBLE PRECISION for **dlarrv/zlarrv**
 The absolute error tolerance for the eigenvalues/eigenvectors.
 Errors in the input eigenvalues must be bounded by *tol*. The eigenvectors output have residual norms bounded by *tol*, and the dot products between different eigenvectors are bounded by *tol*. *tol* must be at least $n * eps * |T|$, where *eps* is the machine precision and $|T|$ is the 1-norm of the tridiagonal matrix.
- m* **INTEGER**. The total number of eigenvalues found.
 $0 \leq m \leq n$. If *range* = 'A', $m = n$, and if *range* = 'I', $m = iu - il + 1$.
- w* **REAL** for **slarrv/clarrv**
DOUBLE PRECISION for **dlarrv/zlarrv**
 Array, **DIMENSION** (*n*). The first *m* elements of *w* contain the eigenvalues for which eigenvectors are to be computed. The eigenvalues should be grouped by

split-off block and ordered from smallest to largest within the block (The output array *w* from `?larre` is expected here). Errors in *w* must be bounded by *tol*.

<i>iblock</i>	<p>INTEGER.</p> <p>Array, DIMENSION (<i>n</i>). The submatrix indices associated with the corresponding eigenvalues in <i>w</i>; <i>iblock</i>(<i>i</i>)=1 if eigenvalue <i>w</i>(<i>i</i>) belongs to the first submatrix from the top, =2 if <i>w</i>(<i>i</i>) belongs to the second submatrix, etc.</p>
<i>ldz</i>	<p>INTEGER. The leading dimension of the output array <i>z</i>. $ldz \geq 1$, and if <i>jobz</i> = 'V', $ldz \geq \max(1,n)$.</p>
<i>work</i>	<p>REAL for <code>slarrv/clarrv</code> DOUBLE PRECISION for <code>dlarrv/zlarrv</code> Workspace array, DIMENSION (13<i>n</i>).</p>
<i>iwork</i>	<p>INTEGER. Workspace array, DIMENSION (6<i>n</i>).</p>

Output Parameters

<i>d</i>	On exit, <i>d</i> may be overwritten.
<i>l</i>	On exit, <i>l</i> is overwritten.
<i>z</i>	<p>REAL for <code>slarrv</code> DOUBLE PRECISION for <code>dlarrv</code> COMPLEX for <code>clarrv</code> COMPLEX*16 for <code>zlarrv</code> Array, DIMENSION (<i>ldz</i>, max(1,<i>m</i>)). If <i>jobz</i> = 'V', then if <i>info</i> = 0, the first <i>m</i> columns of <i>z</i> contain the orthonormal eigenvectors of the matrix <i>T</i> corresponding to the selected eigenvalues, with the <i>i</i>-th column of <i>z</i> holding the eigenvector associated</p>

with $w(i)$.
 If $jobz = 'N'$, then z is not referenced.



NOTE. The user must ensure that at least $\max(1,m)$ columns are supplied in the array z ; if $range = 'V'$, the exact value of m is not known in advance and an upper bound must be used.

isuppz **INTEGER** .
 Array, **DIMENSION** (2* $\max(1,m)$). The support of the eigenvectors in z , i.e., the indices indicating the nonzero elements in z . The i -th eigenvector is nonzero only in elements $isuppz(2i-1)$ through $isuppz(2i)$.

info **INTEGER**.
 If $info = 0$: successful exit
 If $info = -i < 0$: the i -th argument had an illegal value
 $info > 0$: if $info = 1$, there is an internal error in
 ?larrb;
 if $info = 2$, there is an internal error in ?stein.

?lartg

Generates a plane rotation with real cosine and real/complex sine.

```
call slartg ( f, g, cs, sn, r )
call dlartg ( f, g, cs, sn, r )
call clartg ( f, g, cs, sn, r )
call zlartg ( f, g, cs, sn, r )
```

Discussion

The routine generates a plane rotation so that

$$\begin{bmatrix} cs & sn \\ -\text{conjg}(sn) & cs \end{bmatrix} \cdot \begin{bmatrix} f \\ g \end{bmatrix} = \begin{bmatrix} r \\ 0 \end{bmatrix}$$

where $cs^2 + |sn|^2 = 1$

This is a slower, more accurate version of the BLAS Level 1 routine `?rotg`, except for the following differences.

For `slartg/dlartg`:

`f` and `g` are unchanged on return;

If `g=0`, then `cs=1` and `sn=0`;

If `f=0` and `g ≠ 0`, then `cs=0` and `sn=1` without doing any floating point operations (saves work in `?bdsqr` when there are zeros on the diagonal);

If `f` exceeds `g` in magnitude, `cs` will be positive.

For `clartg/zlartg`:

`f` and `g` are unchanged on return;

If `g=0`, then `cs=1` and `sn=0`;

If `f=0`, then `cs=0` and `sn` is chosen so that `r` is real.

Input Parameters

`f, g` REAL for `slartg`
 DOUBLE PRECISION for `dlartg`
 COMPLEX for `clartg`
 COMPLEX*16 for `zlartg`
 The first and second component of vector to be rotated.

Output Parameters

`cs` REAL for `slartg/clartg`
 DOUBLE PRECISION for `dlartg/zlartg`
 The cosine of the rotation.

sn REAL for slartg
 DOUBLE PRECISION for dlartg
 COMPLEX for clartg
 COMPLEX*16 for zartg
 The sine of the rotation.

r REAL for slartg
 DOUBLE PRECISION for dlartg
 COMPLEX for clartg
 COMPLEX*16 for zartg
 The nonzero component of the rotated vector.

?lartv

Applies a vector of plane rotations with real cosines and real/complex sines to the elements of a pair of vectors.

```

call slartv ( n, x, incx, y, incy, c, s, incc )
call dlartv ( n, x, incx, y, incy, c, s, incc )
call clartv ( n, x, incx, y, incy, c, s, incc )
call zartv ( n, x, incx, y, incy, c, s, incc )

```

Discussion

The routine applies a vector of real/complex plane rotations with real cosines to elements of the real/complex vectors x and y . For $i = 1, 2, \dots, n$

$$\begin{bmatrix} x_i \\ y_i \end{bmatrix} = \begin{bmatrix} c(i) & s(i) \\ -\text{conjg}(s(i)) & c(i) \end{bmatrix} \begin{bmatrix} x_i \\ y_i \end{bmatrix}$$

Input Parameters

n **INTEGER**. The number of plane rotations to be applied.

<i>x, y</i>	REAL for <code>slartv</code> DOUBLE PRECISION for <code>dlartv</code> COMPLEX for <code>clartv</code> COMPLEX*16 for <code>zlartv</code> Arrays, DIMENSION $(1+(n-1)*incx)$ and $(1+(n-1)*incy)$, respectively. The input vectors <i>x</i> and <i>y</i> .
<i>incx</i>	INTEGER. The increment between elements of <i>x</i> . <i>incx</i> > 0.
<i>incy</i>	INTEGER. The increment between elements of <i>y</i> . <i>incy</i> > 0.
<i>c</i>	REAL for <code>slartv/clartv</code> DOUBLE PRECISION for <code>dlartv/zlartv</code> Array, DIMENSION $(1+(n-1)*incc)$. The cosines of the plane rotations.
<i>s</i>	REAL for <code>slartv</code> DOUBLE PRECISION for <code>dlartv</code> COMPLEX for <code>clartv</code> COMPLEX*16 for <code>zlartv</code> Array, DIMENSION $(1+(n-1)*incc)$. The sines of the plane rotations.
<i>incc</i>	INTEGER. The increment between elements of <i>c</i> and <i>s</i> . <i>incc</i> > 0.

Output Parameters

x, y The rotated vectors *x* and *y*.

?laruv

Returns a vector of *n* random real numbers from a uniform distribution.

```
call slaruv ( iseed, n, x )
call dlaruv ( iseed, n, x )
```

Discussion

The routine `?laruv` returns a vector of n random real numbers from a uniform (0,1) distribution ($n \leq 28$).

This is an auxiliary routine called by `?larnv`.

Input Parameters

`iseed` **INTEGER**.
Array, **DIMENSION** (4). On entry, the seed of the random number generator; the array elements must be between 0 and 4095, and `iseed(4)` must be odd.

`n` **INTEGER**. The number of random numbers to be generated. $n \leq 28$.

Output Parameters

`x` **REAL** for `slaruv`
DOUBLE PRECISION for `dlaruv`
Array, **DIMENSION** (n). The generated random numbers.

`seed` On exit, the seed is updated.

?larz

Applies an elementary reflector (as returned by `?tzzrzf`) to a general matrix.

```
call slarz ( side, m, n, l, v, incv, tau, c, ldc, work )
call dlarz ( side, m, n, l, v, incv, tau, c, ldc, work )
call clarz ( side, m, n, l, v, incv, tau, c, ldc, work )
call zlarz ( side, m, n, l, v, incv, tau, c, ldc, work )
```

Discussion

The routine `?larz` applies a real/complex elementary reflector H to a real/complex m -by- n matrix C , from either the left or the right.

H is represented in the form

$$H = I - \tau * v * v',$$

where τ is a real/complex scalar and v is a real/complex vector.

If $\tau = 0$, then H is taken to be the unit matrix.

For complex flavors, to apply H' (the conjugate transpose of H), supply `conjg(τ)` instead of τ .

H is a product of k elementary reflectors as returned by `?tzzrzf`.

Input Parameters

<code>side</code>	<p>CHARACTER*1.</p> <p>If <code>side = 'L'</code>: form $H * C$</p> <p>If <code>side = 'R'</code>: form $C * H$</p>
<code>m</code>	INTEGER. The number of rows of the matrix C .
<code>n</code>	INTEGER. The number of columns of the matrix C .
<code>l</code>	<p>INTEGER. The number of entries of the vector v containing the meaningful part of the Householder vectors.</p> <p>If <code>side = 'L'</code>, $m \geq l \geq 0$, if <code>side = 'R'</code>, $n \geq l \geq 0$.</p>
<code>v</code>	<p>REAL for <code>slarz</code></p> <p>DOUBLE PRECISION for <code>dlarz</code></p> <p>COMPLEX for <code>clarz</code></p> <p>COMPLEX*16 for <code>zlarz</code></p> <p>Array, DIMENSION $(1+(l-1)*\text{abs}(\text{incv}))$. The vector v in the representation of H as returned by <code>?tzzrzf</code>. v is not used if $\tau = 0$.</p>
<code>incv</code>	<p>INTEGER. The increment between elements of v.</p> <p>$\text{incv} \neq 0$.</p>
<code>tau</code>	<p>REAL for <code>slarz</code></p> <p>DOUBLE PRECISION for <code>dlarz</code></p> <p>COMPLEX for <code>clarz</code></p> <p>COMPLEX*16 for <code>zlarz</code></p> <p>The value τ in the representation of H.</p>

c REAL for `slarz`
 DOUBLE PRECISION for `dlarz`
 COMPLEX for `clarz`
 COMPLEX*16 for `zlarz`
 Array, DIMENSION (*ldc*,*n*).
 On entry, the *m*-by-*n* matrix *C*.

ldc INTEGER. The leading dimension of the array *c*.
ldc ≥ max(1,*m*).

work REAL for `slarz`
 DOUBLE PRECISION for `dlarz`
 COMPLEX for `clarz`
 COMPLEX*16 for `zlarz`
 Workspace array, DIMENSION
 (*n*) if *side* = 'L' or
 (*m*) if *side* = 'R'.

Output Parameters

c On exit, *c* is overwritten by the matrix $H*C$ if *side* = 'L', or $C*H$ if *side* = 'R'.

?larzb

Applies a block reflector or its transpose/conjugate-transpose to a general matrix.

```

call slarzb ( side, trans, direct, storev, m, n, k, l,
             v, ldv, t, ldt, c, ldc, work, ldwork )
call dlarzb ( side, trans, direct, storev, m, n, k, l,
             v, ldv, t, ldt, c, ldc, work, ldwork )
call clarzb ( side, trans, direct, storev, m, n, k, l,
             v, ldv, t, ldt, c, ldc, work, ldwork )
call zlarzb ( side, trans, direct, storev, m, n, k, l,
             v, ldv, t, ldt, c, ldc, work, ldwork )

```

Discussion

The routine applies a real/complex block reflector H or its transpose H^T (or H^H for complex flavors) to a real/complex distributed m -by- n matrix C from the left or the right.

Currently, only `storev = 'R'` and `direct = 'B'` are supported.

Input Parameters

<code>side</code>	<p>CHARACTER*1.</p> <p>If <code>side = 'L'</code>: apply H or H' from the left</p> <p>If <code>side = 'R'</code>: apply H or H' from the right</p>
<code>trans</code>	<p>CHARACTER*1.</p> <p>If <code>trans = 'N'</code>: apply H (No transpose)</p> <p>If <code>trans='C'</code>: apply H' (Transpose/conjugate transpose)</p>
<code>direct</code>	<p>CHARACTER*1. Indicates how H is formed from a product of elementary reflectors</p> <p>= 'F': $H = H(1) H(2) \dots H(k)$ (forward, not supported yet)</p> <p>= 'B': $H = H(k) \dots H(2) H(1)$ (backward)</p>
<code>storev</code>	<p>CHARACTER*1. Indicates how the vectors which define the elementary reflectors are stored:</p> <p>= 'C': Column-wise (not supported yet)</p> <p>= 'R': Row-wise.</p>
<code>m</code>	INTEGER. The number of rows of the matrix C .
<code>n</code>	INTEGER. The number of columns of the matrix C .
<code>k</code>	INTEGER. The order of the matrix T (equal to the number of elementary reflectors whose product defines the block reflector).
<code>l</code>	<p>INTEGER. The number of columns of the matrix V containing the meaningful part of the Householder reflectors.</p> <p>If <code>side = 'L'</code>, $m \geq l \geq 0$, if <code>side = 'R'</code>, $n \geq l \geq 0$.</p>
<code>v</code>	<p>REAL for <code>slarzb</code></p> <p>DOUBLE PRECISION for <code>dlarzb</code></p> <p>COMPLEX for <code>clarzb</code></p>

COMPLEX*16 for `zlarzb`
 Array, DIMENSION (ldv, nv).
 If `storev = 'C'`, $nv = k$; if `storev = 'R'`, $nv = 1$.

`ldv` INTEGER. The leading dimension of the array `v`.
 If `storev = 'C'`, $ldv \geq 1$; if `storev = 'R'`, $ldv \geq k$.

`t` REAL for `slarzb`
 DOUBLE PRECISION for `dlarzb`
 COMPLEX for `clarzb`
 COMPLEX*16 for `zlarzb`
 Array, DIMENSION (ldt, k). The triangular k -by- k matrix T in the representation of the block reflector.

`ldt` INTEGER. The leading dimension of the array `t`.
 $ldt \geq k$.

`c` REAL for `slarzb`
 DOUBLE PRECISION for `dlarzb`
 COMPLEX for `clarzb`
 COMPLEX*16 for `zlarzb`
 Array, DIMENSION (ldc, n). On entry, the m -by- n matrix C .

`ldc` INTEGER. The leading dimension of the array `c`.
 $ldc \geq \max(1, m)$.

`work` REAL for `slarzb`
 DOUBLE PRECISION for `dlarzb`
 COMPLEX for `clarzb`
 COMPLEX*16 for `zlarzb`
 Workspace array, DIMENSION ($ldwork, k$).

`ldwork` INTEGER. The leading dimension of the array `work`.
 If `side = 'L'`, $ldwork \geq \max(1, n)$;
 if `side = 'R'`, $ldwork \geq \max(1, m)$.

Output Parameters

`c` On exit, `c` is overwritten by H^*C or $H'*C$ or C^*H or C^*H' .

?larzt

Forms the triangular factor T of a block reflector $H = I - VTV^H$.

```
call slarzt ( direct, storev, n, k, v, ldv, tau, t, ldt )
call dlarzt ( direct, storev, n, k, v, ldv, tau, t, ldt )
call clarzt ( direct, storev, n, k, v, ldv, tau, t, ldt )
call zlarzt ( direct, storev, n, k, v, ldv, tau, t, ldt )
```

Discussion

The routine forms the triangular factor T of a real/complex block reflector H of order $> n$, which is defined as a product of k elementary reflectors.

If $direct = 'F'$, $H = H(1) H(2) \dots H(k)$ and T is upper triangular.

If $direct = 'B'$, $H = H(k) \dots H(2) H(1)$ and T is lower triangular.

If $storev = 'C'$, the vector which defines the elementary reflector $H(i)$ is stored in the i -th column of the array v , and

$$H = I - V^*T^*V'$$

If $storev = 'R'$, the vector which defines the elementary reflector $H(i)$ is stored in the i -th row of the array v , and

$$H = I - V'^*T^*V$$

Currently, only $storev = 'R'$ and $direct = 'B'$ are supported.

Input Parameters

direct CHARACTER*1. Specifies the order in which the elementary reflectors are multiplied to form the block reflector:
 If $direct = 'F'$: $H = H(1) H(2) \dots H(k)$ (forward, not supported yet)
 If $direct = 'B'$: $H = H(k) \dots H(2) H(1)$ (backward)

storev CHARACTER*1. Specifies how the vectors which define the elementary reflectors are stored (see also *Application Notes* below):
 If $storev = 'C'$: column-wise (not supported yet)
 If $storev = 'R'$: row-wise

<i>n</i>	INTEGER . The order of the block reflector H . $n \geq 0$.
<i>k</i>	INTEGER . The order of the triangular factor T (equal to the number of elementary reflectors). $k \geq 1$.
<i>v</i>	REAL for slarzt DOUBLE PRECISION for dlarzt COMPLEX for clarzt COMPLEX*16 for zlarzt Array, DIMENSION (<i>ldv</i> , <i>k</i>) if <i>storev</i> = 'C' (<i>ldv</i> , <i>n</i>) if <i>storev</i> = 'R' The matrix V .
<i>ldv</i>	INTEGER . The leading dimension of the array <i>v</i> . If <i>storev</i> = 'C', $ldv \geq \max(1, n)$; if <i>storev</i> = 'R', $ldv \geq k$.
<i>tau</i>	REAL for slarzt DOUBLE PRECISION for dlarzt COMPLEX for clarzt COMPLEX*16 for zlarzt Array, DIMENSION (<i>k</i>). <i>tau</i> (<i>i</i>) must contain the scalar factor of the elementary reflector $H(i)$.
<i>ldt</i>	INTEGER . The leading dimension of the output array <i>t</i> . $ldt \geq k$.

Output Parameters

<i>t</i>	REAL for slarzt DOUBLE PRECISION for dlarzt COMPLEX for clarzt COMPLEX*16 for zlarzt Array, DIMENSION (<i>ldt</i> , <i>k</i>). The k -by- k triangular factor T of the block reflector. If <i>direct</i> = 'F', T is upper triangular; if <i>direct</i> = 'B', T is lower triangular. The rest of the array is not used.
<i>v</i>	The matrix V . See <i>Application Notes</i> below.

Application Notes

The shape of the matrix V and the storage of the vectors which define the $H(i)$ is best illustrated by the following example with $n = 5$ and $k = 3$. The elements equal to 1 are not stored; the corresponding array elements are modified but restored on exit. The rest of the array is not used.

direct = 'F' and *storev* = 'C': *direct* = 'F' and *storev* = 'R':

$$V = \begin{bmatrix} v_1 & v_2 & v_3 \\ v_1 & v_2 & v_3 \\ v_1 & v_2 & v_3 \\ v_1 & v_2 & v_3 \\ v_1 & v_2 & v_3 \\ \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \\ 1 & \cdot & \cdot \\ & 1 & \cdot \\ & & 1 \end{bmatrix}$$

$$\begin{array}{c} \text{---}V\text{---} \\ / \quad \quad \backslash \\ \left[\begin{array}{cccccccc} v_1 & v_1 & v_1 & v_1 & v_1 & \cdot & \cdot & \cdot & 1 \\ v_2 & v_2 & v_2 & v_2 & v_2 & \cdot & \cdot & \cdot & 1 \\ v_3 & v_3 & v_3 & v_3 & v_3 & \cdot & \cdot & \cdot & 1 \end{array} \right] \end{array}$$

direct = 'B' and *storev* = 'C': *direct* = 'B' and *storev* = 'R':

$$V = \begin{bmatrix} 1 \\ \cdot & 1 \\ \cdot & \cdot & 1 \\ \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \\ v_1 & v_2 & v_3 \\ v_1 & v_2 & v_3 \\ v_1 & v_2 & v_3 \\ v_1 & v_2 & v_3 \\ v_1 & v_2 & v_3 \end{bmatrix}$$

$$\begin{array}{c} \text{---}V\text{---} \\ / \quad \quad \backslash \\ \left[\begin{array}{cccccccc} 1 & \cdot & \cdot & \cdot & \cdot & v_1 & v_1 & v_1 & v_1 & v_1 \\ \cdot & 1 & \cdot & \cdot & \cdot & v_2 & v_2 & v_2 & v_2 & v_2 \\ \cdot & \cdot & 1 & \cdot & \cdot & v_3 & v_3 & v_3 & v_3 & v_3 \end{array} \right] \end{array}$$

?las2

Computes singular values of a 2-by-2 triangular matrix.

```
call slas2 ( f, g, h, ssmín, ssmáx )
call dlas2 ( f, g, h, ssmín, ssmáx )
```

Discussion

The routine `?las2` computes the singular values of the 2-by-2 matrix

$$\begin{bmatrix} f & g \\ 0 & h \end{bmatrix}$$

On return, `ssmín` is the smaller singular value and `ssmáx` is the larger singular value.

Input Parameters

`f, g, h` REAL for `slas2`
 DOUBLE PRECISION for `dlas2`
The (1,1), (1,2) and (2,2) elements of the 2-by-2 matrix, respectively.

Output Parameters

`ssmín, ssmáx` REAL for `slas2`
 DOUBLE PRECISION for `dlas2`
The smaller and the larger singular values, respectively.

Application Notes

Barring over/underflow, all output quantities are correct to within a few units in the last place (*ulps*), even in the absence of a guard digit in addition/subtraction.

In IEEE arithmetic, the code works correctly if one matrix element is infinite.

Overflow will not occur unless the largest singular value itself overflows, or is within a few ulps of overflow. (On machines with partial overflow, like the Cray, overflow may occur if the largest singular value is within a factor of 2 of overflow.)

Underflow is harmless if underflow is gradual. Otherwise, results may correspond to a matrix modified by perturbations of size near the underflow threshold.

?lascl

Multiplies a general rectangular matrix by a real scalar defined as c_{to}/c_{from}

```
call slascl ( type, kl, ku, cfrom, cto, m, n, a, lda, info )
call dlascl ( type, kl, ku, cfrom, cto, m, n, a, lda, info )
call clascl ( type, kl, ku, cfrom, cto, m, n, a, lda, info )
call zlascl ( type, kl, ku, cfrom, cto, m, n, a, lda, info )
```

Discussion

The routine `?lascl` multiplies the m -by- n real/complex matrix A by the real scalar $cto/cfrom$. The operation is performed without over/underflow as long as the final result $cto*A(i,j)/cfrom$ does not over/underflow.

`type` specifies that A may be full, upper triangular, lower triangular, upper Hessenberg, or banded.

Input Parameters

`type` CHARACTER*1. `type` indices the storage `type` of the input matrix.

- = 'G': A is a full matrix.
- = 'L': A is a lower triangular matrix.
- = 'U': A is an upper triangular matrix.
- = 'H': A is an upper Hessenberg matrix.
- = 'B': A is a symmetric band matrix with lower bandwidth kl and upper bandwidth ku and with the

only the lower half stored
 = 'Q': A is a symmetric band matrix with lower bandwidth kl and upper bandwidth ku and with the only the upper half stored.
 = 'Z': A is a band matrix with lower bandwidth kl and upper bandwidth ku .

kl **INTEGER**. The lower bandwidth of A . Referenced only if *type* = 'B', 'Q' or 'Z'.

ku **INTEGER**. The upper bandwidth of A . Referenced only if *type* = 'B', 'Q' or 'Z'.

cfrom, cto **REAL** for `slascl/clascl`
 DOUBLE PRECISION for `dlascl/zlascl`
 The matrix A is multiplied by *cto/cfrom*. $A(i,j)$ is computed without over/underflow if the final result $cto * A(i,j) / cfrom$ can be represented without over/underflow. *cfrom* must be nonzero.

m **INTEGER**. The number of rows of the matrix A . $m \geq 0$.

n **INTEGER**. The number of columns of the matrix A . $n \geq 0$.

a **REAL** for `slascl`
 DOUBLE PRECISION for `dlascl`
 COMPLEX for `clascl`
 COMPLEX*16 for `zlascl`
 Array, **DIMENSION** (*lda, m*). The matrix to be multiplied by *cto/cfrom*. See *type* for the storage type.

lda **INTEGER**. The leading dimension of the array *a*.
 $lda \geq \max(1,m)$.

Output Parameters

a The multiplied matrix A .

info **INTEGER**.
 If *info* = 0 - successful exit
 If *info* = -*i* < 0, the *i*-th argument had an illegal value.

?lasd0

Computes the singular values of a real upper bidiagonal n -by- m matrix B with diagonal d and off-diagonal e .

Used by ?bdsdc.

```
call slasd0 ( n, sqre, d, e, u, ldu, vt, ldvt, smlsiz,
             iwork, work, info )
call dlasd0 ( n, sqre, d, e, u, ldu, vt, ldvt, smlsiz,
             iwork, work, info )
```

Discussion

Using a divide and conquer approach, the routine ?lasd0 computes the singular value decomposition (SVD) of a real upper bidiagonal n -by- m matrix B with diagonal d and offdiagonal e , where $m = n + sqre$.

The algorithm computes orthogonal matrices U and VT such that $B = U*S*VT$. The singular values S are overwritten on d .

A related subroutine, ?lasda, computes only the singular values, and optionally, the singular vectors in compact form.

Input Parameters

n **INTEGER**. On entry, the row dimension of the upper bidiagonal matrix. This is also the dimension of the main diagonal array d .

sqre **INTEGER**. Specifies the column dimension of the bidiagonal matrix.
 If $sqre = 0$: The bidiagonal matrix has column dimension $m = n$;
 If $sqre = 1$: The bidiagonal matrix has column dimension $m = n+1$;

<i>d</i>	REAL for <code>slasd0</code> DOUBLE PRECISION for <code>dlasd0</code> Array, DIMENSION (<i>n</i>). On entry, <i>d</i> contains the main diagonal of the bidiagonal matrix.
<i>e</i>	REAL for <code>slasd0</code> DOUBLE PRECISION for <code>dlasd0</code> Array, DIMENSION (<i>m</i> -1). Contains the subdiagonal entries of the bidiagonal matrix. On exit, <i>e</i> is destroyed.
<i>ldu</i>	INTEGER. On entry, leading dimension of the output array <i>u</i> .
<i>ldvt</i>	INTEGER. On entry, leading dimension of the output array <i>vt</i> .
<i>smlsiz</i>	INTEGER. On entry, maximum size of the subproblems at the bottom of the computation tree.
<i>iwork</i>	INTEGER. Workspace array, DIMENSION must be at least $(8n)$.
<i>work</i>	REAL for <code>slasd0</code> DOUBLE PRECISION for <code>dlasd0</code> Workspace array, DIMENSION must be at least $(3m^2 + 2m)$.

Output Parameters

<i>d</i>	On exit <i>d</i> , if <i>info</i> = 0, contains singular values of the bidiagonal matrix.
<i>u</i>	REAL for <code>slasd0</code> DOUBLE PRECISION for <code>dlasd0</code> Array, DIMENSION at least (<i>ldu</i> , <i>n</i>). On exit, <i>u</i> contains the left singular vectors.
<i>vt</i>	REAL for <code>slasd0</code> DOUBLE PRECISION for <code>dlasd0</code> Array, DIMENSION at least (<i>ldvt</i> , <i>m</i>). On exit, <i>vt</i> ' contains the right singular vectors.

info INTEGER.
 If *info* = 0: successful exit.
 If *info* = -*i* < 0, the *i*-th argument had an illegal value.
 If *info* = 1, an singular value did not converge.

?lasd1

Computes the SVD of an upper bidiagonal matrix *B* of the specified size. Used by ?bdsdc.

```
call slasd1 ( nl, nr, sqre, d, alpha, beta, u, ldu, vt,
             ldvt, idxq, iwork, work, info )
call dlasd1 ( nl, nr, sqre, d, alpha, beta, u, ldu, vt,
             ldvt, idxq, iwork, work, info )
```

Discussion

This routine computes the SVD of an upper bidiagonal *n*-by-*m* matrix *B*, where $n = nl + nr + 1$ and $m = n + sqre$. The routine ?lasd1 is called from ?lasd0.

A related subroutine ?lasd7 handles the case in which the singular values (and the singular vectors in factored form) are desired.

?lasd1 computes the SVD as follows:

$$\begin{aligned}
 B &= U(in) * \begin{bmatrix} D1(in) & 0 & 0 & 0 \\ Z1' & a & Z2' & b \\ 0 & 0 & D2(in) & 0 \end{bmatrix} * VT(in) \\
 &= U(out) * (D(out) \ 0) * VT(out)
 \end{aligned}$$

where $Z = (Z1' \ a \ Z2' \ b) = u' \ VT'$, and *u* is a vector of dimension *m* with *alpha* and *beta* in the *nl*+1 and *nl*+2 -th entries and zeros elsewhere; and the entry *b* is empty if *sqre* = 0.

The left singular vectors of the original matrix are stored in u , and the transpose of the right singular vectors are stored in vt , and the singular values are in d . The algorithm consists of three stages:

The first stage consists of deflating the size of the problem when there are multiple singular values or when there are zeros in the Z vector. For each such occurrence the dimension of the secular equation problem is reduced by one. This stage is performed by the routine `?lasd2`.

The second stage consists of calculating the updated singular values. This is done by finding the square roots of the roots of the secular equation via the routine `?lasd4` (as called by `?lasd3`). This routine also calculates the singular vectors of the current problem.

The final stage consists of computing the updated singular vectors directly using the updated singular values. The singular vectors for the current problem are multiplied with the singular vectors from the overall problem.

Input Parameters

nl **INTEGER**. The row dimension of the upper block.
nl ≥ 1.

nr **INTEGER**. The row dimension of the lower block.
nr ≥ 1.

sqre **INTEGER**.
 If *sqre* = 0: the lower block is an *nr*-by-*nr* square matrix.
 If *sqre* = 1: the lower block is an *nr*-by-(*nr*+1) rectangular matrix. The bidiagonal matrix has row dimension $n = nl + nr + 1$, and column dimension $m = n + sqre$.

d **REAL** for `slasd1`
 DOUBLE PRECISION for `dlasd1`
 Array, **DIMENSION** ($n = nl + nr + 1$). On entry $d(1:nl, 1:nl)$ contains the singular values of the upper block; and $d(nl+2:n)$ contains the singular values of the lower block.

<i>alpha</i>	<p>REAL for <code>slasd1</code> DOUBLE PRECISION for <code>dlasd1</code> Contains the diagonal element associated with the added row.</p>
<i>beta</i>	<p>REAL for <code>slasd1</code> DOUBLE PRECISION for <code>dlasd1</code> Contains the off-diagonal element associated with the added row.</p>
<i>u</i>	<p>REAL for <code>slasd1</code> DOUBLE PRECISION for <code>dlasd1</code> Array, DIMENSION (<i>ldu</i>, <i>n</i>). On entry <code>u(1:nl, 1:nl)</code> contains the left singular vectors of the upper block; <code>u(nl+2:n, nl+2:n)</code> contains the left singular vectors of the lower block.</p>
<i>ldu</i>	<p>INTEGER. The leading dimension of the array <code>u</code>. $ldu \geq \max(1, n)$.</p>
<i>vt</i>	<p>REAL for <code>slasd1</code> DOUBLE PRECISION for <code>dlasd1</code> Array, DIMENSION (<i>ldvt</i>, <i>m</i>), where $m = n + sqre$. On entry <code>vt(1:nl+1, 1:nl+1)</code> contains the right singular vectors of the upper block; <code>vt(nl+2:m, nl+2:m)</code> contains the right singular vectors of the lower block.</p>
<i>ldvt</i>	<p>INTEGER. The leading dimension of the array <code>vt</code>. $ldvt \geq \max(1, m)$.</p>
<i>iwork</i>	<p>INTEGER. Workspace array, DIMENSION ($4n$).</p>
<i>work</i>	<p>REAL for <code>slasd1</code> DOUBLE PRECISION for <code>dlasd1</code> Workspace array, DIMENSION ($3m^2 + 2m$).</p>

Output Parameters

<i>d</i>	<p>On exit <code>d(1:n)</code> contains the singular values of the modified matrix.</p>
----------	---

<i>u</i>	On exit <i>u</i> contains the left singular vectors of the bidiagonal matrix.
<i>vt</i>	On exit <i>vt'</i> contains the right singular vectors of the bidiagonal matrix.
<i>idxq</i>	INTEGER Array, DIMENSION (<i>n</i>). Contains the permutation which will reintegrate the subproblem just solved back into sorted order, that is, $d(\text{idxq}(i = 1, n))$ will be in ascending order.
<i>info</i>	INTEGER . If <i>info</i> = 0: successful exit. If <i>info</i> = - <i>i</i> < 0, the <i>i</i> -th argument had an illegal value. If <i>info</i> = 1, an singular value did not converge.

?lasd2

Merges the two sets of singular values together into a single sorted set.

Used by ?bdsdc.

```
call slasd2 ( nl, nr, sqre, k, d, z, alpha, beta, u, ldu,
             vt, ldvt, dsigma, u2, ldu2, vt2, ldvt2,
             idxp, idx, idxc, idxq, coltyp, info )
call dlasd2 ( nl, nr, sqre, k, d, z, alpha, beta, u, ldu,
             vt, ldvt, dsigma, u2, ldu2, vt2, ldvt2,
             idxp, idx, idxc, idxq, coltyp, info )
```

Discussion

The routine ?lasd2 merges the two sets of singular values together into a single sorted set. Then it tries to deflate the size of the problem. There are two ways in which deflation can occur: when two or more singular values are close together or if there is a tiny entry in the Z vector. For each such occurrence the order of the related secular equation problem is reduced by one.

The routine `?lasd2` is called from `?lasd1`.

Input Parameters

<i>nl</i>	INTEGER . The row dimension of the upper block. $nl \geq 1$.
<i>nr</i>	INTEGER . The row dimension of the lower block. $nr \geq 1$.
<i>sqre</i>	INTEGER . If <i>sqre</i> = 0: the lower block is an <i>nr</i> -by- <i>nr</i> square matrix If <i>sqre</i> = 1: the lower block is an <i>nr</i> -by- $(nr+1)$ rectangular matrix. The bidiagonal matrix has $n = nl + nr + 1$ rows and $m = n + sqre \geq n$ columns.
<i>d</i>	REAL for <code>slasd2</code> DOUBLE PRECISION for <code>dlasd2</code> Array, DIMENSION (<i>n</i>). On entry <i>d</i> contains the singular values of the two submatrices to be combined.
<i>alpha</i>	REAL for <code>slasd2</code> DOUBLE PRECISION for <code>dlasd2</code> Contains the diagonal element associated with the added row.
<i>beta</i>	REAL for <code>slasd2</code> DOUBLE PRECISION for <code>dlasd2</code> Contains the off-diagonal element associated with the added row.
<i>u</i>	REAL for <code>slasd2</code> DOUBLE PRECISION for <code>dlasd2</code> Array, DIMENSION (<i>ldu</i> , <i>n</i>). On entry <i>u</i> contains the left singular vectors of two submatrices in the two square blocks with corners at (1,1), (<i>nl</i> , <i>nl</i>), and (<i>nl</i> +2, <i>nl</i> +2), (<i>n</i> , <i>n</i>).
<i>ldu</i>	INTEGER . The leading dimension of the array <i>u</i> . $ldu \geq n$.

ldu2 **INTEGER**. The leading dimension of the output array *u2*.
ldu2 $\geq n$.

vt **REAL** for *slasd2*
DOUBLE PRECISION for *dlsasd2*
Array, **DIMENSION** (*ldvt*, *m*). On entry *vt*' contains the right singular vectors of two submatrices in the two square blocks with corners at (1,1), (*n1*+1, *n1*+1), and (*n1*+2, *n1*+2), (*m*,*m*).

ldvt **INTEGER**. The leading dimension of the array *vt*.
ldvt $\geq m$.

ldvt2 **INTEGER**. The leading dimension of the output array *vt2*. *ldvt2* $\geq m$.

idxp **INTEGER**.
Workspace array, **DIMENSION** (*n*). This will contain the permutation used to place deflated values of *d* at the end of the array. On output *idxp*(2:*k*) points to the nondeflated *d*-values and *idxp*(*k*+1:*n*) points to the deflated singular values.

idx **INTEGER**.
Workspace array, **DIMENSION** (*n*). This will contain the permutation used to sort the contents of *d* into ascending order.

coltyp **INTEGER**.
Workspace array, **DIMENSION** (*n*). As workspace, this will contain a label which will indicate which of the following types a column in the *u2* matrix or a row in the *vt2* matrix is:
1 : non-zero in the upper half only
2 : non-zero in the lower half only
3 : dense
4 : deflated.

idxq **INTEGER**.
Array, **DIMENSION** (*n*). This contains the permutation which separately sorts the two sub-problems in *d* into ascending order. Note that entries in the first half of this

permutation must first be moved one position backward; and entries in the second half must first have $n1+1$ added to their values.

Output Parameters

- k* **INTEGER**. Contains the dimension of the non-deflated matrix, This is the order of the related secular equation. $1 \leq k \leq n$.
- d* On exit *d* contains the trailing ($n-k$) updated singular values (those which were deflated) sorted into increasing order.
- u* On exit *u* contains the trailing ($n-k$) updated left singular vectors (those which were deflated) in its last $n-k$ columns.
- z* **REAL** for **slasd2**
DOUBLE PRECISION for **dlsasd2**
Array, **DIMENSION** (n). On exit *z* contains the updating row vector in the secular equation.
- dsigma* **REAL** for **slasd2**
DOUBLE PRECISION for **dlsasd2**
Array, **DIMENSION** (n). Contains a copy of the diagonal elements ($k-1$ singular values and one zero) in the secular equation.
- u2* **REAL** for **slasd2**
DOUBLE PRECISION for **dlsasd2**
Array, **DIMENSION** ($ldu2, n$). Contains a copy of the first $k-1$ left singular vectors which will be used by **?lasd3** in a matrix multiply (**?gemm**) to solve for the new left singular vectors. *u2* is arranged into four blocks. The first block contains a column with 1 at $n1+1$ and zero everywhere else; the second block contains non-zero entries only at and above $n1$; the third contains non-zero entries only below $n1+1$; and the fourth is dense.

- vt* On exit *vt*' contains the trailing (*n-k*) updated right singular vectors (those which were deflated) in its last *n-k* columns. In case *sgre* =1, the last row of *vt* spans the right null space.
- vt2* REAL for *slasd2*
DOUBLE PRECISION for *dlasd2*
Array, DIMENSION (*ldvt2*, *n*). *vt2*' contains a copy of the first *k* right singular vectors which will be used by *?lasd3* in a matrix multiply (*?gemm*) to solve for the new right singular vectors. *vt2* is arranged into three blocks. The first block contains a row that corresponds to the special 0 diagonal element in *sigma*; the second block contains non-zeros only at and before *nl* +1; the third block contains non-zeros only at and after *nl* +2.
- idxc* INTEGER.
Array, DIMENSION (*n*). This will contain the permutation used to arrange the columns of the deflated *U* matrix into three groups: the first group contains non-zero entries only at and above *nl*, the second contains non-zero entries only below *nl*+2, and the third is dense.
- coltyp* On exit, it is an array of dimension 4, with *coltyp*(*i*) being the dimension of the *i*-th type columns.
- info* INTEGER.
If *info* = 0: successful exit
If *info* = -*i* < 0, the *i*-th argument had an illegal value.

?lasd3

Finds all square roots of the roots of the secular equation, as defined by the values in D and Z , and then updates the singular vectors by matrix multiplication. Used by ?bdsdc.

```
call slasd3 ( nl, nr, sqre, k, d, q, ldq, dsigma, u, ldu,
             u2, ldu2, vt, ldvt, vt2, ldvt2, idxc, ctot,
             z, info )
call dlasd3 ( nl, nr, sqre, k, d, q, ldq, dsigma, u, ldu,
             u2, ldu2, vt, ldvt, vt2, ldvt2, idxc, ctot,
             z, info )
```

Discussion

The routine ?lasd3 finds all the square roots of the roots of the secular equation, as defined by the values in D and Z . It makes the appropriate calls to ?lasd4 and then updates the singular vectors by matrix multiplication.

The routine ?lasd3 is called from ?lasd1.

Input Parameters

nl **INTEGER.** The row dimension of the upper block.
nl ≥ 1.

nr **INTEGER.** The row dimension of the lower block.
nr ≥ 1.

sqre **INTEGER.**
 If *sqre* = 0: the lower block is an *nr*-by-*nr* square matrix.
 If *sqre* = 1: the lower block is an *nr*-by-(*nr*+1) rectangular matrix. The bidiagonal matrix has $n = nl + nr + 1$ rows and $m = n + sqre ≥ n$ columns.

k **INTEGER.** The size of the secular equation, $1 ≤ k ≤ n$.

q REAL for `slasd3`
DOUBLE PRECISION for `dlasd3`
Workspace array, DIMENSION at least (ldq, k).

ldq INTEGER. The leading dimension of the array *q*.
 $ldq \geq k$.

dsigma REAL for `slasd3`
DOUBLE PRECISION for `dlasd3`
Array, DIMENSION (k). The first k elements of this array contain the old roots of the deflated updating problem. These are the poles of the secular equation.

u REAL for `slasd3`
DOUBLE PRECISION for `dlasd3`
Array, DIMENSION (ldu, n). The last $n - k$ columns of this matrix contain the deflated left singular vectors.

ldu INTEGER. The leading dimension of the array *u*.
 $ldu \geq n$.

u2 REAL for `slasd3`
DOUBLE PRECISION for `dlasd3`
Array, DIMENSION ($ldu2, n$). The first k columns of this matrix contain the non-deflated left singular vectors for the split problem.

ldu2 INTEGER. The leading dimension of the array *u2*.
 $ldu2 \geq n$.

vt REAL for `slasd3`
DOUBLE PRECISION for `dlasd3`
Array, DIMENSION ($ldvt, m$). The last $m - k$ columns of *vt'* contain the deflated right singular vectors.

ldvt INTEGER. The leading dimension of the array *vt*.
 $ldvt \geq n$.

vt2 REAL for `slasd3`
DOUBLE PRECISION for `dlasd3`
Array, DIMENSION ($ldvt2, n$). The first k columns of *vt2'* contain the non-deflated right singular vectors for the split problem.

<i>ldvt2</i>	INTEGER. The leading dimension of the array <i>vt2</i> . $ldvt2 \geq n$.
<i>idxc</i>	INTEGER. Array, DIMENSION (<i>n</i>). The permutation used to arrange the columns of <i>u</i> (and rows of <i>vt</i>) into three groups: the first group contains non-zero entries only at and above (or before) <i>n1</i> + 1; the second contains non-zero entries only at and below (or after) <i>n1</i> + 2; and the third is dense. The first column of <i>u</i> and the row of <i>vt</i> are treated separately, however. The rows of the singular vectors found by <i>?lasd4</i> must be likewise permuted before the matrix multiplies can take place.
<i>ctot</i>	INTEGER. Array, DIMENSION (4). A count of the total number of the various types of columns in <i>u</i> (or rows in <i>vt</i>), as described in <i>idxc</i> . The fourth column type is any column which has been deflated.
<i>z</i>	REAL for <i>slasd3</i> DOUBLE PRECISION for <i>dlasd3</i> Array, DIMENSION (<i>k</i>). The first <i>k</i> elements of this array contain the components of the deflation-adjusted updating row vector.

Output Parameters

<i>d</i>	REAL for <i>slasd3</i> DOUBLE PRECISION for <i>dlasd3</i> Array, DIMENSION (<i>k</i>). On exit the square roots of the roots of the secular equation, in ascending order.
<i>info</i>	INTEGER. If <i>info</i> = 0: successful exit. If <i>info</i> = - <i>i</i> < 0, the <i>i</i> -th argument had an illegal value. If <i>info</i> = 1, an singular value did not converge.

Application Notes

This code makes very mild assumptions about floating point arithmetic. It will work on machines with a guard digit in add/subtract, or on those binary machines without guard digits which subtract like the Cray XMP, Cray YMP, Cray C 90, or Cray 2. It could conceivably fail on hexadecimal or decimal machines without guard digits, but we know of none.

?lasd4

Computes the square root of the i -th updated eigenvalue of a positive symmetric rank-one modification to a positive diagonal matrix.

Used by ?bdsdc.

```
call slasd4 ( n, i, d, z, delta, rho, sigma, work, info )
call dlasd4 ( n, i, d, z, delta, rho, sigma, work, info )
```

Discussion

This routine computes the square root of the i -th updated eigenvalue of a positive symmetric rank-one modification to a positive diagonal matrix whose entries are given as the squares of the corresponding entries in the array d , and that $0 \leq d(i) < d(j)$ for $i < j$ and that $rho > 0$. This is arranged by the calling routine, and is no loss in generality. The rank-one modified system is thus

$$\text{diag}(d) * \text{diag}(d) + rho * Z * Z_{\text{transpose}}$$

where we assume the Euclidean norm of Z is 1. The method consists of approximating the rational functions in the secular equation by simpler interpolating rational functions.

Input Parameters

n **INTEGER**. The length of all arrays.

<i>i</i>	INTEGER. The index of the eigenvalue to be computed. $1 \leq i \leq n$.
<i>d</i>	REAL for <code>slasd4</code> DOUBLE PRECISION for <code>dlasd4</code> Array, DIMENSION (<i>n</i>). The original eigenvalues. It is assumed that they are in order, $0 \leq d(i) < d(j)$ for $i < j$.
<i>z</i>	REAL for <code>slasd4</code> DOUBLE PRECISION for <code>dlasd4</code> Array, DIMENSION (<i>n</i>). The components of the updating vector.
<i>rho</i>	REAL for <code>slasd4</code> DOUBLE PRECISION for <code>dlasd4</code> The scalar in the symmetric updating formula.
<i>work</i>	REAL for <code>slasd4</code> DOUBLE PRECISION for <code>dlasd4</code> Workspace array, DIMENSION (<i>n</i>). If $n \neq 1$, <i>work</i> contains $d(j) + \text{sigma}_i$ in its <i>j</i> -th component. If $n = 1$, then <i>work</i> (1) = 1.

Output Parameters

<i>delta</i>	REAL for <code>slasd4</code> DOUBLE PRECISION for <code>dlasd4</code> Array, DIMENSION (<i>n</i>). If $n \neq 1$, <i>delta</i> contains $d(j) - \text{sigma}_i$ in its <i>j</i> -th component. If $n = 1$, then <i>delta</i> (1) = 1. The vector <i>delta</i> contains the information necessary to construct the (singular) eigenvectors.
<i>sigma</i>	REAL for <code>slasd4</code> DOUBLE PRECISION for <code>dlasd4</code> The computed λ_i , the <i>i</i> -th updated eigenvalue.
<i>info</i>	INTEGER. = 0: successful exit > 0: if <i>info</i> = 1, the updating process failed.

?lasd5

Computes the square root of the i -th eigenvalue of a positive symmetric rank-one modification of a 2-by-2 diagonal matrix. Used by ?bdsdc.

```
call slasd5 ( i, d, z, delta, rho, dsigma, work )
call dlasd5 ( i, d, z, delta, rho, dsigma, work )
```

Discussion

This routine computes the square root of the i -th eigenvalue of a positive symmetric rank-one modification of a 2-by-2 diagonal matrix

$$\text{diag}(d) * \text{diag}(d) + rho * Z * Z_transpose$$

The diagonal entries in the array d are assumed to satisfy $0 \leq d(i) < d(j)$ for $i < j$. We also assume $rho > 0$ and that the Euclidean norm of the vector Z is one.

Input Parameters

i **INTEGER**. The index of the eigenvalue to be computed.
 $i = 1$ or $i = 2$.

d **REAL** for **slasd5**
DOUBLE PRECISION for **dlasd5**
Array, **DIMENSION** (2).
The original eigenvalues. We assume $0 \leq d(1) < d(2)$.

z **REAL** for **slasd5**
DOUBLE PRECISION for **dlasd5**
Array, **DIMENSION** (2).
The components of the updating vector.

rho **REAL** for **slasd5**
DOUBLE PRECISION for **dlasd5**
The scalar in the symmetric updating formula.

work REAL for `slasd5`
 DOUBLE PRECISION for `dlasd5`.
 Workspace array, DIMENSION (2).
 Contains ($d(j) + \text{sigma}_i$) in its j -th component.

Output Parameters

delta REAL for `slasd5`
 DOUBLE PRECISION for `dlasd5`.
 Array, DIMENSION (2).
 Contains ($d(j) - \lambda_i$) in its j -th component. The vector *delta* contains the information necessary to construct the eigenvectors.

dsigma REAL for `slasd5`
 DOUBLE PRECISION for `dlasd5`.
 The computed λ_i , the i -th updated eigenvalue.

?lasd6

Computes the SVD of an updated upper bidiagonal matrix obtained by merging two smaller ones by appending a row.
 Used by `?bdsdc`.

```
call slasd6 (  icompq, nl, nr, sqre, d, vf, vl, alpha,
              beta, idxq, perm, givptr, givcol, ldgcol,
              givnum, ldgnum, poles, difl, difr, z, k, c,
              s, work, iwork, info)
call dlasd6 (  icompq, nl, nr, sqre, d, vf, vl, alpha,
              beta, idxq, perm, givptr, givcol, ldgcol,
              givnum, ldgnum, poles, difl, difr, z, k, c,
              s, work, iwork, info)
```

Discussion

The routine `?lasd6` computes the *SVD* of an updated upper bidiagonal matrix B obtained by merging two smaller ones by appending a row. This routine is used only for the problem which requires all singular values and optionally singular vector matrices in factored form. B is an n -by- m matrix with

$n = nl + nr + 1$ and $m = n + sqre$. A related subroutine, `?lasd1`, handles the case in which all singular values and singular vectors of the bidiagonal matrix are desired. `?lasd6` computes the *SVD* as follows:

$$B = U(in) * \begin{bmatrix} D1(in) & 0 & 0 & 0 \\ Z1' & a & Z2' & b \\ 0 & 0 & D2(in) & 0 \end{bmatrix} * VT(in)$$

$$= U(out) * (D(out) \ 0) * VT(out)$$

where $Z = (Z1' \ a \ Z2' \ b) = u' \ VT'$, and u is a vector of dimension m with `alpha` and `beta` in the $nl+1$ and $nl+2$ -th entries and zeros elsewhere; and the entry b is empty if `sqre` = 0.

The singular values of B can be computed using $D1$, $D2$, the first components of all the right singular vectors of the lower block, and the last components of all the right singular vectors of the upper block. These components are stored and updated in `vf` and `v1`, respectively, in `?lasd6`. Hence U and VT are not explicitly referenced.

The singular values are stored in D . The algorithm consists of two stages: the first stage consists of deflating the size of the problem when there are multiple singular values or if there is a zero in the Z vector. For each such occurrence the dimension of the secular equation problem is reduced by one. This stage is performed by the routine `?lasd7`.

The second stage consists of calculating the updated singular values. This is done by finding the roots of the secular equation via the routine `?lasd4` (as called by `?lasd8`). This routine also updates `vf` and `v1` and computes the distances between the updated singular values and the old singular values. `?lasd6` is called from `?lasda`.

Input Parameters

<i>icompq</i>	INTEGER . Specifies whether singular vectors are to be computed in factored form: = 0: Compute singular values only = 1: Compute singular vectors in factored form as well.
<i>nl</i>	INTEGER . The row dimension of the upper block. $nl \geq 1$.
<i>nr</i>	INTEGER . The row dimension of the lower block. $nr \geq 1$.
<i>sqre</i>	INTEGER . = 0: the lower block is an <i>nr</i> -by- <i>nr</i> square matrix. = 1: the lower block is an <i>nr</i> -by-(<i>nr</i> +1) rectangular matrix. The bidiagonal matrix has row dimension $n=nl+nr+1$, and column dimension $m = n + sqre$.
<i>d</i>	REAL for slasd6 DOUBLE PRECISION for dlsasd6 Array, DIMENSION ($nl+nr+1$). On entry $d(1:nl,1:nl)$ contains the singular values of the upper block, and $d(nl+2:n)$ contains the singular values of the lower block.
<i>vf</i>	REAL for slasd6 DOUBLE PRECISION for dlsasd6 Array, DIMENSION (m). On entry, $vf(1:nl+1)$ contains the first components of all right singular vectors of the upper block; and $vf(nl+2:m)$ contains the first components of all right singular vectors of the lower block.
<i>vl</i>	REAL for slasd6 DOUBLE PRECISION for dlsasd6 Array, DIMENSION (m). On entry, $vl(1:nl+1)$ contains the last components of all right singular vectors of the upper block; and $vl(nl+2:m)$ contains the last components of all right singular vectors of the lower block.

<i>alpha</i>	REAL for <code>slasd6</code> DOUBLE PRECISION for <code>dlasd6</code> Contains the diagonal element associated with the added row.
<i>beta</i>	REAL for <code>slasd6</code> DOUBLE PRECISION for <code>dlasd6</code> Contains the off-diagonal element associated with the added row.
<i>ldgcol</i>	INTEGER. The leading dimension of the output array <i>givcol</i> , must be at least <i>n</i> .
<i>ldgnum</i>	INTEGER. The leading dimension of the output arrays <i>givnum</i> and <i>poles</i> , must be at least <i>n</i> .
<i>work</i>	REAL for <code>slasd6</code> DOUBLE PRECISION for <code>dlasd6</code> Workspace array, DIMENSION ($4m$).
<i>iwork</i>	INTEGER Workspace array, DIMENSION ($3n$).

Output Parameters

<i>d</i>	On exit <i>d</i> (1: <i>n</i>) contains the singular values of the modified matrix.
<i>vf</i>	On exit, <i>vf</i> contains the first components of all right singular vectors of the bidiagonal matrix.
<i>vl</i>	On exit, <i>vl</i> contains the last components of all right singular vectors of the bidiagonal matrix.
<i>idxq</i>	INTEGER. Array, DIMENSION (<i>n</i>). This contains the permutation which will reintegrate the subproblem just solved back into sorted order, that is, <i>d</i> (<i>idxq</i> (<i>i</i> = 1, <i>n</i>)) will be in ascending order.
<i>perm</i>	INTEGER. Array, DIMENSION (<i>n</i>). The permutations (from deflation and sorting) to be applied to each block. Not referenced if <i>icompq</i> = 0.

<i>givptr</i>	INTEGER. The number of Givens rotations which took place in this subproblem. Not referenced if <i>icompg</i> = 0.
<i>givcol</i>	INTEGER. Array, DIMENSION (<i>ldgcol</i> , 2). Each pair of numbers indicates a pair of columns to take place in a Givens rotation. Not referenced if <i>icompg</i> = 0.
<i>givnum</i>	REAL for slasd6 DOUBLE PRECISION for dlasd6 Array, DIMENSION (<i>ldgnum</i> , 2). Each number indicates the <i>C</i> or <i>S</i> value to be used in the corresponding Givens rotation. Not referenced if <i>icompg</i> = 0.
<i>poles</i>	REAL for slasd6 DOUBLE PRECISION for dlasd6 Array, DIMENSION (<i>ldgnum</i> , 2). On exit, <i>poles</i> (1,*) is an array containing the new singular values obtained from solving the secular equation, and <i>poles</i> (2,*) is an array containing the poles in the secular equation. Not referenced if <i>icompg</i> = 0.
<i>difl</i>	REAL for slasd6 DOUBLE PRECISION for dlasd6 Array, DIMENSION (<i>n</i>). On exit, <i>difl</i> (<i>i</i>) is the distance between <i>i</i> -th updated (undeflated) singular value and the <i>i</i> -th (undeflated) old singular value.
<i>difr</i>	REAL for slasd6 DOUBLE PRECISION for dlasd6 Array, DIMENSION (<i>ldgnum</i> , 2) if <i>icompg</i> = 1 and DIMENSION (<i>n</i>) if <i>icompg</i> = 0. On exit, <i>difr</i> (<i>i</i> , 1) is the distance between <i>i</i> -th updated (undeflated) singular value and the <i>i</i> +1-th (undeflated) old singular value. If <i>icompg</i> = 1, <i>difr</i> (1: <i>k</i> , 2) is an array containing the normalizing factors for the right singular vector matrix. See ? lasd8 for details on <i>difl</i> and <i>difr</i> .

<i>z</i>	REAL for <code>slasd6</code> DOUBLE PRECISION for <code>dlasd6</code> Array, DIMENSION (<i>m</i>). The first elements of this array contain the components of the deflation-adjusted updating row vector.
<i>k</i>	INTEGER. Contains the dimension of the non-deflated matrix. This is the order of the related secular equation. $1 \leq k \leq n$.
<i>c</i>	REAL for <code>slasd6</code> DOUBLE PRECISION for <code>dlasd6</code> <i>c</i> contains garbage if <code>sqre</code> = 0 and the C-value of a Givens rotation related to the right null space if <code>sqre</code> = 1.
<i>s</i>	REAL for <code>slasd6</code> DOUBLE PRECISION for <code>dlasd6</code> <i>s</i> contains garbage if <code>sqre</code> = 0 and the S-value of a Givens rotation related to the right null space if <code>sqre</code> = 1.
<i>info</i>	INTEGER. = 0: successful exit. < 0: if <i>info</i> = - <i>i</i> , the <i>i</i> -th argument had an illegal value. > 0: if <i>info</i> = 1, an singular value did not converge

?lasd7

Merges the two sets of singular values together into a single sorted set. Then it tries to deflate the size of the problem.

Used by ?bdsdc.

```
call slasd7 ( icompg, nl, nr, sqre, k, d, z, zw, vf, vfw,
             vl, vlw, alpha, beta, dsigma, idx, idxp,
             idxq, perm, givptr, givcol, ldgcol, givnum,
             ldgnum, c, s, info )
```

```
call dlasd7 (  icompq, nl, nr, sqre, k, d, z, zw, vf, vfw,
              vl, vlw, alpha, beta, dsigma, idx, idxp,
              idxq, perm, givptr, givcol, ldgcol, givnum,
              ldgnum, c, s, info )
```

Discussion

The routine `?lasd7` merges the two sets of singular values together into a single sorted set. Then it tries to deflate the size of the problem. There are two ways in which deflation can occur: when two or more singular values are close together or if there is a tiny entry in the Z vector. For each such occurrence the order of the related secular equation problem is reduced by one. `?lasd7` is called from `?lasd6`.

Input Parameters

<i>icompq</i>	INTEGER. Specifies whether singular vectors are to be computed in compact form, as follows: = 0: Compute singular values only. = 1: Compute singular vectors of upper bidiagonal matrix in compact form.
<i>nl</i>	INTEGER. The row dimension of the upper block. $nl \geq 1$.
<i>nr</i>	INTEGER. The row dimension of the lower block. $nr \geq 1$.
<i>sqre</i>	INTEGER. = 0: the lower block is an nr -by- nr square matrix. = 1: the lower block is an nr -by- $(nr+1)$ rectangular matrix. The bidiagonal matrix has $n = nl + nr + 1$ rows and $m = n + sqre \geq n$ columns.
<i>d</i>	REAL for <code>slasd7</code> DOUBLE PRECISION for <code>dlasd7</code> Array, DIMENSION (n). On entry <i>d</i> contains the singular values of the two submatrices to be combined.
<i>zw</i>	REAL for <code>slasd7</code> DOUBLE PRECISION for <code>dlasd7</code> Array, DIMENSION (m). Workspace for <i>z</i> .

vf REAL for *slasd7*
DOUBLE PRECISION for *dlasd7*
Array, DIMENSION (*m*). On entry, *vf*(1:*nl*+1) contains the first components of all right singular vectors of the upper block; and *vf*(*nl*+2:*m*) contains the first components of all right singular vectors of the lower block.

vfw REAL for *slasd7*
DOUBLE PRECISION for *dlasd7*
Array, DIMENSION (*m*). Workspace for *vf*.

vl REAL for *slasd7*
DOUBLE PRECISION for *dlasd7*
Array, DIMENSION (*m*). On entry, *vl*(1:*nl*+1) contains the last components of all right singular vectors of the upper block; and *vl*(*nl*+2:*m*) contains the last components of all right singular vectors of the lower block.

vlw REAL for *slasd7*
DOUBLE PRECISION for *dlasd7*
Array, DIMENSION (*m*). Workspace for *vl*.

alpha REAL for *slasd7*
DOUBLE PRECISION for *dlasd7*.
Contains the diagonal element associated with the added row.

beta REAL for *slasd7*
DOUBLE PRECISION for *dlasd7*
Contains the off-diagonal element associated with the added row.

idx INTEGER.
Workspace array, DIMENSION (*n*). This will contain the permutation used to sort the contents of *d* into ascending order.

<i>idxp</i>	INTEGER. Workspace array, DIMENSION (<i>n</i>). This will contain the permutation used to place deflated values of <i>d</i> at the end of the array.
<i>idxq</i>	INTEGER. Array, DIMENSION (<i>n</i>). This contains the permutation which separately sorts the two sub-problems in <i>d</i> into ascending order. Note that entries in the first half of this permutation must first be moved one position backward; and entries in the second half must first have <i>n</i> +1 added to their values.
<i>ldgcol</i>	INTEGER. The leading dimension of the output array <i>givcol</i> , must be at least <i>n</i> .
<i>ldgnum</i>	INTEGER. The leading dimension of the output array <i>givnum</i> , must be at least <i>n</i> .

Output Parameters

<i>k</i>	INTEGER. Contains the dimension of the non-deflated matrix, this is the order of the related secular equation. $1 \leq k \leq n$.
<i>d</i>	On exit, <i>d</i> contains the trailing (<i>n-k</i>) updated singular values (those which were deflated) sorted into increasing order.
<i>z</i>	REAL for <i>slasd7</i> DOUBLE PRECISION for <i>dlsasd7</i> . Array, DIMENSION (<i>m</i>). On exit, <i>z</i> contains the updating row vector in the secular equation.
<i>vf</i>	On exit, <i>vf</i> contains the first components of all right singular vectors of the bidiagonal matrix.
<i>vl</i>	On exit, <i>vl</i> contains the last components of all right singular vectors of the bidiagonal matrix.

dsigma REAL for *slasd7*
DOUBLE PRECISION for *dlsasd7*.
Array, DIMENSION (*n*). Contains a copy of the diagonal elements (*k*-1 singular values and one zero) in the secular equation.

idxp On output, *idxp*(2:*k*) points to the nondeflated *d*-values and *idxp*(*k*+1:*n*) points to the deflated singular values.

perm INTEGER.
Array, DIMENSION (*n*). The permutations (from deflation and sorting) to be applied to each singular block. Not referenced if *icomprq* = 0.

givptr INTEGER. The number of Givens rotations which took place in this subproblem. Not referenced if *icomprq* = 0.

givcol INTEGER.
Array, DIMENSION (*ldgcol*, 2). Each pair of numbers indicates a pair of columns to take place in a Givens rotation. Not referenced if *icomprq* = 0.

givnum REAL for *slasd7*
DOUBLE PRECISION for *dlsasd7*.
Array, DIMENSION (*ldgnum*, 2). Each number indicates the *C* or *S* value to be used in the corresponding Givens rotation. Not referenced if *icomprq* = 0.

c REAL for *slasd7*.
DOUBLE PRECISION for *dlsasd7*.
c contains garbage if *sqre* = 0 and the *C*-value of a Givens rotation related to the right null space if *sqre* = 1.

s REAL for *slasd7*.
DOUBLE PRECISION for *dlsasd7*.
s contains garbage if *sqre* = 0 and the *S*-value of a Givens rotation related to the right null space if *sqre* = 1.

info **INTEGER**.
 = 0: successful exit.
 < 0: if *info* = -*i*, the *i*-th argument had an illegal value.

?lasd8

Finds the square roots of the roots of the secular equation, and stores, for each element in D, the distance to its two nearest poles. Used by ?bdsdc.

```
call slasd8 ( icompg, k, d, z, vf, vl, difl, difr,
             lddifr, dsigma, work, info )
call dlasd8 ( icompg, k, d, z, vf, vl, difl, difr,
             lddifr, dsigma, work, info )
```

Discussion

The routine `?lasd8` finds the square roots of the roots of the secular equation, as defined by the values in *dsigma* and *z*. It makes the appropriate calls to `?lasd4`, and stores, for each element in *d*, the distance to its two nearest poles (elements in *dsigma*). It also updates the arrays *vf* and *vl*, the first and last components of all the right singular vectors of the original bidiagonal matrix. `?lasd8` is called from `?lasd6`.

Input Parameters

icompg **INTEGER**. Specifies whether singular vectors are to be computed in factored form in the calling routine:
 = 0: Compute singular values only.
 = 1: Compute singular vectors in factored form as well.

k **INTEGER**. The number of terms in the rational function to be solved by `?lasd4`. $k \geq 1$.

<i>z</i>	REAL for <i>slasd8</i> DOUBLE PRECISION for <i>dlasd8</i> . Array, DIMENSION (<i>k</i>). The first <i>k</i> elements of this array contain the components of the deflation-adjusted updating row vector.
<i>vf</i>	REAL for <i>slasd8</i> DOUBLE PRECISION for <i>dlasd8</i> . Array, DIMENSION (<i>k</i>). On entry, <i>vf</i> contains information passed through <i>dbede8</i> .
<i>v1</i>	REAL for <i>slasd8</i> DOUBLE PRECISION for <i>dlasd8</i> . Array, DIMENSION (<i>k</i>). On entry, <i>v1</i> contains information passed through <i>dbede8</i> .
<i>lddifr</i>	INTEGER. The leading dimension of the output array <i>difr</i> , must be at least <i>k</i> .
<i>dsigma</i>	REAL for <i>slasd8</i> DOUBLE PRECISION for <i>dlasd8</i> . Array, DIMENSION (<i>k</i>). The first <i>k</i> elements of this array contain the old roots of the deflated updating problem. These are the poles of the secular equation.
<i>work</i>	REAL for <i>slasd8</i> DOUBLE PRECISION for <i>dlasd8</i> . Workspace array, DIMENSION at least $(3k)$.

Output Parameters

<i>d</i>	REAL for <i>slasd8</i> DOUBLE PRECISION for <i>dlasd8</i> . Array, DIMENSION (<i>k</i>). On output, <i>d</i> contains the updated singular values.
<i>vf</i>	On exit, <i>vf</i> contains the first <i>k</i> components of the first components of all right singular vectors of the bidiagonal matrix.
<i>v1</i>	On exit, <i>v1</i> contains the first <i>k</i> components of the last components of all right singular vectors of the bidiagonal matrix.

<i>difl</i>	<p>REAL for <code>slasd8</code> DOUBLE PRECISION for <code>dlasd8</code>. Array, <code>DIMENSION (k)</code>. On exit, $difl(i) = d(i) - dsigma(i)$.</p>
<i>difr</i>	<p>REAL for <code>slasd8</code> DOUBLE PRECISION for <code>dlasd8</code>. Array, <code>DIMENSION (lddifr, 2)</code> if <code>icompq = 1</code> and <code>DIMENSION (k)</code> if <code>icompq = 0</code>. On exit, $difr(i,1) = d(i) - dsigma(i+1)$, $difr(k,1)$ is not defined and will not be referenced. If <code>icompq = 1</code>, $difr(1:k,2)$ is an array containing the normalizing factors for the right singular vector matrix.</p>
<i>info</i>	<p>INTEGER. = 0: successful exit. < 0: if $info = -i$, the i-th argument had an illegal value. > 0: if $info = 1$, a singular value did not converge.</p>

?lasd9

Finds the square roots of the roots of the secular equation, and stores, for each element in D , the distance to its two nearest poles. Used by ?bdsdc.

```
call slasd9 ( icompq, ldu, k, d, z, vf, vl, difl, difr,
             dsigma, work, info )
call dlasd9 ( icompq, ldu, k, d, z, vf, vl, difl, difr,
             dsigma, work, info )
```

Discussion

The routine ?lasd9 finds the square roots of the roots of the secular equation, as defined by the values in *dsigma* and *z*. It makes the appropriate calls to ?lasd4, and stores, for each element in *d*, the distance

to its two nearest poles (elements in *dsigma*). It also updates the arrays *vf* and *vl*, the first and last components of all the right singular vectors of the original bidiagonal matrix. `?lasd9` is called from `?lasd7`.

Input Parameters

<i>icompg</i>	INTEGER . Specifies whether singular vectors are to be computed in factored form in the calling routine: If <i>icompg</i> = 0, compute singular values only; If <i>icompg</i> = 1, compute singular vector matrices in factored form also.
<i>k</i>	INTEGER . The number of terms in the rational function to be solved by <code>slasd4</code> . $k \geq 1$.
<i>dsigma</i>	REAL for <code>slasd9</code> DOUBLE PRECISION for <code>dlasd9</code> . Array, DIMENSION (<i>k</i>). The first <i>k</i> elements of this array contain the old roots of the deflated updating problem. These are the poles of the secular equation.
<i>z</i>	REAL for <code>slasd9</code> DOUBLE PRECISION for <code>dlasd9</code> . Array, DIMENSION (<i>k</i>). The first <i>k</i> elements of this array contain the components of the deflation-adjusted updating row vector.
<i>vf</i>	REAL for <code>slasd9</code> DOUBLE PRECISION for <code>dlasd9</code> . Array, DIMENSION (<i>k</i>). On entry, <i>vf</i> contains information passed through <code>sbede8</code> .
<i>vl</i>	REAL for <code>slasd9</code> DOUBLE PRECISION for <code>dlasd9</code> . Array, DIMENSION (<i>k</i>). On entry, <i>vl</i> contains information passed through <code>sbede8</code> .
<i>work</i>	REAL for <code>slasd9</code> DOUBLE PRECISION for <code>dlasd9</code> . Workspace array, DIMENSION at least (3 <i>k</i>).

Output Parameters

<i>d</i>	<p>REAL for <code>slasd9</code> DOUBLE PRECISION for <code>dlasd9</code>. Array, DIMENSION(<i>k</i>). <i>d</i>(<i>i</i>) contains the updated singular values.</p>
<i>vf</i>	<p>On exit, <i>vf</i> contains the first <i>k</i> components of the first components of all right singular vectors of the bidiagonal matrix.</p>
<i>vl</i>	<p>On exit, <i>vl</i> contains the first <i>k</i> components of the last components of all right singular vectors of the bidiagonal matrix.</p>
<i>difl</i>	<p>REAL for <code>slasd9</code> DOUBLE PRECISION for <code>dlasd9</code>. Array, DIMENSION (<i>k</i>). On exit, $difl(i) = d(i) - dsigma(i)$.</p>
<i>difr</i>	<p>REAL for <code>slasd9</code> DOUBLE PRECISION for <code>dlasd9</code>. Array, DIMENSION (<i>ldu</i>, 2) if <i>icompq</i> = 1 and DIMENSION (<i>k</i>) if <i>icompq</i> = 0. On exit, $difr(i, 1) = d(i) - dsigma(i+1)$, $difr(k, 1)$ is not defined and will not be referenced. If <i>icompq</i> = 1, $difr(1:k, 2)$ is an array containing the normalizing factors for the right singular vector matrix.</p>
<i>info</i>	<p>INTEGER. = 0: successful exit. < 0: if <i>info</i> = -<i>i</i>, the <i>i</i>-th argument had an illegal value. > 0: if <i>info</i> = 1, an singular value did not converge</p>

?lasda

Computes the singular value decomposition (SVD) of a real upper bidiagonal matrix with diagonal d and off-diagonal e . Used by ?bdsdc.

```
call slasda ( icompg, smlsiz, n, sqre, d, e, u, ldu, vt,  
            k, difl, difr, z, poles, givptr, givcol,  
            ldgcol, perm, givnum, c, s, work, iwork,  
            info )  
call dlasda ( icompg, smlsiz, n, sqre, d, e, u, ldu, vt,  
            k, difl, difr, z, poles, givptr, givcol,  
            ldgcol, perm, givnum, c, s, work, iwork,  
            info )
```

Discussion

Using a divide and conquer approach, ?lasda computes the singular value decomposition (SVD) of a real upper bidiagonal n -by- m matrix B with diagonal d and off-diagonal e , where $m = n + sqre$. The algorithm computes the singular values in the SVD $B = U*S*VT$. The orthogonal matrices U and VT are optionally computed in compact form. A related subroutine, ?lasd0, computes the singular values and the singular vectors in explicit form.

Input Parameters

icompg **INTEGER**. Specifies whether singular vectors are to be computed in compact form, as follows:
= 0: Compute singular values only.
= 1: Compute singular vectors of upper bidiagonal matrix in compact form.

smlsiz **INTEGER**. The maximum size of the subproblems at the bottom of the computation tree.

<i>n</i>	INTEGER. The row dimension of the upper bidiagonal matrix. This is also the dimension of the main diagonal array <i>d</i> .
<i>sqre</i>	INTEGER. Specifies the column dimension of the bidiagonal matrix. If <i>sqre</i> = 0: The bidiagonal matrix has column dimension $m = n$; If <i>sqre</i> = 1: The bidiagonal matrix has column dimension $m = n + 1$.
<i>d</i>	REAL for <i>slasda</i> DOUBLE PRECISION for <i>dlasda</i> . Array, DIMENSION (<i>n</i>). On entry <i>d</i> contains the main diagonal of the bidiagonal matrix.
<i>e</i>	REAL for <i>slasda</i> DOUBLE PRECISION for <i>dlasda</i> . Array, DIMENSION (<i>m</i> - 1). Contains the subdiagonal entries of the bidiagonal matrix. On exit, <i>e</i> has been destroyed.
<i>ldu</i>	INTEGER. The leading dimension of arrays <i>u</i> , <i>vt</i> , <i>difl</i> , <i>difr</i> , <i>poles</i> , <i>givnum</i> , and <i>z</i> . $ldu \geq n$.
<i>ldgcol</i>	INTEGER. The leading dimension of arrays <i>givcol</i> and <i>perm</i> . $ldgcol \geq n$.
<i>work</i>	REAL for <i>slasda</i> DOUBLE PRECISION for <i>dlasda</i> . Workspace array, DIMENSION ($6n + (smlsiz + 1)^2$).
<i>iwork</i>	INTEGER. Workspace array, DIMENSION must be at least $(7n)$.

Output Parameters

<i>d</i>	On exit <i>d</i> , if <i>info</i> = 0, contains the singular values of the bidiagonal matrix.
<i>u</i>	REAL for <i>slasda</i> DOUBLE PRECISION for <i>dlasda</i> . Array, DIMENSION (<i>ldu</i> , <i>smlsiz</i>) if <i>icompg</i> = 1.

Not referenced if *icompq* = 0.

If *icompq* = 1, on exit, *u* contains the left singular vector matrices of all subproblems at the bottom level.

<i>vt</i>	<p>REAL for <i>slasda</i> DOUBLE PRECISION for <i>dlasda</i>. Array, DIMENSION (<i>ldu</i>, <i>smlsiz</i>+1) if <i>icompq</i> = 1, and not referenced if <i>icompq</i> = 0. If <i>icompq</i> = 1, on exit, <i>vt</i> contains the right singular vector matrices of all subproblems at the bottom level.</p>
<i>k</i>	<p>INTEGER. Array, DIMENSION (<i>n</i>) if <i>icompq</i> = 1 and DIMENSION (1) if <i>icompq</i> = 0. If <i>icompq</i> = 1, on exit, <i>k</i>(<i>i</i>) is the dimension of the <i>i</i>-th secular equation on the computation tree.</p>
<i>difl</i>	<p>REAL for <i>slasda</i> DOUBLE PRECISION for <i>dlasda</i>. Array, DIMENSION (<i>ldu</i>, <i>nlvl</i>), where <i>nlvl</i> = floor (log₂ (<i>n</i>/<i>smlsiz</i>)).</p>
<i>difr</i>	<p>REAL for <i>slasda</i> DOUBLE PRECISION for <i>dlasda</i>. Array, DIMENSION (<i>ldu</i>, 2 <i>nlvl</i>) if <i>icompq</i> = 1 and DIMENSION (<i>n</i>) if <i>icompq</i> = 0. If <i>icompq</i> = 1, on exit, <i>difl</i>(1:<i>n</i>, <i>i</i>) and <i>difr</i>(1:<i>n</i>,2<i>i</i>-1) record distances between singular values on the <i>i</i>-th level and singular values on the (<i>i</i> -1)-th level, and <i>difr</i>(1:<i>n</i>, 2<i>i</i>) contains the normalizing factors for the right singular vector matrix. See ?<i>lasd8</i> for details.</p>
<i>z</i>	<p>REAL for <i>slasda</i> DOUBLE PRECISION for <i>dlasda</i>. Array, DIMENSION (<i>ldu</i>, <i>nlvl</i>) if <i>icompq</i> = 1 and DIMENSION (<i>n</i>) if <i>icompq</i> = 0.</p>

The first k elements of $z(1, i)$ contain the components of the deflation-adjusted updating row vector for subproblems on the i -th level.

- poles* REAL for *slasda*
 DOUBLE PRECISION for *dlasda*
 Array, DIMENSION (*ldu*, $2 * n_{lvl}$) if *icompq* = 1, and not referenced if *icompq* = 0. If *icompq* = 1, on exit, *poles*(1, $2i - 1$) and *poles*(1, $2i$) contain the new and old singular values involved in the secular equations on the i -th level.
- givptr* INTEGER.
 Array, DIMENSION (*n*) if *icompq* = 1, and not referenced if *icompq* = 0. If *icompq* = 1, on exit, *givptr*(i) records the number of Givens rotations performed on the i -th problem on the computation tree.
- givcol* INTEGER .
 Array, DIMENSION (*ldgcol*, $2 * n_{lvl}$) if *icompq* = 1, and not referenced if *icompq* = 0. If *icompq* = 1, on exit, for each i , *givcol*(1, $2i - 1$) and *givcol*(1, $2i$) record the locations of Givens rotations performed on the i -th level on the computation tree.
- perm* INTEGER .
 Array, DIMENSION (*ldgcol*, *nlvl*) if *icompq* = 1, and not referenced if *icompq* = 0. If *icompq* = 1, on exit, *perm*($1, i$) records permutations done on the i -th level of the computation tree.
- givnum* REAL for *slasda*
 DOUBLE PRECISION for *dlasda*.
 Array DIMENSION (*ldu*, $2 * n_{lvl}$) if *icompq* = 1, and not referenced if *icompq* = 0. If *icompq* = 1, on exit, for each i , *givnum*(1, $2i - 1$) and *givnum*(1, $2i$) record the C- and S-values of Givens rotations performed on the i -th level on the computation tree.

c REAL for `slasda`
DOUBLE PRECISION for `dlasda`.
Array,
DIMENSION (*n*) if `icompg` = 1, and
DIMENSION (1) if `icompg` = 0.
If `icompg` = 1 and the *i*-th subproblem is not square, on
exit, *c*(*i*) contains the *C*-value of a Givens rotation
related to the right null space of the *i*-th subproblem.

s REAL for `slasda`
DOUBLE PRECISION for `dlasda`.
Array,
DIMENSION (*n*) if `icompg` = 1, and
DIMENSION (1) if `icompg` = 0.
If `icompg` = 1 and the *i*-th subproblem is not square, on
exit, *s*(*i*) contains the *S*-value of a Givens rotation
related to the right null space of the *i*-th subproblem.

info INTEGER.
= 0: successful exit.
< 0: if *info* = -*i*, the *i*-th argument had an illegal value
> 0: if *info* = 1, an singular value did not converge

?lasdq

Computes the SVD of a real bidiagonal
matrix with diagonal *d* and off-diagonal *e*.

Used by `?bdsdc`.

```
call slasdq ( uplo, sqre, n, ncv, nru, ncc, d, e, vt,
             ldvt, u, ldu, c, ldc, work, info )
call dlasdq ( uplo, sqre, n, ncv, nru, ncc, d, e, vt,
             ldvt, u, ldu, c, ldc, work, info )
```

Discussion

The routine `?lasdq` computes the singular value decomposition (SVD) of a real (upper or lower) bidiagonal matrix with diagonal d and off-diagonal e , accumulating the transformations if desired. Letting B denote the input bidiagonal matrix, the algorithm computes orthogonal matrices Q and P such that $B = Q S P'$ (P' denotes the transpose of P). The singular values S are overwritten on d .

The input matrix U is changed to UQ if desired.

The input matrix VT is changed to $P' VT$ if desired.

The input matrix C is changed to $Q' C$ if desired.

Input Parameters

- uplo* CHARACTER*1. On entry, *uplo* specifies whether the input bidiagonal matrix is upper or lower bidiagonal. If *uplo* = 'U' or 'u', B is upper bidiagonal; If *uplo* = 'L' or 'l', B is lower bidiagonal.
- sqre* INTEGER.
 = 0: then the input matrix is n -by- n .
 = 1: then the input matrix is n -by- $(n+1)$ if *uplu* = 'U' and $(n+1)$ -by- n if *uplu* = 'L'. The bidiagonal matrix has $n = nl + nr + 1$ rows and $m = n + sqre \geq n$ columns.
- n* INTEGER. On entry, *n* specifies the number of rows and columns in the matrix. *n* must be at least 0.
- ncvt* INTEGER. On entry, *ncvt* specifies the number of columns of the matrix VT . *ncvt* must be at least 0.
- nru* INTEGER. On entry, *nru* specifies the number of rows of the matrix U . *nru* must be at least 0.
- ncc* INTEGER. On entry, *ncc* specifies the number of columns of the matrix C . *ncc* must be at least 0.
- d* REAL for `sldsq`
 DOUBLE PRECISION for `dlasdq`.
 Array, DIMENSION (n). On entry, *d* contains the diagonal entries of the bidiagonal matrix whose SVD is desired.

- e* REAL for `slasdq`
DOUBLE PRECISION for `dlasdq`.
Array, DIMENSION is $(n-1)$ if `sqre` = 0 and n if `sqre` = 1. On entry, the entries of *e* contain the off-diagonal entries of the bidiagonal matrix whose SVD is desired.
- vt* REAL for `slasdq`
DOUBLE PRECISION for `dlasdq`.
Array, DIMENSION (*ldvt*, *ncvt*). On entry, contains a matrix which on exit has been premultiplied by P' , dimension n -by-*ncvt* if `sqre` = 0 and $(n+1)$ -by-*ncvt* if `sqre` = 1 (not referenced if *ncvt*=0).
- ldvt* INTEGER. On entry, *ldvt* specifies the leading dimension of *vt* as declared in the calling (sub) program. *ldvt* must be at least 1. If *ncvt* is nonzero, *ldvt* must also be at least n .
- u* REAL for `slasdq`
DOUBLE PRECISION for `dlasdq`.
Array, DIMENSION (*ldu*, n). On entry, contains a matrix which on exit has been postmultiplied by Q , dimension *nru*-by- n if `sqre` = 0 and *nru*-by- $(n+1)$ if `sqre` = 1 (not referenced if *nru*=0).
- ldu* INTEGER. On entry, *ldu* specifies the leading dimension of *u* as declared in the calling (sub) program. *ldu* must be at least $\max(1, nru)$.
- c* REAL for `slasdq`
DOUBLE PRECISION for `dlasdq`.
Array, DIMENSION (*ldc*, *ncc*). On entry, contains an n -by-*ncc* matrix which on exit has been premultiplied by Q' , dimension n -by-*ncc* if `sqre` = 0 and $(n+1)$ -by-*ncc* if `sqre` = 1 (not referenced if *ncc*=0).
- ldc* INTEGER. On entry, *ldc* specifies the leading dimension of *c* as declared in the calling (sub) program. *ldc* must be at least 1. If *ncc* is non-zero, *ldc* must also be at least n .

work REAL for `slasdq`
 DOUBLE PRECISION for `dlasdq`.
 Array, DIMENSION (4*n*). This is a workspace array. Only referenced if one of *ncvt*, *nru*, or *ncc* is nonzero, and if *n* is at least 2.

Output Parameters

d On normal exit, *d* contains the singular values in ascending order.

e On normal exit, *e* will contain 0. If the algorithm does not converge, *d* and *e* will contain the diagonal and superdiagonal entries of a bidiagonal matrix orthogonally equivalent to the one given as input.

vt On exit, the matrix has been premultiplied by *P*'.

u On exit, the matrix has been postmultiplied by *Q*.

c On exit, the matrix has been premultiplied by *Q*'.

info INTEGER. On exit, a value of 0 indicates a successful exit. If *info* < 0, argument number *-info* is illegal. If *info* > 0, the algorithm did not converge, and *info* specifies how many superdiagonals did not converge.

?lasdt

Creates a tree of subproblems for bidiagonal divide and conquer.

Used by ?bdsdc.

```
call slasdt ( n, lvl, nd, inode, ndiml, ndimr, msub )
call dlasdt ( n, lvl, nd, inode, ndiml, ndimr, msub )
```

Discussion

The routine creates a tree of subproblems for bidiagonal divide and conquer.

Input Parameters

- n* **INTEGER**. On entry, the number of diagonal elements of the bidiagonal matrix.
- msub* **INTEGER**. On entry, the maximum row dimension each subproblem at the bottom of the tree can be of.

Output Parameters

- lvl* **INTEGER**. On exit, the number of levels on the computation tree.
- nd* **INTEGER**. On exit, the number of nodes on the tree.
- inode* **INTEGER**.
Array, **DIMENSION** (*n*). On exit, centers of subproblems.
- ndiml* **INTEGER**.
Array, **DIMENSION** (*n*). On exit, row dimensions of left children.
- ndimr* **INTEGER**.
Array, **DIMENSION** (*n*). On exit, row dimensions of right children.

?laset

*Initializes the off-diagonal elements
and the diagonal elements of a matrix to
given values.*

```
call slaset ( uplo, m, n, alpha, beta, a, lda )
call dlaset ( uplo, m, n, alpha, beta, a, lda )
call claset ( uplo, m, n, alpha, beta, a, lda )
call zlaset ( uplo, m, n, alpha, beta, a, lda )
```

Discussion

The routine initializes an m -by- n matrix A to $beta$ on the diagonal and $alpha$ on the off-diagonals .

Input parameters

<i>uplo</i>	<p>CHARACTER*1. Specifies the part of the matrix A to be set.</p> <p>If <i>uplo</i> = 'U', upper triangular part is set; the strictly lower triangular part of A is not changed.</p> <p>If <i>uplo</i> = 'L': lower triangular part is set; the strictly upper triangular part of A is not changed.</p> <p>Otherwise: all of the matrix A is set.</p>
<i>m</i>	INTEGER. The number of rows of the matrix A . $m \geq 0$.
<i>n</i>	INTEGER. The number of columns of the matrix A . $n \geq 0$.
<i>alpha, beta</i>	<p>REAL for <code>saset</code></p> <p>DOUBLE PRECISION for <code>dlaset</code></p> <p>COMPLEX for <code>claset</code></p> <p>COMPLEX*16 for <code>zlaset</code>.</p> <p>The constants to which the off-diagonal and diagonal elements are to be set, respectively.</p>
<i>a</i>	<p>REAL for <code>saset</code></p> <p>DOUBLE PRECISION for <code>dlaset</code></p> <p>COMPLEX for <code>claset</code></p> <p>COMPLEX*16 for <code>zlaset</code>.</p> <p>Array, DIMENSION (<i>lda</i>, <i>n</i>).</p> <p>On entry, the m-by-n matrix A.</p>
<i>lda</i>	<p>INTEGER. The leading dimension of the array A.</p> <p>$lda \geq \max(1,m)$.</p>

Output Parameters

<i>a</i>	<p>On exit, the leading m-by-n submatrix of A is set as follows:</p> <p>if <i>uplo</i> = 'U', $A(i,j) = alpha$, $1 \leq i \leq j-1$, $1 \leq j \leq n$,</p> <p>if <i>uplo</i> = 'L', $A(i,j) = alpha$, $j+1 \leq i \leq m$, $1 \leq j \leq n$,</p>
----------	---

otherwise, $A(i,j) = \text{alpha}$, $1 \leq i \leq m$, $1 \leq j \leq n$, $i \neq j$,

and, for all uplo , $A(i,i) = \text{beta}$, $1 \leq i \leq \min(m, n)$.

?lasq1

Computes the singular values of a real square bidiagonal matrix. Used by

?bdsqr.

```
call slasq1 ( n, d, e, work, info )
call dlasq1 ( n, d, e, work, info )
```

Discussion

The routine `?lasq1` computes the singular values of a real n -by- n bidiagonal matrix with diagonal d and off-diagonal e . The singular values are computed to high relative accuracy, in the absence of denormalization, underflow and overflow.

Input Parameters

n **INTEGER**. The number of rows and columns in the matrix. $n \geq 0$.

d **REAL** for `slasq1`
DOUBLE PRECISION for `dlasq1`.
 Array, **DIMENSION** (n). On entry, d contains the diagonal elements of the bidiagonal matrix whose *SVD* is desired.

e **REAL** for `slasq1`
DOUBLE PRECISION for `dlasq1`.
 Array, **DIMENSION** (n). On entry, elements $e(1:n-1)$ contain the off-diagonal elements of the bidiagonal matrix whose *SVD* is desired.

work REAL for `slasq1`
 DOUBLE PRECISION for `dlasq1`.
 Workspace array, `DIMENSION` ($4n$).

Output Parameters

d On normal exit, *d* contains the singular values in decreasing order.

e On exit, *e* is overwritten.

info INTEGER.
 = 0: successful exit;
 < 0: if *info* = -*i*, the *i*-th argument had an illegal value;
 > 0: the algorithm failed:
 = 1, a split was marked by a positive value in *e*;
 = 2, current block of *z* not diagonalized after $30 * n$ iterations (in inner while loop);
 = 3, termination criterion of outer while loop not met (program created more than *n* unreduced blocks).

?lasq2

*Computes all the eigenvalues of the symmetric positive definite tridiagonal matrix associated with the *qd* array *z* to high relative accuracy. Used by ?bdsqr and ?stegr.*

```
call slasq2 ( n, z, info )
call dlasq2 ( n, z, info )
```

Discussion

The routine `?lasq2` computes all the eigenvalues of the symmetric positive definite tridiagonal matrix associated with the *qd* array *z* to high relative accuracy, in the absence of denormalization, underflow and overflow.

To see the relation of z to the tridiagonal matrix, let L be a unit lower bidiagonal matrix with subdiagonals $z(2,4,6,\dots)$ and let U be an upper bidiagonal matrix with 1's above and diagonal $z(1,3,5,\dots)$. The tridiagonal is LU or, if you prefer, the symmetric tridiagonal to which it is similar.

Input Parameters

n **INTEGER**. The number of rows and columns in the matrix. $n \geq 0$.

z **REAL** for `slasq2`
DOUBLE PRECISION for `dlasq2`.
 Array, **DIMENSION** ($4n$). On entry, z holds the qd array.

Output Parameters

z On exit, entries 1 to n hold the eigenvalues in decreasing order, $z(2n+1)$ holds the trace, and $z(2n+2)$ holds the sum of the eigenvalues. If $n > 2$, then $z(2n+3)$ holds the iteration count, $z(2n+4)$ holds $n\text{divs}/n\text{in}^2$, and $z(2n+5)$ holds the percentage of shifts that failed.

$info$ **INTEGER**.
 = 0: successful exit;
 < 0: if the i -th argument is a scalar and had an illegal value, then $info = -i$, if the i -th argument is an array and the j -entry had an illegal value, then $info = -(i*100+j)$;
 > 0: the algorithm failed:
 = 1, a split was marked by a positive value in e ;
 = 2, current block of z not diagonalized after $30*n$ iterations (in inner while loop);
 = 3, termination criterion of outer while loop not met (program created more than n unreduced blocks).

Application Notes

The routine `?lasq2` defines a logical variable, `ieee`, which is `.TRUE.` on machines which follow IEEE-754 floating-point standard in their handling of infinities and NaNs, and `.FALSE.` otherwise. This variable is passed to `?lasq3`.

?lasq3

Checks for deflation, computes a shift and calls `dqds`. Used by `?bdsqr`.

```
call slasq3 ( i0, n0, z, pp, dmin, sigma, desig, qmax,
             nfail, iter, ndiv, ieee )
call dlasq3 ( i0, n0, z, pp, dmin, sigma, desig, qmax,
             nfail, iter, ndiv, ieee )
```

Discussion

The routine `?lasq3` checks for deflation, computes a shift (`tau`) and calls `dqds`. In case of failure, it changes shifts, and tries again until output is positive.

Input Parameters

<code>i0</code>	INTEGER. First index.
<code>n0</code>	INTEGER. Last index.
<code>z</code>	REAL for <code>slasq3</code> DOUBLE PRECISION for <code>dlasq3</code> . Array, DIMENSION (4n). <code>z</code> holds the <code>qd</code> array.
<code>pp</code>	INTEGER. <code>pp=0</code> for ping, <code>pp=1</code> for pong.
<code>desig</code>	REAL for <code>slasq3</code> DOUBLE PRECISION for <code>dlasq3</code> . Lower order part of <code>sigma</code> .

qmax REAL for `slasq3`
 DOUBLE PRECISION for `dlasq3`.
 Maximum value of q .

ieee LOGICAL. Flag for IEEE or non-IEEE arithmetic
 (passed to `?lasq5`).

Output Parameters

dmin REAL for `slasq3`
 DOUBLE PRECISION for `dlasq3`.
 Minimum value of d .

sigma REAL for `slasq3`
 DOUBLE PRECISION for `dlasq3`.
 Sum of shifts used in current segment.

desig Lower order part of *sigma*.

nfail INTEGER. Number of times shift was too big.

iter INTEGER. Number of iterations.

ndiv INTEGER. Number of divisions.

ttype INTEGER. Shift type.

?lasq4

*Computes an approximation to the
 smallest eigenvalue using values of d
 from the previous transform.*

Used by ?bdsqr.

```
call slasq4 ( i0, n0, z, pp, n0in, dmin, dmin1, dmin2,
             dn, dn1, dn2, tau, ttype )
call dlasq4 ( i0, n0, z, pp, n0in, dmin, dmin1, dmin2,
             dn, dn1, dn2, tau, ttype )
```

Discussion

The routine computes an approximation τ to the smallest eigenvalue using values of d from the previous transform.

Input Parameters

<i>i0</i>	INTEGER. First index.
<i>n0</i>	INTEGER. Last index.
<i>z</i>	REAL for <code>slasq4</code> DOUBLE PRECISION for <code>dlasq4</code> . Array, DIMENSION (4 <i>n</i>). <i>z</i> holds the <i>qd</i> array.
<i>pp</i>	INTEGER. <i>pp</i> =0 for ping, <i>pp</i> =1 for pong.
<i>noin</i>	INTEGER. The value of <i>n0</i> at start of <code>eigtest</code> .
<i>dmin</i>	REAL for <code>slasq4</code> DOUBLE PRECISION for <code>dlasq4</code> . Minimum value of <i>d</i> .
<i>dmin1</i>	REAL for <code>slasq4</code> DOUBLE PRECISION for <code>dlasq4</code> . Minimum value of <i>d</i> , excluding $d(n0)$.
<i>dmin2</i>	REAL for <code>slasq4</code> DOUBLE PRECISION for <code>dlasq4</code> . Minimum value of <i>d</i> , excluding $d(n0)$ and $d(n0-1)$.
<i>dn</i>	REAL for <code>slasq4</code> DOUBLE PRECISION for <code>dlasq4</code> . Contains $d(n)$.
<i>dn1</i>	REAL for <code>slasq4</code> DOUBLE PRECISION for <code>dlasq4</code> . Contains $d(n-1)$.
<i>dn2</i>	REAL for <code>slasq4</code> DOUBLE PRECISION for <code>dlasq4</code> . Contains $d(n-2)$.

Output Parameters

tau REAL for `slasq4`
DOUBLE PRECISION for `dlasq4`.
This is the shift.

ttype INTEGER. Shift type.

?lasq5

Computes one *dqds* transform in ping-pong form. Used by `?bdsqr` and `?stegr`.

```
call slasq5 ( i0, n0, z, pp, tau, dmin, dmin1, dmin2,
             dn, dnml, dnm2, ieee )
call dlasq5 ( i0, n0, z, pp, tau, dmin, dmin1, dmin2,
             dn, dnml, dnm2, ieee )
```

Discussion

The routine computes one *dqds* transform in ping-pong form, one version for IEEE machines another for non-IEEE machines.

Input Parameters

i0 INTEGER First index.

n0 INTEGER Last index.

z REAL for `slasq5`
DOUBLE PRECISION for `dlasq5`.
Array, DIMENSION (4*n*). *z* holds the *qd* array. *emin* is stored in *z*(4**n0*) to avoid an extra argument.

pp INTEGER. *pp*=0 for ping, *pp*=1 for pong.

tau REAL for `slasq5`
DOUBLE PRECISION for `dlasq5`.
This is the shift.

ieee LOGICAL. Flag for IEEE or non-IEEE arithmetic.

Output Parameters

dmin REAL for `slasq5`
DOUBLE PRECISION for `dlasq5`.
Minimum value of d .

dmin1 REAL for `slasq5`
DOUBLE PRECISION for `dlasq5`.
Minimum value of d , excluding $d(n0)$.

dmin2 REAL for `slasq5`
DOUBLE PRECISION for `dlasq5`.
Minimum value of d , excluding $d(n0)$ and $d(n0-1)$.

dn REAL for `slasq5`
DOUBLE PRECISION for `dlasq5`.
Contains $d(n0)$, the last value of d .

dnm1 REAL for `slasq5`
DOUBLE PRECISION for `dlasq5`.
Contains $d(n0-1)$.

dnm2 REAL for `slasq5`
DOUBLE PRECISION for `dlasq5`.
Contains $d(n0-2)$.

?lasq6

Computes one dqds transform in ping-pong form. Used by ?bdsqr and ?stegr.

```
call slasq6 ( i0, n0, z, pp, dmin, dmin1, dmin2, dn,
             dnm1, dnm2 )
call dlasq6 ( i0, n0, z, pp, dmin, dmin1, dmin2, dn,
             dnm1, dnm2 )
```

Discussion

The routine `?lasq6` computes one dqd (shift equal to zero) transform in ping-pong form, with protection against underflow and overflow.

Input Parameters

i0 **INTEGER**. First index.
n0 **INTEGER**. Last index.
z **REAL** for `slasq6`
 DOUBLE PRECISION for `dlasq6`.
 Array, **DIMENSION** (4*n*). *z* holds the qd array. *emin* is
 stored in $z(4*n0)$ to avoid an extra argument.
pp **INTEGER**. *pp*=0 for ping, *pp*=1 for pong.

Output Parameters

dmin **REAL** for `slasq6`
 DOUBLE PRECISION for `dlasq6`.
 Minimum value of d .
dmin1 **REAL** for `slasq6`
 DOUBLE PRECISION for `dlasq6`.
 Minimum value of d , excluding $d(n0)$.
dmin2 **REAL** for `slasq6`
 DOUBLE PRECISION for `dlasq6`.
 Minimum value of d , excluding $d(n0)$ and $d(n0-1)$.
dn **REAL** for `slasq6`
 DOUBLE PRECISION for `dlasq6`.
 Contains $d(n0)$, the last value of d .
dnm1 **REAL** for `slasq6`
 DOUBLE PRECISION for `dlasq6`.
 Contains $d(n0-1)$.
dnm2 **REAL** for `slasq6`
 DOUBLE PRECISION for `dlasq6`.
 Contains $d(n0-2)$.

?lasr

Applies a sequence of plane rotations to a general rectangular matrix.

```
call slasr ( side, pivot, direct, m, n, c, s, a, lda )
call dlasr ( side, pivot, direct, m, n, c, s, a, lda )
call clasr ( side, pivot, direct, m, n, c, s, a, lda )
call zlasr ( side, pivot, direct, m, n, c, s, a, lda )
```

Discussion

The routine performs the transformation:

$A := P A$, when *side* = 'L' or 'l' (Left-hand side)

$A := A P'$, when *side* = 'R' or 'r' (Right-hand side)

where A is an m -by- n real matrix and P is an orthogonal matrix, consisting of a sequence of plane rotations determined by the parameters *pivot* and *direct* as follows ($z = m$ when *side* = 'L' or 'l' and $z = n$ when *side* = 'R' or 'r'):

When *direct* = 'F' or 'f' (Forward sequence) then

$$P = P(z - 1) \dots P(2) P(1),$$

and when *direct* = 'B' or 'b' (Backward sequence) then

$$P = P(1) P(2) \dots P(z - 1),$$

where $P(k)$ is a plane rotation matrix for the following planes:

when *pivot* = 'V' or 'v' (Variable pivot), the plane $(k, k + 1)$

when *pivot* = 'T' or 't' (Top pivot), the plane $(1, k + 1)$

when *pivot* = 'B' or 'b' (Bottom pivot), the plane (k, z)

$c(k)$ and $s(k)$ must contain the cosine and sine that define the matrix $P(k)$. The 2-by-2 plane rotation part of the matrix $P(k)$, $R(k)$, is assumed to be of the form:

$$R(k) = \begin{bmatrix} c(k) & s(k) \\ -s(k) & c(k) \end{bmatrix}$$

Input Parameters

<i>side</i>	CHARACTER*1 . Specifies whether the plane rotation matrix P is applied to A on the left or the right. = 'L': Left, compute $A := P A$ = 'R': Right, compute $A := A P'$
<i>direct</i>	CHARACTER*1 . Specifies whether P is a forward or backward sequence of plane rotations. = 'F': Forward, $P = P(z-1) \dots P(2) P(1)$ = 'B': Backward, $P = P(1) P(2) \dots P(z-1)$
<i>pivot</i>	CHARACTER*1 . Specifies the plane for which $P(k)$ is a plane rotation matrix. = 'V': Variable pivot, the plane $(k, k+1)$ = 'T': Top pivot, the plane $(1, k+1)$ = 'B': Bottom pivot, the plane (k, z)
<i>m</i>	INTEGER . The number of rows of the matrix A . If $m \leq 1$, an immediate return is effected.
<i>n</i>	INTEGER . The number of columns of the matrix A . If $n \leq 1$, an immediate return is effected.
<i>c, s</i>	REAL for slasr/clasr DOUBLE PRECISION for dlasr/zlasr . Arrays, DIMENSION $(m-1)$ if <i>side</i> = 'L', $(n-1)$ if <i>side</i> = 'R'. $c(k)$ and $s(k)$ contain the cosine and sine that define the matrix $P(k)$ as described above.
<i>a</i>	REAL for slasr DOUBLE PRECISION for dlasr COMPLEX for clasr COMPLEX*16 for zlasr . Array, DIMENSION (<i>lda</i> , <i>n</i>). The m -by- n matrix A .
<i>lda</i>	INTEGER . The leading dimension of the array A . $lda \geq \max(1, m)$.

Output Parameters

a On exit, *A* is overwritten by *PA* if *side* = 'R' or by *AP* if *side* = 'L'.

?lasrt

Sorts numbers in increasing or decreasing order.

```
call slasrt ( id, n, d, info )
call dlasrt ( id, n, d, info )
```

Discussion

The routine `?lasrt` sorts the numbers in *d* in increasing order (if *id* = 'I') or in decreasing order (if *id* = 'D'). It uses Quick Sort, reverting to Insertion Sort on arrays of size ≤ 20 . Dimension of *stack* limits *n* to about 2^{32} .

Input Parameters

id CHARACTER*1.
= 'I': sort *d* in increasing order;
= 'D': sort *d* in decreasing order.

n INTEGER. The length of the array *d*.

d REAL for `slasrt`
DOUBLE PRECISION for `dlasrt`.
On entry, the array to be sorted.

Output Parameters

d On exit, *d* has been sorted into increasing order ($d(1) \leq \dots \leq d(n)$) or into decreasing order ($d(1) \geq \dots \geq d(n)$), depending on *id*.

info INTEGER.
= 0: successful exit
< 0: if *info* = -*i*, the *i*-th argument had an illegal value.

?lassq

Updates a sum of squares represented in scaled form.

```
call slassq ( n, x, incx, scale, sumsq )
call dlassq ( n, x, incx, scale, sumsq )
call classq ( n, x, incx, scale, sumsq )
call zlassq ( n, x, incx, scale, sumsq )
```

Discussion

The real routines `slassq/dlassq` return the values `scl` and `smsq` such that

$$scl^2 * smsq = x(1)^2 + \dots + x(n)^2 + scale^2 * sumsq,$$

where $x(i) = x(1 + (i - 1) incx)$.

The value of `sumsq` is assumed to be non-negative and `scl` returns the value

$$scl = \max(scale, \max_i \text{abs}(x(i))).$$

Values `scale` and `sumsq` must be supplied in `scale` and `sumsq`, and `scl` and `smsq` are overwritten on `scale` and `sumsq`, respectively.

The complex routines `classq/zlassq` return the values `scl` and `ssq` such that

$$scl^2 * ssq = x(1)^2 + \dots + x(n)^2 + scale^2 * sumsq,$$

where $x(i) = \text{abs}(x(1 + (i - 1) incx))$.

The value of `sumsq` is assumed to be at least unity and the value of `ssq` will then satisfy

$$1.0 \leq ssq \leq sumsq + 2n$$

`scale` is assumed to be non-negative and `scl` returns the value

$$scl = \max_i (scale, \text{abs}(\text{real}(x(i))), \text{abs}(\text{aimag}(x(i)))).$$

Values `scale` and `sumsq` must be supplied in `scale` and `sumsq`, and `scl` and `ssq` are overwritten on `scale` and `sumsq`, respectively.

All routines `?lassq` make only one pass through the vector `x`.

Input Parameters

<i>n</i>	INTEGER. The number of elements to be used from the vector <i>x</i> .
<i>x</i>	REAL for <code>slassq</code> DOUBLE PRECISION for <code>dlassq</code> COMPLEX for <code>classq</code> COMPLEX*16 for <code>zlassq</code> . The vector for which a scaled sum of squares is computed: $x(i) = x(1 + (i - 1) \textit{incx})$, $1 \leq i \leq n$.
<i>incx</i>	INTEGER. The increment between successive values of the vector <i>x</i> . <i>incx</i> > 0.
<i>scale</i>	REAL for <code>slassq/classq</code> DOUBLE PRECISION for <code>dlassq/zlassq</code> . On entry, the value <i>scale</i> in the equation above.
<i>sumsq</i>	REAL for <code>slassq/classq</code> DOUBLE PRECISION for <code>dlassq/zlassq</code> . On entry, the value <i>sumsq</i> in the equation above.

Output Parameters

<i>scale</i>	On exit, <i>scale</i> is overwritten with <i>scl</i> , the scaling factor for the sum of squares.
<i>sumsq</i>	<i>For real flavors:</i> On exit, <i>sumsq</i> is overwritten with the value <i>sumsq</i> in the equation above. <i>For complex flavors:</i> On exit, <i>sumsq</i> is overwritten with the value <i>ssq</i> in the equation above.

?lasv2

Computes the singular value decomposition of a 2-by-2 triangular matrix

```
call slasv2 ( f, g, h, ssmín, ssmáx, snr, csr, snl, csl )
call dlasv2 ( f, g, h, ssmín, ssmáx, snr, csr, snl, csl )
```

Discussion

The routine ?lasv2 computes the singular value decomposition of a 2-by-2 triangular matrix

$$\begin{bmatrix} f & g \\ 0 & h \end{bmatrix}$$

On return, $\text{abs}(ssmax)$ is the larger singular value, $\text{abs}(ssmin)$ is the smaller singular value, and (csl,snl) and (csr,snr) are the left and right singular vectors for $\text{abs}(ssmax)$, giving the decomposition

$$\begin{bmatrix} csl & snl \\ -snl & csl \end{bmatrix} \begin{bmatrix} f & g \\ 0 & h \end{bmatrix} \begin{bmatrix} csr & -snr \\ snr & csr \end{bmatrix} = \begin{bmatrix} ssmáx & 0 \\ 0 & ssmín \end{bmatrix}$$

Input Parameters

f, g, h REAL for slasv2
 DOUBLE PRECISION for dlasv2.
The (1,1), (1,2) and (2,2) elements of the 2-by-2 matrix, respectively.

Output Parameters

ssmin, ssmáx REAL for slasv2
 DOUBLE PRECISION for dlasv2.
 $\text{abs}(ssmin)$ and $\text{abs}(ssmáx)$ is the smaller and the larger singular value, respectively.

<i>snl, csl</i>	<p>REAL for <code>slasv2</code></p> <p>DOUBLE PRECISION for <code>dlasv2</code>.</p> <p>The vector (<i>csl</i>, <i>snl</i>) is a unit left singular vector for the singular value <code>abs(ssmax)</code>.</p>
<i>snr, csr</i>	<p>REAL for <code>slasv2</code></p> <p>DOUBLE PRECISION for <code>dlasv2</code>.</p> <p>The vector (<i>csr</i>, <i>snr</i>) is a unit right singular vector for the singular value <code>abs(ssmax)</code>.</p>

Application Notes

Any input parameter may be aliased with any output parameter. Barring over/underflow and assuming a guard digit in subtraction, all output quantities are correct to within a few units in the last place (ulps).

In IEEE arithmetic, the code works correctly if one matrix element is infinite.

Overflow will not occur unless the largest singular value itself overflows or is within a few ulps of overflow. (On machines with partial overflow, like the Cray, overflow may occur if the largest singular value is within a factor of 2 of overflow.)

Underflow is harmless if underflow is gradual. Otherwise, results may correspond to a matrix modified by perturbations of size near the underflow threshold.

?laswp

Performs a series of row interchanges on a general rectangular matrix.

```
call slaswp ( n, a, lda, k1, k2, ipiv, incx )
call dlaswp ( n, a, lda, k1, k2, ipiv, incx )
call claswp ( n, a, lda, k1, k2, ipiv, incx )
call zlaswp ( n, a, lda, k1, k2, ipiv, incx )
```

Discussion

The routine performs a series of row interchanges on the matrix A . One row interchange is initiated for each of rows $k1$ through $k2$ of A .

Input Parameters

n **INTEGER**. The number of columns of the matrix A .

a **REAL** for `slaswp`
DOUBLE PRECISION for `dlaswp`
COMPLEX for `claswp`
COMPLEX*16 for `zlaswp`.
Array, **DIMENSION** (lda, n).
On entry, the matrix of column dimension n to which the row interchanges will be applied.

lda **INTEGER**. The leading dimension of the array a .

$k1$ **INTEGER**. The first element of $ipiv$ for which a row interchange will be done.

$k2$ **INTEGER**. The last element of $ipiv$ for which a row interchange will be done.

$ipiv$ **INTEGER**.
Array, **DIMENSION** ($m * \text{abs}(incx)$).
The vector of pivot indices. Only the elements in positions $k1$ through $k2$ of $ipiv$ are accessed.
 $ipiv(k) = l$ implies rows k and l are to be interchanged.

$incx$ **INTEGER**. The increment between successive values of $ipiv$. If $ipiv$ is negative, the pivots are applied in reverse order.

Output Parameters

a On exit, the permuted matrix.

?lasy2

Solves the Sylvester matrix equation
where the matrices are of order 1 or 2.

```
call slasy2 ( ltranl, ltranr, isgn, n1, n2, tl, ldtl,
             tr, ldtr, b, ldb, scale, x, ldx, xnorm, info )
call dlasy2 ( ltranl, ltranr, isgn, n1, n2, tl, ldtl,
             tr, ldtr, b, ldb, scale, x, ldx, xnorm, info )
```

Discussion

The routine solves for the $n1$ -by- $n2$ matrix X , $1 \leq n1, n2 \leq 2$, in

$$\text{op}(TL) * X + \text{isgn} * X * \text{op}(TR) = \text{scale} * B,$$

where

TL is $n1$ -by- $n1$,

TR is $n2$ -by- $n2$,

B is $n1$ -by- $n2$,

and $\text{isgn} = 1$ or -1 . $\text{op}(T) = T$ or T' , where T' denotes the transpose of T .

Input Parameters

ltranl LOGICAL.
On entry, **ltranl** specifies the $\text{op}(TL)$:
= **.FALSE.**, $\text{op}(TL) = TL$,
= **.TRUE.**, $\text{op}(TL) = TL'$.

ltranr LOGICAL.
On entry, **ltranr** specifies the $\text{op}(TR)$:
= **.FALSE.**, $\text{op}(TR) = TR$,
= **.TRUE.**, $\text{op}(TR) = TR'$.

isgn INTEGER. On entry, **isgn** specifies the sign of the equation as described before. **isgn** may only be 1 or -1.

n1 INTEGER. On entry, **n1** specifies the order of matrix TL . **n1** may only be 0, 1 or 2.

n2 **INTEGER.**
On entry, *n2* specifies the order of matrix *TR*.
n2 may only be 0, 1 or 2.

t1 **REAL** for *slasy2*
DOUBLE PRECISION for *dlasy2*.
Array, **DIMENSION** (*ldt1*,2). On entry, *t1* contains an
n1-by-*n1* matrix *TL*.

ldt1 **INTEGER.**The leading dimension of the matrix *t1*.
ldt1 \geq $\max(1, n1)$.

tr **REAL** for *slasy2*
DOUBLE PRECISION for *dlasy2*.
Array, **DIMENSION** (*ldtr*,2). On entry, *tr* contains an
n2-by-*n2* matrix *TR*.

ldtr **INTEGER.**
The leading dimension of the matrix *tr*.
ldtr \geq $\max(1, n2)$.

b **REAL** for *slasy2*
DOUBLE PRECISION for *dlasy2*.
Array, **DIMENSION** (*ldb*,2). On entry, the *n1*-by-*n2*
matrix *b* contains the right-hand side of the equation.

ldb **INTEGER.**
The leading dimension of the matrix *b*.
ldb \geq $\max(1, n1)$.

ldx **INTEGER.**
The leading dimension of the output matrix *x*.
ldx \geq $\max(1, n1)$.

Output Parameters

scale **REAL** for *slasy2*
DOUBLE PRECISION for *dlasy2*.
On exit, *scale* contains the scale factor.
scale is chosen less than or equal to 1 to prevent the
solution overflowing.

<code>x</code>	REAL for <code>slasy2</code> DOUBLE PRECISION for <code>dlasy2</code> . Array, DIMENSION (<code>ldx</code> ,2). On exit, <code>x</code> contains the <code>n1</code> -by- <code>n2</code> solution.
<code>xnorm</code>	REAL for <code>slasy2</code> DOUBLE PRECISION for <code>dlasy2</code> . On exit, <code>xnorm</code> is the infinity-norm of the solution.
<code>info</code>	INTEGER. On exit, <code>info</code> is set to 0: successful exit. 1: <code>TL</code> and <code>TR</code> have too close eigenvalues, so <code>TL</code> or <code>TR</code> is perturbed to get a nonsingular equation.



NOTE. *In the interests of speed, this routine does not check the inputs for errors.*

?lasyf

Computes a partial factorization of a real/complex symmetric matrix, using the diagonal pivoting method.

```
call slasyf ( uplo, n, nb, kb, a, lda, ipiv, w, ldw, info)
call dlasyf ( uplo, n, nb, kb, a, lda, ipiv, w, ldw, info)
call clasyf ( uplo, n, nb, kb, a, lda, ipiv, w, ldw, info)
call zlasyf ( uplo, n, nb, kb, a, lda, ipiv, w, ldw, info)
```

Discussion

The routine `?lasyf` computes a partial factorization of a real/complex symmetric matrix `A` using the Bunch-Kaufman diagonal pivoting method. The partial factorization has the form:

$$A = \begin{bmatrix} I & U_{12} \\ 0 & U_{22} \end{bmatrix} \begin{bmatrix} A_{11} & 0 \\ 0 & D \end{bmatrix} \begin{bmatrix} I & 0 \\ U_{12}' & U_{22}' \end{bmatrix} \quad \text{if } uplo = 'U', \text{ or}$$

$$A = \begin{bmatrix} L_{11} & 0 \\ L_{21} & I \end{bmatrix} \begin{bmatrix} D & 0 \\ 0 & A_{22} \end{bmatrix} \begin{bmatrix} L_{11}' & L_{21}' \\ 0 & I \end{bmatrix} \quad \text{if } uplo = 'L'$$

where the order of D is at most nb . The actual order is returned in the argument kb , and is either nb or $nb-1$, or n if $n \leq nb$.

This is an auxiliary routine called by `?sytrf`. It uses blocked code (calling Level 3 BLAS) to update the submatrix A_{11} (if $uplo = 'U'$) or A_{22} (if $uplo = 'L'$).

Input Parameters

<i>uplo</i>	CHARACTER*1. Specifies whether the upper or lower triangular part of the symmetric matrix A is stored: = 'U': Upper triangular = 'L': Lower triangular
<i>n</i>	INTEGER. The order of the matrix A . $n \geq 0$.
<i>nb</i>	INTEGER. The maximum number of columns of the matrix A that should be factored. nb should be at least 2 to allow for 2-by-2 pivot blocks.
<i>a</i>	REAL for <code>slasyf</code> DOUBLE PRECISION for <code>dlasyf</code> COMPLEX for <code>clasyf</code> COMPLEX*16 for <code>zlasyf</code> . Array, DIMENSION (lda, n). On entry, the symmetric matrix A . If $uplo = 'U'$, the leading n -by- n upper triangular part of a contains the upper triangular part of the matrix A , and the strictly lower triangular part of a is not referenced. If $uplo = 'L'$, the leading n -by- n lower

triangular part of *a* contains the lower triangular part of the matrix *A*, and the strictly upper triangular part of *a* is not referenced.

lda **INTEGER.**
 The leading dimension of the array *a*. $lda \geq \max(1,n)$.

w **REAL** for **slasyf**
DOUBLE PRECISION for **dlasyf**
COMPLEX for **clasyf**
COMPLEX*16 for **zlasyf**.
 Workspace array, **DIMENSION** (*ldw*, *nb*).

ldw **INTEGER.**
 The leading dimension of the array *w*. $ldw \geq \max(1,n)$.

Output Parameters

kb **INTEGER.**
 The number of columns of *A* that were actually factored
kb is either *nb*-1 or *nb*, or *n* if $n \leq nb$.

a On exit, *a* contains details of the partial factorization.

ipiv **INTEGER.**
 Array, **DIMENSION** (*n*). Details of the interchanges and the block structure of *D*.
 If *uplo* = 'U', only the last *kb* elements of *ipiv* are set;
 if *uplo* = 'L', only the first *kb* elements are set.
 If *ipiv*(*k*) > 0, then rows and columns *k* and *ipiv*(*k*) were interchanged and *D*(*k*,*k*) is a 1-by-1 diagonal block.
 If *uplo* = 'U' and *ipiv*(*k*) = *ipiv*(*k*-1) < 0, then rows and columns *k*-1 and -*ipiv*(*k*) were interchanged and *D*(*k*-1:*k*, *k*-1:*k*) is a 2-by-2 diagonal block.
 If *uplo* = 'L' and *ipiv*(*k*) = *ipiv*(*k*+1) < 0, then rows and columns *k*+1 and -*ipiv*(*k*) were interchanged and *D*(*k*:*k*+1, *k*:*k*+1) is a 2-by-2 diagonal block.

info **INTEGER**.
 = 0: successful exit
 > 0: if *info* = *k*, $D(k,k)$ is exactly zero. The factorization has been completed, but the block diagonal matrix D is exactly singular.

?lahef

Computes a partial factorization of a complex Hermitian indefinite matrix, using the diagonal pivoting method.

```
call clahef ( uplo, n, nb, kb, a, lda, ipiv, w, ldw, info)
call zlahef ( uplo, n, nb, kb, a, lda, ipiv, w, ldw, info)
```

Discussion

The routine **?lahef** computes a partial factorization of a complex Hermitian matrix A , using the Bunch-Kaufman diagonal pivoting method. The partial factorization has the form:

$$A = \begin{bmatrix} I & U_{12} \\ 0 & U_{22} \end{bmatrix} \begin{bmatrix} A_{11} & 0 \\ 0 & D \end{bmatrix} \begin{bmatrix} I & 0 \\ U_{12}' & U_{22}' \end{bmatrix} \quad \text{if } uplo = 'U', \text{ or}$$

$$A = \begin{bmatrix} L_{11} & 0 \\ L_{21} & I \end{bmatrix} \begin{bmatrix} D & 0 \\ 0 & A_{22} \end{bmatrix} \begin{bmatrix} L_{11}' & L_{21}' \\ 0 & I \end{bmatrix} \quad \text{if } uplo = 'L'$$

where the order of D is at most *nb*. The actual order is returned in the argument *kb*, and is either *nb* or *nb*-1, or *n* if $n \leq nb$. Note that U' denotes the conjugate transpose of U .

This is an auxiliary routine called by **?hetrf**. It uses blocked code (calling Level 3 BLAS) to update the submatrix A_{11} (if *uplo* = 'U') or A_{22} (if *uplo* = 'L').

Input Parameters

<i>uplo</i>	<p>CHARACTER*1.</p> <p>Specifies whether the upper or lower triangular part of the Hermitian matrix <i>A</i> is stored:</p> <p>= 'U': Upper triangular</p> <p>= 'L': Lower triangular</p>
<i>n</i>	<p>INTEGER.</p> <p>The order of the matrix <i>A</i>. $n \geq 0$.</p>
<i>nb</i>	<p>INTEGER.</p> <p>The maximum number of columns of the matrix <i>A</i> that should be factored. <i>nb</i> should be at least 2 to allow for 2-by-2 pivot blocks.</p>
<i>a</i>	<p>COMPLEX for <code>clahef</code></p> <p>COMPLEX*16 for <code>zlahef</code>.</p> <p>Array, DIMENSION (<i>lda</i>, <i>n</i>).</p> <p>On entry, the Hermitian matrix <i>A</i>. If <i>uplo</i> = 'U', the leading <i>n</i>-by-<i>n</i> upper triangular part of <i>a</i> contains the upper triangular part of the matrix <i>A</i>, and the strictly lower triangular part of <i>a</i> is not referenced.</p> <p>If <i>uplo</i> = 'L', the leading <i>n</i>-by-<i>n</i> lower triangular part of <i>a</i> contains the lower triangular part of the matrix <i>A</i>, and the strictly upper triangular part of <i>a</i> is not referenced.</p>
<i>lda</i>	<p>INTEGER.</p> <p>The leading dimension of the array <i>a</i>. $lda \geq \max(1, n)$.</p>
<i>w</i>	<p>COMPLEX for <code>clahef</code></p> <p>COMPLEX*16 for <code>zlahef</code>.</p> <p>Workspace array, DIMENSION (<i>ldw</i>, <i>nb</i>).</p>
<i>ldw</i>	<p>INTEGER.</p> <p>The leading dimension of the array <i>w</i>. $ldw \geq \max(1, nb)$.</p>

Output Parameters

<i>kb</i>	<p>INTEGER.</p> <p>The number of columns of <i>A</i> that were actually factored</p> <p><i>kb</i> is either <i>nb</i>-1 or <i>nb</i>, or <i>n</i> if $n \leq nb$.</p>
-----------	--

a On exit, *a* contains details of the partial factorization.

ipiv **INTEGER.**
 Array, **DIMENSION** (*n*). Details of the interchanges and the block structure of *D*.
 If *uplo* = 'U', only the last *kb* elements of *ipiv* are set;
 if *uplo* = 'L', only the first *kb* elements are set.
 If *ipiv*(*k*) > 0, then rows and columns *k* and *ipiv*(*k*) were interchanged and *D*(*k*,*k*) is a 1-by-1 diagonal block.
 If *uplo* = 'U' and *ipiv*(*k*) = *ipiv*(*k*-1) < 0, then rows and columns *k*-1 and -*ipiv*(*k*) were interchanged and *D*(*k*-1:*k*, *k*-1:*k*) is a 2-by-2 diagonal block.
 If *uplo* = 'L' and *ipiv*(*k*) = *ipiv*(*k*+1) < 0, then rows and columns *k*+1 and -*ipiv*(*k*) were interchanged and *D*(*k*:*k*+1, *k*:*k*+1) is a 2-by-2 diagonal block.

info **INTEGER.**
 = 0: successful exit
 > 0: if *info* = *k*, *D*(*k*,*k*) is exactly zero. The factorization has been completed, but the block diagonal matrix *D* is exactly singular.

?latbs

Solves a triangular banded system of equations.

```

call slatbs ( uplo, trans, diag, normin, n, kd, ab,
              ldab, x, scale, cnorm, info )
call dlatbs ( uplo, trans, diag, normin, n, kd, ab,
              ldab, x, scale, cnorm, info )
call clatbs ( uplo, trans, diag, normin, n, kd, ab,
              ldab, x, scale, cnorm, info )
call zlatbs ( uplo, trans, diag, normin, n, kd, ab,
              ldab, x, scale, cnorm, info )

```

Discussion

The routine solves one of the triangular systems

$Ax = s b$ or $A^T x = s b$ or $A^H x = s b$ (for complex flavors)

with scaling to prevent overflow, where A is an upper or lower triangular band matrix. Here A^T denotes the transpose of A , A^H denotes the conjugate transpose of A , x and b are n -element vectors, and s is a scaling factor, usually less than or equal to 1, chosen so that the components of x will be less than the overflow threshold. If the unscaled problem will not cause overflow, the Level 2 BLAS routine `?tbsv` is called. If the matrix A is singular ($A(j, j) = 0$ for some j), then s is set to 0 and a non-trivial solution to $Ax = 0$ is returned.

Input Parameters

<i>uplo</i>	CHARACTER*1. Specifies whether the matrix A is upper or lower triangular. = 'U': Upper triangular = 'L': Lower triangular
<i>trans</i>	CHARACTER*1. Specifies the operation applied to A . = 'N': Solve $Ax = s b$ (no transpose) = 'T': Solve $A^T x = s b$ (transpose) = 'C': Solve $A^H x = s b$ (conjugate transpose)
<i>diag</i>	CHARACTER*1. Specifies whether or not the matrix A is unit triangular = 'N': Non-unit triangular = 'U': Unit triangular
<i>normin</i>	CHARACTER*1. Specifies whether <i>cnorm</i> has been set or not. = 'Y': <i>cnorm</i> contains the column norms on entry; = 'N': <i>cnorm</i> is not set on entry. On exit, the norms will be computed and stored in <i>cnorm</i> .
<i>n</i>	INTEGER. The order of the matrix A . $n \geq 0$.

kd INTEGER.
The number of subdiagonals or superdiagonals in the triangular matrix *A*. $kd \geq 0$.

ab REAL for *slatbs*
DOUBLE PRECISION for *dlatbs*
COMPLEX for *clatbs*
COMPLEX*16 for *zlatbs*.
Array, DIMENSION (*ldab*, *n*). The upper or lower triangular band matrix *A*, stored in the first $kd+1$ rows of the array. The *j*-th column of *A* is stored in the *j*-th column of the array *ab* as follows:
if *uplo* = 'U', $ab(kd+1+i-j, j) = A(i, j)$ for $\max(1, j-kd) \leq i \leq j$;
if *uplo* = 'L', $ab(1+i-j, j) = A(i, j)$ for $j \leq i \leq \min(n, j+kd)$.

ldab INTEGER.
The leading dimension of the array *ab*. $ldab \geq kd+1$.

x REAL for *slatbs*
DOUBLE PRECISION for *dlatbs*
COMPLEX for *clatbs*
COMPLEX*16 for *zlatbs*.
Array, DIMENSION (*n*).
On entry, the right hand side *b* of the triangular system.

cnorm REAL for *slatbs/clatbs*
DOUBLE PRECISION for *dlatbs/zlatbs*.
Array, DIMENSION (*n*).
If *normin* = 'Y', *cnorm* is an input argument and *cnorm*(*j*) contains the norm of the off-diagonal part of the *j*-th column of *A*. If *trans* = 'N', *cnorm*(*j*) must be greater than or equal to the infinity-norm, and if *trans* = 'T' or 'C', *cnorm*(*j*) must be greater than or equal to the 1-norm.

Output Parameters

<i>scale</i>	<p>REAL for <code>slatbs/clatbs</code> DOUBLE PRECISION for <code>dlatbs/zlatbs</code>.</p> <p>The scaling factor s for the triangular system as described above.</p> <p>If <i>scale</i> = 0, the matrix A is singular or badly scaled, and the vector x is an exact or approximate solution to $Ax = 0$.</p>
<i>cnorm</i>	<p>If <i>normin</i> = 'N', <i>cnorm</i> is an output argument and <i>cnorm</i>(j) returns the 1-norm of the off-diagonal part of the j-th column of A.</p>
<i>info</i>	<p>INTEGER.</p> <p>= 0: successful exit < 0: if <i>info</i> = $-k$, the k-th argument had an illegal value</p>

?latdf

Uses the LU factorization of the n -by- n matrix computed by `?getc2` and computes a contribution to the reciprocal Dif-estimate.

```
call slatdf ( ijob, n, z, ldz, rhs, rdsum, rdscal, ipiv, jpiv )
call dlatdf ( ijob, n, z, ldz, rhs, rdsum, rdscal, ipiv, jpiv )
call clatdf ( ijob, n, z, ldz, rhs, rdsum, rdscal, ipiv, jpiv )
call zlatdf ( ijob, n, z, ldz, rhs, rdsum, rdscal, ipiv, jpiv )
```

Discussion

The routine `?latdf` uses the LU factorization of the n -by- n matrix Z computed by `?getc2` and computes a contribution to the reciprocal Dif-estimate by solving $Zx = b$ for x , and choosing the right-hand side b such that the norm of x is as large as possible. On entry *rhs* = b holds the contribution from earlier solved sub-systems, and on return *rhs* = x .

The factorization of Z returned by `?getc2` has the form $Z = P L U Q$, where P and Q are permutation matrices. L is lower triangular with unit diagonal elements and U is upper triangular.

Input Parameters

<i>ijob</i>	<p>INTEGER.</p> <p><i>ijob</i> = 2: First compute an approximative null-vector e of Z using <code>?gecon</code>, e is normalized, and solve for $Zx = \pm e - f$ with the sign giving the greater value of $2\text{-norm}(x)$. This option is about 5 times as expensive as default.</p> <p><i>ijob</i> \neq 2 (default): Local look ahead strategy where all entries of the right-hand side b is chosen as either +1 or -1 .</p>
<i>n</i>	<p>INTEGER.</p> <p>The number of columns of the matrix Z.</p>
<i>z</i>	<p>REAL for <code>slatdf/clatdf</code> DOUBLE PRECISION for <code>dlatdf/zlatdf</code>.</p> <p>Array, DIMENSION (<i>ldz</i>, <i>n</i>)</p> <p>On entry, the LU part of the factorization of the n-by-n matrix Z computed by <code>?getc2</code>: $Z = P L U Q$.</p>
<i>ldz</i>	<p>INTEGER.</p> <p>The leading dimension of the array z. $lda \geq \max(1, n)$.</p>
<i>rhs</i>	<p>REAL for <code>slatdf/clatdf</code> DOUBLE PRECISION for <code>dlatdf/zlatdf</code>.</p> <p>Array, DIMENSION (<i>n</i>).</p> <p>On entry, <i>rhs</i> contains contributions from other subsystems.</p>
<i>rdsum</i>	<p>REAL for <code>slatdf/clatdf</code> DOUBLE PRECISION for <code>dlatdf/zlatdf</code>.</p> <p>On entry, the sum of squares of computed contributions to the Dif-estimate under computation by <code>?tgsyl</code>, where the scaling factor <i>rdscal</i> has been factored out.</p>

If *trans* = 'T', *rdsum* is not touched.

Note that *rdsum* only makes sense when *?tgsy2* is called by *?tgsyl*.

<i>rdscal</i>	<p>REAL for <i>slatdf/clatdf</i> DOUBLE PRECISION for <i>dlatdf/zlatdf</i>. On entry, scaling factor used to prevent overflow in <i>rdsum</i>. If <i>trans</i> = 'T', <i>rdscal</i> is not touched. Note that <i>rdscal</i> only makes sense when <i>?tgsy2</i> is called by <i>?tgsyl</i>.</p>
<i>ipiv</i>	<p>INTEGER. Array, DIMENSION (<i>n</i>). The pivot indices; for $1 \leq i \leq n$, row <i>i</i> of the matrix has been interchanged with row <i>ipiv</i>(<i>i</i>).</p>
<i>jpiv</i>	<p>INTEGER. Array, DIMENSION (<i>n</i>). The pivot indices; for $1 \leq j \leq n$, column <i>j</i> of the matrix has been interchanged with column <i>jpiv</i>(<i>j</i>).</p>

Output Parameters

<i>rhs</i>	<p>On exit, <i>rhs</i> contains the solution of the subsystem with entries according to the value of <i>ijob</i>.</p>
<i>rdsum</i>	<p>On exit, the corresponding sum of squares updated with the contributions from the current sub-system. If <i>trans</i> = 'T', <i>rdsum</i> is not touched.</p>
<i>rdscal</i>	<p>On exit, <i>rdscal</i> is updated with respect to the current contributions in <i>rdsum</i>. If <i>trans</i> = 'T', <i>rdscal</i> is not touched.</p>

?latps

*Solves a triangular system of equations
with the matrix held in packed storage.*

```
call slatps (uplo, trans, diag, normin, n, ap, x, scale, cnorm, info)
call dlatps (uplo, trans, diag, normin, n, ap, x, scale, cnorm, info)
call clatps (uplo, trans, diag, normin, n, ap, x, scale, cnorm, info)
call zlatps (uplo, trans, diag, normin, n, ap, x, scale, cnorm, info)
```

Discussion

The routine `?latps` solves one of the triangular systems

$Ax = s b$ or $A^T x = s b$ or $A^H x = s b$ (for complex flavors)

with scaling to prevent overflow, where A is an upper or lower triangular matrix stored in packed form. Here A^T denotes the transpose of A , A^H denotes the conjugate transpose of A , x and b are n -element vectors, and s is a scaling factor, usually less than or equal to 1, chosen so that the components of x will be less than the overflow threshold. If the unscaled problem will not cause overflow, the Level 2 BLAS routine `?tpsv` is called. If the matrix A is singular ($A(j, j) = 0$ for some j), then s is set to 0 and a non-trivial solution to $Ax = 0$ is returned.

Input Parameters

`uplo` CHARACTER*1.
Specifies whether the matrix A is upper or lower triangular.
= 'U': Upper triangular
= 'L': Lower triangular

`trans` CHARACTER*1.
Specifies the operation applied to A .
= 'N': Solve $Ax = s b$ (no transpose)
= 'T': Solve $A^T x = s b$ (transpose)
= 'C': Solve $A^H x = s b$ (conjugate transpose)

<i>diag</i>	<p>CHARACTER*1.</p> <p>Specifies whether or not the matrix <i>A</i> is unit triangular.</p> <p>= 'N': Non-unit triangular</p> <p>= 'U': Unit triangular</p>
<i>normin</i>	<p>CHARACTER*1.</p> <p>Specifies whether <i>cnorm</i> has been set or not.</p> <p>= 'Y': <i>cnorm</i> contains the column norms on entry;</p> <p>= 'N': <i>cnorm</i> is not set on entry. On exit, the norms will be computed and stored in <i>cnorm</i>.</p>
<i>n</i>	<p>INTEGER.</p> <p>The order of the matrix <i>A</i>. $n \geq 0$.</p>
<i>ap</i>	<p>REAL for <i>slatps</i></p> <p>DOUBLE PRECISION for <i>dlatps</i></p> <p>COMPLEX for <i>clatps</i></p> <p>COMPLEX*16 for <i>zlatps</i>.</p> <p>Array, DIMENSION ($n(n+1)/2$). The upper or lower triangular matrix <i>A</i>, packed columnwise in a linear array. The <i>j</i>-th column of <i>A</i> is stored in the array <i>ap</i> as follows:</p> <p>if <i>uplo</i> = 'U', $ap(i + (j-1)j/2) = A(i, j)$ for $1 \leq i \leq j$;</p> <p>if <i>uplo</i> = 'L', $ap(i + (j-1)(2n-j)/2) = A(i, j)$ for $j \leq i \leq n$.</p>
<i>x</i>	<p>REAL for <i>slatps</i></p> <p>DOUBLE PRECISION for <i>dlatps</i></p> <p>COMPLEX for <i>clatps</i></p> <p>COMPLEX*16 for <i>zlatps</i>.</p> <p>Array, DIMENSION (<i>n</i>)</p> <p>On entry, the right hand side <i>b</i> of the triangular system.</p>
<i>cnorm</i>	<p>REAL for <i>slatps/clatps</i></p> <p>DOUBLE PRECISION for <i>dlatps/zlatps</i>.</p> <p>Array, DIMENSION (<i>n</i>).</p> <p>If <i>normin</i> = 'Y', <i>cnorm</i> is an input argument and <i>cnorm</i>(<i>j</i>) contains the norm of the off-diagonal part of the <i>j</i>-th column of <i>A</i>. If <i>trans</i> = 'N', <i>cnorm</i>(<i>j</i>) must be greater than or equal to the infinity-norm, and if <i>trans</i> = 'T' or 'C', <i>cnorm</i>(<i>j</i>) must be greater than or equal to the 1-norm.</p>

Output Parameters

<code>x</code>	On exit, <code>x</code> is overwritten by the solution vector <code>x</code> .
<code>scale</code>	<code>REAL</code> for <code>slatps/clatps</code> <code>DOUBLE PRECISION</code> for <code>dlatps/zlatps</code> . The scaling factor <code>s</code> for the triangular system as described above. If <code>scale = 0</code> , the matrix <code>A</code> is singular or badly scaled, and the vector <code>x</code> is an exact or approximate solution to $Ax = 0$.
<code>cnorm</code>	If <code>normin = 'N'</code> , <code>cnorm</code> is an output argument and <code>cnorm(j)</code> returns the 1-norm of the off-diagonal part of the <code>j</code> -th column of <code>A</code> .
<code>info</code>	<code>INTEGER</code> . = 0: successful exit < 0: if <code>info = -k</code> , the <code>k</code> -th argument had an illegal value

?latrd

Reduces the first `nb` rows and columns of a symmetric/Hermitian matrix `A` to real tridiagonal form by an orthogonal/unitary similarity transformation.

```
call slatrd ( uplo, n, nb, a, lda, e, tau, w, ldw )
call dlatrd ( uplo, n, nb, a, lda, e, tau, w, ldw )
call clatrd ( uplo, n, nb, a, lda, e, tau, w, ldw )
call zlatrd ( uplo, n, nb, a, lda, e, tau, w, ldw )
```

Discussion

The routine `?latrd` reduces `nb` rows and columns of a real symmetric or complex Hermitian matrix `A` to symmetric/Hermitian tridiagonal form by an orthogonal/unitary similarity transformation $Q' A Q$, and returns the

matrices V and W which are needed to apply the transformation to the unreduced part of A .

If $uplo = 'U'$, `?latrd` reduces the last nb rows and columns of a matrix, of which the upper triangle is supplied;

if $uplo = 'L'$, `?latrd` reduces the first nb rows and columns of a matrix, of which the lower triangle is supplied.

This is an auxiliary routine called by `?sytrd/?hetrd`.

Input Parameters

<code>uplo</code>	<p>CHARACTER</p> <p>Specifies whether the upper or lower triangular part of the symmetric/Hermitian matrix A is stored:</p> <p>= 'U': Upper triangular</p> <p>= 'L': Lower triangular</p>
<code>n</code>	<p>INTEGER.</p> <p>The order of the matrix A.</p>
<code>nb</code>	<p>INTEGER.</p> <p>The number of rows and columns to be reduced.</p>
<code>a</code>	<p>REAL for <code>slatrd</code></p> <p>DOUBLE PRECISION for <code>dlatrd</code></p> <p>COMPLEX for <code>clatrd</code></p> <p>COMPLEX*16 for <code>zlatrd</code>.</p> <p>Array, DIMENSION (lda, n).</p> <p>On entry, the symmetric/Hermitian matrix A</p> <p>If $uplo = 'U'$, the leading n-by-n upper triangular part of a contains the upper triangular part of the matrix A, and the strictly lower triangular part of a is not referenced.</p> <p>If $uplo = 'L'$, the leading n-by-n lower triangular part of a contains the lower triangular part of the matrix A, and the strictly upper triangular part of a is not referenced.</p>
<code>lda</code>	<p>INTEGER.</p> <p>The leading dimension of the array a. $lda \geq (1,n)$.</p>

ldw INTEGER.
 The leading dimension of the output array *w*.
ldw ≥ max(1,*n*).

Output Parameters

a On exit, if *uplo* = 'U', the last *nb* columns have been reduced to tridiagonal form, with the diagonal elements overwriting the diagonal elements of *a*; the elements above the diagonal with the array *tau*, represent the orthogonal/unitary matrix *Q* as a product of elementary reflectors;
 if *uplo* = 'L', the first *nb* columns have been reduced to tridiagonal form, with the diagonal elements overwriting the diagonal elements of *a*; the elements below the diagonal with the array *tau*, represent the orthogonal/unitary matrix *Q* as a product of elementary reflectors.

e REAL for *slatrd/clatrd*
 DOUBLE PRECISION for *dlatrd/zlatrd*.
 If *uplo* = 'U', *e*(*n-nb*:*n-1*) contains the superdiagonal elements of the last *nb* columns of the reduced matrix;
 if *uplo* = 'L', *e*(1:*nb*) contains the subdiagonal elements of the first *nb* columns of the reduced matrix.

tau REAL for *slatrd*
 DOUBLE PRECISION for *dlatrd*
 COMPLEX for *clatrd*
 COMPLEX*16 for *zlatrd*.
 Array, DIMENSION (*lda*, *n*).
 The scalar factors of the elementary reflectors, stored in *tau*(*n-nb*:*n-1*) if *uplo* = 'U', and in *tau*(1:*nb*) if *uplo* = 'L'.

w REAL for *slatrd*
 DOUBLE PRECISION for *dlatrd*
 COMPLEX for *clatrd*
 COMPLEX*16 for *zlatrd*.

Array, `DIMENSION (lda, n)`.

The n -by- nb matrix W required to update the unreduced part of A .

Application Notes

If `uplo = 'U'`, the matrix Q is represented as a product of elementary reflectors

$$Q = H(n) H(n-1) \dots H(n-nb+1)$$

Each $H(i)$ has the form

$$H(i) = I - \tau v v^*$$

where τ is a real/complex scalar, and v is a real/complex vector with $v(i:n) = 0$ and $v(i-1) = 1$; $v(1:i-1)$ is stored on exit in `a(1:i-1, i)`, and τ in `tau(i-1)`.

If `uplo = 'L'`, the matrix Q is represented as a product of elementary reflectors

$$Q = H(1) H(2) \dots H(nb)$$

Each $H(i)$ has the form

$$H(i) = I - \tau v v^*$$

where τ is a real/complex scalar, and v is a real/complex vector with $v(1:i) = 0$ and $v(i+1) = 1$; $v(i+1:n)$ is stored on exit in `a(i+1:n, i)`, and τ in `tau(i)`.

The elements of the vectors v together form the n -by- nb matrix V which is needed, with W , to apply the transformation to the unreduced part of the matrix, using a symmetric/Hermitian rank- $2k$ update of the form:

$$A := A - VW^* - WV^*$$

The contents of `a` on exit are illustrated by the following examples with $n = 5$ and $nb = 2$:

if `uplo = 'U'`:

$$\begin{bmatrix} a & a & a & v_4 & v_5 \\ & a & a & v_4 & v_5 \\ & & a & 1 & v_5 \\ & & & d & 1 \\ & & & & d \end{bmatrix}$$

if `uplo = 'L'`:

$$\begin{bmatrix} d \\ 1 & d \\ v_1 & 1 & a \\ v_1 & v_2 & a & a \\ v_1 & v_2 & a & a & a \end{bmatrix}$$

where d denotes a diagonal element of the reduced matrix, a denotes an element of the original matrix that is unchanged, and v_i denotes an element of the vector defining $H(i)$.

?latrs

Solves a triangular system of equations with the scale factor set to prevent overflow.

```
call slatrs ( uplo, trans, diag, normin, n, a, lda, x,
              scale, cnorm, info )
call dlatrs ( uplo, trans, diag, normin, n, a, lda, x,
              scale, cnorm, info )
call clatrs ( uplo, trans, diag, normin, n, a, lda, x,
              scale, cnorm, info )
call zlatrs ( uplo, trans, diag, normin, n, a, lda, x,
              scale, cnorm, info )
```

Discussion

The routine solves one of the triangular systems

$Ax = s b$ or $A^T x = s b$ or $A^H x = s b$ (for complex flavors)

with scaling to prevent overflow. Here A is an upper or lower triangular matrix, A^T denotes the transpose of A , A^H denotes the conjugate transpose of A , x and b are n -element vectors, and s is a scaling factor, usually less than or equal to 1, chosen so that the components of x will be less than the overflow threshold. If the unscaled problem will not cause overflow, the Level 2 BLAS routine `?trsv` is called. If the matrix A is singular ($A(j,j) = 0$ for some j), then s is set to 0 and a non-trivial solution to $Ax = 0$ is returned.

Input Parameters

<i>uplo</i>	<p>CHARACTER*1.</p> <p>Specifies whether the matrix <i>A</i> is upper or lower triangular.</p> <p>= 'U': Upper triangular</p> <p>= 'L': Lower triangular</p>
<i>trans</i>	<p>CHARACTER*1.</p> <p>Specifies the operation applied to <i>A</i>.</p> <p>= 'N': Solve $Ax = s b$ (no transpose)</p> <p>= 'T': Solve $A^T x = s b$ (transpose)</p> <p>= 'C': Solve $A^H x = s b$ (conjugate transpose)</p>
<i>diag</i>	<p>CHARACTER*1.</p> <p>Specifies whether or not the matrix <i>A</i> is unit triangular.</p> <p>= 'N': Non-unit triangular</p> <p>= 'U': Unit triangular</p>
<i>normin</i>	<p>CHARACTER*1.</p> <p>Specifies whether <i>cnorm</i> has been set or not.</p> <p>= 'Y': <i>cnorm</i> contains the column norms on entry;</p> <p>= 'N': <i>cnorm</i> is not set on entry. On exit, the norms will be computed and stored in <i>cnorm</i>.</p>
<i>n</i>	<p>INTEGER.</p> <p>The order of the matrix <i>A</i>. $n \geq 0$</p>
<i>a</i>	<p>REAL for <i>slatrs</i></p> <p>DOUBLE PRECISION for <i>dlatrs</i></p> <p>COMPLEX for <i>clatrs</i></p> <p>COMPLEX*16 for <i>zlatrs</i>.</p> <p>Array, DIMENSION (<i>lda</i>, <i>n</i>). Contains the triangular matrix <i>A</i>. If <i>uplo</i> = 'U', the leading <i>n</i>-by-<i>n</i> upper triangular part of the array <i>a</i> contains the upper triangular matrix, and the strictly lower triangular part of <i>a</i> is not referenced. If <i>uplo</i> = 'L', the leading <i>n</i>-by-<i>n</i> lower triangular part of the array <i>a</i> contains the lower triangular matrix, and the strictly upper triangular part of <i>a</i> is not referenced. If <i>diag</i> = 'U', the diagonal elements of <i>a</i> are also not referenced and are assumed to be 1.</p>

lda INTEGER.
The leading dimension of the array *a*. $lda \geq \max(1, n)$.

x REAL for *slatrs*
DOUBLE PRECISION for *dlatrs*
COMPLEX for *clatrs*
COMPLEX*16 for *zlatrs*.
Array, DIMENSION (*n*). On entry, the right hand side *b* of the triangular system.

cnorm REAL for *slatrs/clatrs*)
DOUBLE PRECISION for *dlatrs/zlatrs*.
Array, DIMENSION (*n*). If *normin* = 'Y', *cnorm* is an input argument and *cnorm* (*j*) contains the norm of the off-diagonal part of the *j*-th column of *A*. If *trans* = 'N', *cnorm* (*j*) must be greater than or equal to the infinity-norm, and if *trans* = 'T' or 'C', *cnorm*(*j*) must be greater than or equal to the 1-norm.

Output Parameters

x On exit, *x* is overwritten by the solution vector *x*.

scale REAL for *slatrs/clatrs*)
DOUBLE PRECISION for *dlatrs/zlatrs*.
Array, DIMENSION (*lda*, *n*). The scaling factor *s* for the triangular system as described above.
If *scale* = 0, the matrix *A* is singular or badly scaled, and the vector *x* is an exact or approximate solution to $Ax = 0$.

cnorm If *normin* = 'N', *cnorm* is an output argument and *cnorm*(*j*) returns the 1-norm of the off-diagonal part of the *j*-th column of *A*.

info INTEGER.
= 0: successful exit
< 0: if *info* = -*k*, the *k*-th argument had an illegal value

Application Notes

A rough bound on x is computed; if that is less than overflow, `?trsv` is called, otherwise, specific code is used which checks for possible overflow or divide-by-zero at every operation.

A columnwise scheme is used for solving $Ax = b$. The basic algorithm if A is lower triangular is

```

x[1:n] := b[1:n]
for j = 1, ..., n
  x(j) := x(j) / A(j,j)
  x[j+1:n] := x[j+1:n] - x(j)*A[j+1:n,j]
end

```

Define bounds on the components of x after j iterations of the loop:

$M(j)$ = bound on $x[1:j]$

$G(j)$ = bound on $x[j+1:n]$

Initially, let $M(0) = 0$ and $G(0) = \max\{x(i), i=1, \dots, n\}$.

Then for iteration $j+1$ we have

$M(j+1) \leq G(j) / |A(j+1, j+1)|$

$G(j+1) \leq G(j) + M(j+1) * |A[j+2:n, j+1]|$

$\leq G(j) (1 + \mathit{cnorm}(j+1) / |A(j+1, j+1)|)$,

where $\mathit{cnorm}(j+1)$ is greater than or equal to the infinity-norm of column $j+1$ of A , not counting the diagonal. Hence

$$G(j) \leq G(0) \prod_{1 \leq i \leq j} (1 + \mathit{cnorm}(i) / |A(i,i)|)$$

and

$$|x(j)| \leq (G(0) / |A(j,j)|) \prod_{1 \leq i \leq j} (1 + \mathit{cnorm}(i) / |A(i,i)|)$$

Since $|x(j)| \leq M(j)$, we use the Level 2 BLAS routine `?trsv` if the reciprocal of the largest $M(j)$, $j=1, \dots, n$, is larger than $\max(\mathit{underflow}, 1/\mathit{overflow})$.

The bound on $x(j)$ is also used to determine when a step in the columnwise method can be performed without fear of overflow. If the computed bound

is greater than a large constant, x is scaled to prevent overflow, but if the bound overflows, x is set to 0, $x(j)$ to 1, and scale to 0, and a non-trivial solution to $Ax = 0$ is found.

Similarly, a row-wise scheme is used to solve $A^T x = b$ or $A^H x = b$. The basic algorithm for A upper triangular is

```

for  $j = 1, \dots, n$ 
 $x(j) := ( b(j) - A[1:j-1, j]' x[1:j-1] ) / A(j, j)$ 
end

```

We simultaneously compute two bounds

```

 $G(j) = \text{bound on } ( b(i) - A[1:i-1, i]' x[1:i-1] ), \quad 1 \leq i \leq j$ 
 $M(j) = \text{bound on } x(i), \quad 1 \leq i \leq j$ 

```

The initial values are $G(0) = 0$, $M(0) = \max\{b(i), i=1, \dots, n\}$, and we add the constraint $G(j) \geq G(j-1)$ and $M(j) \geq M(j-1)$ for $j \geq 1$.

Then the bound on $x(j)$ is

$$M(j) \leq M(j-1) * (1 + cnorm(j)) / |A(j, j)|$$

$$\leq M(0) \prod_{1 \leq i \leq j} (1 + cnorm(i) / |A(i, i)|)$$

and we can safely call `?trsv` if $1/M(n)$ and $1/G(n)$ are both greater than $\max(\text{underflow}, 1/\text{overflow})$.

?latrz

Factors an upper trapezoidal matrix by means of orthogonal/unitary transformations.

```

call slatz ( m, n, l, a, lda, tau, work )
call dlatrz ( m, n, l, a, lda, tau, work )
call clatz ( m, n, l, a, lda, tau, work )
call zlatrz ( m, n, l, a, lda, tau, work )

```

Discussion

The routine `?latrz` factors the m -by- $(m+1)$ real/complex upper trapezoidal matrix

$$[A1 \ A2] = [A(1:m, 1:m) \ A(1:m, n-l+1:n)]$$

as $(R \ 0) * Z$, by means of orthogonal/unitary transformations. Z is an $(m+1)$ -by- $(m+1)$ orthogonal/unitary matrix and R and $A1$ are m -by- m upper triangular matrices.

Input Parameters

<code>m</code>	INTEGER. The number of rows of the matrix A . $m \geq 0$.
<code>n</code>	INTEGER. The number of columns of the matrix A . $n \geq 0$.
<code>l</code>	INTEGER. The number of columns of the matrix A containing the meaningful part of the Householder vectors. $n-m \geq l \geq 0$.
<code>a</code>	REAL for <code>slatz</code> DOUBLE PRECISION for <code>dlatz</code> COMPLEX for <code>clatz</code> COMPLEX*16 for <code>zlatrz</code> . Array, DIMENSION (lda, n). On entry, the leading m -by- n upper trapezoidal part of the array <code>a</code> must contain the matrix to be factorized.
<code>lda</code>	INTEGER. The leading dimension of the array <code>a</code> . $lda \geq \max(1, m)$.
<code>work</code>	REAL for <code>slatz</code> DOUBLE PRECISION for <code>dlatz</code> COMPLEX for <code>clatz</code> COMPLEX*16 for <code>zlatrz</code> . Workspace array, DIMENSION (m).

Output Parameters

- a** On exit, the leading m -by- m upper triangular part of **a** contains the upper triangular matrix R , and elements $n-l+1$ to n of the first m rows of **a**, with the array **tau**, represent the orthogonal/unitary matrix Z as a product of m elementary reflectors.
- tau** **REAL** for **slatz**
DOUBLE PRECISION for **dlatrz**
COMPLEX for **clatz**
COMPLEX*16 for **zlatrz**.
 Array, **DIMENSION** (m). The scalar factors of the elementary reflectors.

Application Notes

The factorization is obtained by Householder's method. The k -th transformation matrix, $Z(k)$, which is used to introduce zeros into the $(m - k + 1)$ -th row of A , is given in the form

$$Z(k) = \begin{bmatrix} I & 0 \\ 0 & T(k) \end{bmatrix}$$

where

$$T(k) = I - \tau u(k) u(k)', \quad u(k) = \begin{bmatrix} 1 \\ 0 \\ z(k) \end{bmatrix}$$

τ is a scalar and $z(k)$ is an l -element vector. τ and $z(k)$ are chosen to annihilate the elements of the k -th row of A_2 .

The scalar τ is returned in the k -th element of **tau** and the vector $u(k)$ in the k -th row of A_2 , such that the elements of $z(k)$ are in $a(k, l+1)$, ..., $a(k, n)$. The elements of R are returned in the upper triangular part of A_1 .

Z is given by

$$Z = Z(1) Z(2) \dots Z(m).$$

?lauu2

Computes the product UU^H or L^HL , where U and L are upper or lower triangular matrices (unblocked algorithm).

```
call slauu2 ( uplo, n, a, lda, info )
call dlauu2 ( uplo, n, a, lda, info )
call clauu2 ( uplo, n, a, lda, info )
call zlauu2 ( uplo, n, a, lda, info )
```

Discussion

The routine `?lauu2` computes the product UU' or $L'L$, where the triangular factor U or L is stored in the upper or lower triangular part of the array a . If `uplo = 'U'` or `'u'`, then the upper triangle of the result is stored, overwriting the factor U in a . If `uplo = 'L'` or `'l'`, then the lower triangle of the result is stored, overwriting the factor L in a .

This is the unblocked form of the algorithm, calling Level 2 BLAS.

Input Parameters

<code>uplo</code>	CHARACTER*1. Specifies whether the triangular factor stored in the array a is upper or lower triangular: = 'U': Upper triangular = 'L': Lower triangular
<code>n</code>	INTEGER. The order of the triangular factor U or L . $n \geq 0$.
<code>a</code>	REAL for <code>slauu2</code> DOUBLE PRECISION for <code>dlauu2</code> COMPLEX for <code>clauu2</code>

COMPLEX*16 for `zlaau2`.

Array, DIMENSION (`lda`, `n`). On entry, the triangular factor U or L .

`lda` INTEGER.
The leading dimension of the array `a`. $lda \geq \max(1, n)$.

Output Parameters

`a` On exit, if `uplo` = 'U', the upper triangle of `a` is overwritten with the upper triangle of the product UU' ; if `uplo` = 'L', the lower triangle of `a` is overwritten with the lower triangle of the product LL .

`info` INTEGER.
= 0: successful exit
< 0: if `info` = $-k$, the k -th argument had an illegal value

?lauum

Computes the product UU^H or $L^H L$,
where U and L are upper or lower
triangular matrices.

```
call slauum ( uplo, n, a, lda, info )
call dlauum ( uplo, n, a, lda, info )
call clauum ( uplo, n, a, lda, info )
call zlauum ( uplo, n, a, lda, info )
```

Discussion

The routine `?lauum` computes the product UU' or LL , where the triangular factor U or L is stored in the upper or lower triangular part of the array `a`.

If `uplo` = 'U' or 'u', then the upper triangle of the result is stored, overwriting the factor U in `a`.

If `uplo` = 'L' or 'l', then the lower triangle of the result is stored, overwriting the factor L in `a`.

This is the blocked form of the algorithm, calling Level 3 BLAS.

Input Parameters

<i>uplo</i>	CHARACTER*1. Specifies whether the triangular factor stored in the array <i>a</i> is upper or lower triangular: = 'U': Upper triangular = 'L': Lower triangular
<i>n</i>	INTEGER. The order of the triangular factor <i>U</i> or <i>L</i> . $n \geq 0$.
<i>a</i>	REAL for <code>slauum</code> DOUBLE PRECISION for <code>dlauum</code> COMPLEX for <code>clauum</code> COMPLEX*16 for <code>zlauum</code> . Array, DIMENSION (<i>lda</i> , <i>n</i>). On entry, the triangular factor <i>U</i> or <i>L</i> .
<i>lda</i>	INTEGER. The leading dimension of the array <i>a</i> . $lda \geq \max(1, n)$.

Output Parameters

<i>a</i>	On exit, if <i>uplo</i> = 'U', the upper triangle of <i>a</i> is overwritten with the upper triangle of the product <i>UU</i> ; if <i>uplo</i> = 'L', the lower triangle of <i>a</i> is overwritten with the lower triangle of the product <i>LL</i> .
<i>info</i>	INTEGER. = 0: successful exit < 0: if <i>info</i> = - <i>k</i> , the <i>k</i> -th argument had an illegal value

?org2l/?ung2l

Generates all or part of the orthogonal/unitary matrix Q from a QL factorization determined by ?geqlf (unblocked algorithm).

```
call sorg2l ( m, n, k, a, lda, tau, work, info )
call dorg2l ( m, n, k, a, lda, tau, work, info )
call cung2l ( m, n, k, a, lda, tau, work, info )
call zung2l ( m, n, k, a, lda, tau, work, info )
```

Discussion

The routine ?org2l/?ung2l generates an m -by- n real/complex matrix Q with orthonormal columns, which is defined as the last n columns of a product of k elementary reflectors of order m :

$Q = H(k) \dots H(2) H(1)$ as returned by ?geqlf.

Input Parameters

m **INTEGER.**
The number of rows of the matrix Q . $m \geq 0$.

n **INTEGER.**
The number of columns of the matrix Q . $m \geq n \geq 0$.

k **INTEGER.**
The number of elementary reflectors whose product defines the matrix Q . $n \geq k \geq 0$.

a **REAL** for sorg2l
DOUBLE PRECISION for dorg2l
COMPLEX for cung2l
COMPLEX*16 for zung2l.
Array, **DIMENSION** (lda, n).
On entry, the $(n-k+i)$ -th column must contain the vector

which defines the elementary reflector $H(i)$, for $i = 1, 2, \dots, k$, as returned by `?geqlf` in the last k columns of its array argument a .

<code>lda</code>	INTEGER. The first dimension of the array a . $lda \geq \max(1, m)$.
<code>tau</code>	REAL for <code>sorg2l</code> DOUBLE PRECISION for <code>dorg2l</code> COMPLEX for <code>cung2l</code> COMPLEX*16 for <code>zung2l</code> . Array, DIMENSION (k). $\tau(i)$ must contain the scalar factor of the elementary reflector $H(i)$, as returned by <code>?geqlf</code> .
<code>work</code>	REAL for <code>sorg2l</code> DOUBLE PRECISION for <code>dorg2l</code> COMPLEX for <code>cung2l</code> COMPLEX*16 for <code>zung2l</code> . Workspace array, DIMENSION (n).

Output Parameters

<code>a</code>	On exit, the m -by- n matrix Q .
<code>info</code>	INTEGER. = 0: successful exit < 0: if <code>info</code> = $-i$, the i -th argument has an illegal value

?org2r/?ung2r

Generates all or part of the orthogonal/unitary matrix Q from a QR factorization determined by `?geqrf` (unblocked algorithm).

```
call sorg2r ( m, n, k, a, lda, tau, work, info )
call dorg2r ( m, n, k, a, lda, tau, work, info )
```

```
call cung2r ( m, n, k, a, lda, tau, work, info )
call zung2r ( m, n, k, a, lda, tau, work, info )
```

Discussion

The routine `?org2r/?ung2r` generates an m -by- n real/complex matrix Q with orthonormal columns, which is defined as the first n columns of a product of k elementary reflectors of order m

$$Q = H(1)H(2) \dots H(k)$$

as returned by `?geqrf`.

Input Parameters

m **INTEGER.**
The number of rows of the matrix Q . $m \geq 0$.

n **INTEGER.**
The number of columns of the matrix Q . $m \geq n \geq 0$.

k **INTEGER.**
The number of elementary reflectors whose product defines the matrix Q . $n \geq k \geq 0$.

a **REAL** for `sorg2r`
DOUBLE PRECISION for `dorg2r`
COMPLEX for `cung2r`
COMPLEX*16 for `zung2r`.
Array, **DIMENSION** (lda, n).
On entry, the i -th column must contain the vector which defines the elementary reflector $H(i)$, for $i = 1, 2, \dots, k$, as returned by `?geqrf` in the first k columns of its array argument a .

lda **INTEGER.**
The first **DIMENSION** of the array a . $lda \geq \max(1, m)$.

tau **REAL** for `sorg2r`
DOUBLE PRECISION for `dorg2r`
COMPLEX for `cung2r`
COMPLEX*16 for `zung2r`.

Array, `DIMENSION (k)`.
`tau(i)` must contain the scalar factor of the elementary reflector $H(i)$, as returned by `?geqrf`.

`work` `REAL` for `sorg2r`
 `DOUBLE PRECISION` for `dorg2r`
 `COMPLEX` for `cung2r`
 `COMPLEX*16` for `zung2r`.
 Workspace array, `DIMENSION (n)`.

Output Parameters

`a` On exit, the m -by- n matrix Q .
`info` `INTEGER`.
 = 0: successful exit
 < 0: if `info = -i`, the i -th argument has an illegal value

?orgl2/?ungl2

Generates all or part of the orthogonal/unitary matrix Q from an LQ factorization determined by `?gelqf` (unblocked algorithm).

```
call sorgl2 ( m, n, k, a, lda, tau, work, info )
call dorgl2 ( m, n, k, a, lda, tau, work, info )
call cungl2 ( m, n, k, a, lda, tau, work, info )
call zungl2 ( m, n, k, a, lda, tau, work, info )
```

Discussion

The routine `?orgl2/?ungl2` generates a m -by- n real/complex matrix Q with orthonormal rows, which is defined as the first m rows of a product of k elementary reflectors of order n

$$Q = H(k) \dots H(2) H(1) \text{ or } Q = H(k)' \dots H(2)' H(1)'$$

as returned by `?gelqf`.

Input Parameters

<i>m</i>	INTEGER. The number of rows of the matrix Q . $m \geq 0$.
<i>n</i>	INTEGER. The number of columns of the matrix Q . $n \geq m$.
<i>k</i>	INTEGER. The number of elementary reflectors whose product defines the matrix Q . $m \geq k \geq 0$.
<i>a</i>	REAL for <code>sorgl2</code> DOUBLE PRECISION for <code>dorgl2</code> COMPLEX for <code>cungl2</code> COMPLEX*16 for <code>zungl2</code> . Array, DIMENSION (<i>lda</i> , <i>n</i>). On entry, the <i>i</i> -th row must contain the vector which defines the elementary reflector $H(i)$, for $i = 1, 2, \dots, k$, as returned by <code>?gelqf</code> in the first <i>k</i> rows of its array argument <i>a</i> .
<i>lda</i>	INTEGER. The first dimension of the array <i>a</i> . $lda \geq \max(1, m)$.
<i>tau</i>	REAL for <code>sorgl2</code> DOUBLE PRECISION for <code>dorgl2</code> COMPLEX for <code>cungl2</code> COMPLEX*16 for <code>zungl2</code> . Array, DIMENSION (<i>k</i>). <i>tau</i> (<i>i</i>) must contain the scalar factor of the elementary reflector $H(i)$, as returned by <code>?gelqf</code> .
<i>work</i>	REAL for <code>sorgl2</code> DOUBLE PRECISION for <code>dorgl2</code> COMPLEX for <code>cungl2</code> COMPLEX*16 for <code>zungl2</code> . Workspace array, DIMENSION (<i>m</i>).

Output Parameters

<i>a</i>	On exit, the <i>m</i> -by- <i>n</i> matrix Q .
----------	--

info **INTEGER**.
 = 0: successful exit
 < 0: if *info* = -*i*, the *i*-th argument has an illegal value.

?orgr2/?ungr2

Generates all or part of the orthogonal/unitary matrix *Q* from an *RQ* factorization determined by ?gerqf (unblocked algorithm).

```
call sorgr2 ( m, n, k, a, lda, tau, work, info )
call dorgr2 ( m, n, k, a, lda, tau, work, info )
call cungr2 ( m, n, k, a, lda, tau, work, info )
call zungr2 ( m, n, k, a, lda, tau, work, info )
```

Discussion

The routine ?orgr2/?ungr2 generates an *m*-by-*n* real matrix *Q* with orthonormal rows, which is defined as the last *m* rows of a product of *k* elementary reflectors of order *n*

$$Q = H(1)H(2)\dots H(k) \text{ or } Q = H(1)' H(2)' \dots H(k)'$$

as returned by ?gerqf.

Input Parameters

m **INTEGER**. The number of rows of the matrix *Q*. $m \geq 0$.

n **INTEGER**.
 The number of columns of the matrix *Q*. $n \geq m$.

k **INTEGER**.
 The number of elementary reflectors whose product defines the matrix *Q*. $m \geq k \geq 0$.

a **REAL** for sorgr2
DOUBLE PRECISION for dorgr2
COMPLEX for cungr2

COMPLEX*16 for `zungr2`.
 Array, DIMENSION (`lda`, `n`). On entry, the ($m-k+i$)-th row must contain the vector which defines the elementary reflector $H(i)$, for $i = 1, 2, \dots, k$, as returned by `?gerqf` in the last k rows of its array argument `a`.

`lda` INTEGER.
 The first dimension of the array `a`. $lda \geq \max(1, m)$.

`tau` REAL for `sorgr2`
 DOUBLE PRECISION for `dorgr2`
 COMPLEX for `cungr2`
 COMPLEX*16 for `zungr2`.
 Array, DIMENSION (`k`). `tau(i)` must contain the scalar factor of the elementary reflector $H(i)$, as returned by `?gerqf`.

`work` REAL for `sorgr2`
 DOUBLE PRECISION for `dorgr2`
 COMPLEX for `cungr2`
 COMPLEX*16 for `zungr2`.
 Workspace array, DIMENSION (`m`).

Output Parameters

`a` On exit, the m -by- n matrix Q .

`info` INTEGER.
 = 0: successful exit
 < 0: if `info` = $-i$, the i -th argument has an illegal value

?orm2l/?unm2l

Multiplies a general matrix by the orthogonal/unitary matrix from a QL factorization determined by ?geqlf (unblocked algorithm).

```
call sorm2l ( side, trans, m, n, k, a, lda, tau, c, ldc, work, info )
call dorm2l ( side, trans, m, n, k, a, lda, tau, c, ldc, work, info )
call cumm2l ( side, trans, m, n, k, a, lda, tau, c, ldc, work, info )
call zumm2l ( side, trans, m, n, k, a, lda, tau, c, ldc, work, info )
```

Discussion

The routine ?orm2l/?unm2l overwrites the general real/complex m -by- n matrix C with

Q^*C if $side = 'L'$ and $trans = 'N'$, or
 Q^*C if $side = 'L'$ and $trans = 'T'$ (for real flavors) or
 $trans = 'C'$ (for complex flavors), or
 $C*Q$ if $side = 'R'$ and $trans = 'N'$, or
 $C*Q'$ if $side = 'R'$ and $trans = 'T'$ (for real flavors) or
 $trans = 'C'$ (for complex flavors)

where Q is a real orthogonal or complex unitary matrix defined as the product of k elementary reflectors

$$Q = H(k) \dots H(2) H(1)$$

as returned by ?geqlf. Q is of order m if $side = 'L'$ and of order n if $side = 'R'$.

Input Parameters

$side$ CHARACTER*1.
 = 'L': apply Q or Q' from the left
 = 'R': apply Q or Q' from the right

trans CHARACTER*1.
 = 'N': apply Q (No transpose)
 = 'T': apply Q' (Transpose, for real flavors)
 = 'C': apply Q' (Conjugate transpose, for complex flavors)

m INTEGER.
 The number of rows of the matrix C . $m \geq 0$.

n INTEGER.
 The number of columns of the matrix C . $n \geq 0$.

k INTEGER.
 The number of elementary reflectors whose product defines the matrix Q .
 If *side* = 'L', $m \geq k \geq 0$;
 if *side* = 'R', $n \geq k \geq 0$.

a REAL for *sorm2l*
 DOUBLE PRECISION for *dorm2l*
 COMPLEX for *cunm2l*
 COMPLEX*16 for *zunm2l*.
 Array, DIMENSION (*lda*,*k*). The i -th column must contain the vector which defines the elementary reflector $H(i)$, for $i = 1, 2, \dots, k$, as returned by `?geqlf` in the last k columns of its array argument *a*. The array *a* is modified by the routine but restored on exit.

lda INTEGER.
 The leading dimension of the array *a*.
 If *side* = 'L', $lda \geq \max(1, m)$;
 if *side* = 'R', $lda \geq \max(1, n)$.

tau REAL for *sorm2l*
 DOUBLE PRECISION for *dorm2l*
 COMPLEX for *cunm2l*
 COMPLEX*16 for *zunm2l*.
 Array, DIMENSION (*k*). *tau*(i) must contain the scalar factor of the elementary reflector $H(i)$, as returned by `?geqlf`.

c REAL for `sorm2l`
DOUBLE PRECISION for `dorm2l`
COMPLEX for `cunm2l`
COMPLEX*16 for `zunm2l`.
Array, DIMENSION (*ldc*, *n*). On entry, the *m*-by-*n* matrix *C*.

ldc INTEGER.
The leading dimension of the array *C*. $ldc \geq \max(1, m)$.

work REAL for `sorm2l`
DOUBLE PRECISION for `dorm2l`
COMPLEX for `cunm2l`
COMPLEX*16 for `zunm2l`.
Workspace array, DIMENSION:
(*n*) if *side* = 'L',
(*m*) if *side* = 'R'.

Output Parameters

c On exit, *c* is overwritten by *QC* or *Q'C* or *CQ'* or *CQ*.

info INTEGER.
= 0: successful exit
< 0: if *info* = -*i*, the *i*-th argument had an illegal value

?orm2r/?unm2r

Multiplies a general matrix by the orthogonal/unitary matrix from a QR factorization determined by ?geqrf (unblocked algorithm).

```
call sorm2r ( side, trans, m, n, k, a, lda, tau, c, ldc, work, info )
call dorm2r ( side, trans, m, n, k, a, lda, tau, c, ldc, work, info )
call cunm2r ( side, trans, m, n, k, a, lda, tau, c, ldc, work, info )
call zunm2r ( side, trans, m, n, k, a, lda, tau, c, ldc, work, info )
```

Discussion

The routine `?orm2r/?unm2r` overwrites the general real/complex m -by- n matrix C with

Q^*C if `side = 'L'` and `trans = 'N'`, or
 Q^*C if `side = 'L'` and `trans = 'T'` (for real flavors) or
`trans = 'C'` (for complex flavors), or
 C^*Q if `side = 'R'` and `trans = 'N'`, or
 C^*Q if `side = 'R'` and `trans = 'T'` (for real flavors) or
`trans = 'C'` (for complex flavors)

where Q is a real orthogonal or complex unitary matrix defined as the product of k elementary reflectors

$$Q = H(1) H(2) \dots H(k)$$

as returned by `?geqrf`. Q is of order m if `side = 'L'` and of order n if `side = 'R'`.

Input Parameters

`side` CHARACTER*1.
= 'L': apply Q or Q' from the left
= 'R': apply Q or Q' from the right

`trans` CHARACTER*1.
= 'N': apply Q (No transpose)
= 'T': apply Q' (Transpose, for real flavors)
= 'C': apply Q' (Conjugate transpose, for complex flavors)

`m` INTEGER.
The number of rows of the matrix C . $m \geq 0$.

`n` INTEGER.
The number of columns of the matrix C . $n \geq 0$.

`k` INTEGER.
The number of elementary reflectors whose product defines the matrix Q .
If `side = 'L'`, $m \geq k \geq 0$;
if `side = 'R'`, $n \geq k \geq 0$.

<i>a</i>	<p>REAL for <code>sorm2r</code> DOUBLE PRECISION for <code>dorm2r</code> COMPLEX for <code>cunm2r</code> COMPLEX*16 for <code>zunm2r</code>.</p> <p>Array, DIMENSION (<i>lda</i>,<i>k</i>).The <i>i</i>-th column must contain the vector which defines the elementary reflector $H(i)$, for $i = 1, 2, \dots, k$, as returned by <code>?geqrf</code> in the first <i>k</i> columns of its array argument <i>a</i>. The array <i>a</i> is modified by the routine but restored on exit.</p>
<i>lda</i>	<p>INTEGER.</p> <p>The leading dimension of the array <i>a</i>. If <i>side</i> = 'L', $lda \geq \max(1, m)$; if <i>side</i> = 'R', $lda \geq \max(1, n)$.</p>
<i>tau</i>	<p>REAL for <code>sorm2r</code> DOUBLE PRECISION for <code>dorm2r</code> COMPLEX for <code>cunm2r</code> COMPLEX*16 for <code>zunm2r</code>.</p> <p>Array, DIMENSION (<i>k</i>). <i>tau</i>(<i>i</i>) must contain the scalar factor of the elementary reflector $H(i)$, as returned by <code>?geqrf</code>.</p>
<i>c</i>	<p>REAL for <code>sorm2r</code> DOUBLE PRECISION for <code>dorm2r</code> COMPLEX for <code>cunm2r</code> COMPLEX*16 for <code>zunm2r</code>.</p> <p>Array, DIMENSION (<i>ldc</i>, <i>n</i>). On entry, the <i>m</i>-by-<i>n</i> matrix <i>C</i>.</p>
<i>ldc</i>	<p>INTEGER.</p> <p>The leading dimension of the array <i>C</i>. $ldc \geq \max(1, m)$.</p>
<i>work</i>	<p>REAL for <code>sorm2r</code> DOUBLE PRECISION for <code>dorm2r</code> COMPLEX for <code>cunm2r</code> COMPLEX*16 for <code>zunm2r</code>.</p> <p>Workspace array, DIMENSION (<i>n</i>) if <i>side</i> = 'L', (<i>m</i>) if <i>side</i> = 'R'.</p>

Output Parameters

c On exit, *c* is overwritten by QC or $Q'C$ or CQ' or CQ .

info INTEGER.
 = 0: successful exit
 < 0: if *info* = -i, the *i*-th argument had an illegal value

?orml2/?unml2

Multiplies a general matrix by the orthogonal/unitary matrix from a LQ factorization determined by ?gelqf (unblocked algorithm).

```
call sorml2 ( side, trans, m, n, k, a, lda, tau, c, ldc, work, info )
call dorml2 ( side, trans, m, n, k, a, lda, tau, c, ldc, work, info )
call cunml2 ( side, trans, m, n, k, a, lda, tau, c, ldc, work, info )
call zunml2 ( side, trans, m, n, k, a, lda, tau, c, ldc, work, info )
```

Discussion

The routine ?orml2/?unml2 overwrites the general real/complex *m*-by-*n* matrix *C* with

$Q * C$ if *side* = 'L' and *trans* = 'N', or
 $Q' * C$ if *side* = 'L' and *trans* = 'T' (for real flavors) or
 trans = 'C' (for complex flavors), or
 $C * Q$ if *side* = 'R' and *trans* = 'N', or
 $C * Q'$ if *side* = 'R' and *trans* = 'T' (for real flavors) or
 trans = 'C' (for complex flavors)

where *Q* is a real orthogonal or complex unitary matrix defined as the product of *k* elementary reflectors

$$Q = H(k) \dots H(2) H(1) \text{ or } Q = H(k)' \dots H(2)' H(1)'$$

as returned by ?gelqf. *Q* is of order *m* if *side* = 'L' and of order *n* if *side* = 'R'.

Input Parameters

<i>side</i>	<p>CHARACTER*1.</p> <p>= 'L': apply Q or Q' from the left</p> <p>= 'R': apply Q or Q' from the right</p>
<i>trans</i>	<p>CHARACTER*1.</p> <p>= 'N': apply Q (No transpose)</p> <p>= 'T': apply Q' (Transpose, for real flavors)</p> <p>= 'C': apply Q' (Conjugate transpose, for complex flavors)</p>
<i>m</i>	<p>INTEGER.</p> <p>The number of rows of the matrix C. $m \geq 0$.</p>
<i>n</i>	<p>INTEGER.</p> <p>The number of columns of the matrix C. $n \geq 0$.</p>
<i>k</i>	<p>INTEGER.</p> <p>The number of elementary reflectors whose product defines the matrix Q.</p> <p>If <i>side</i> = 'L', $m \geq k \geq 0$;</p> <p>if <i>side</i> = 'R', $n \geq k \geq 0$.</p>
<i>a</i>	<p>REAL for <code>sorml2</code></p> <p>DOUBLE PRECISION for <code>dorml2</code></p> <p>COMPLEX for <code>cunml2</code></p> <p>COMPLEX*16 for <code>zunml2</code>.</p> <p>Array, DIMENSION</p> <p>(<i>lda</i>, <i>m</i>) if <i>side</i> = 'L',</p> <p>(<i>lda</i>, <i>n</i>) if <i>side</i> = 'R'</p> <p>The i-th row must contain the vector which defines the elementary reflector $H(i)$, for $i = 1, 2, \dots, k$, as returned by <code>?gelqf</code> in the first k rows of its array argument <i>a</i>. The array <i>a</i> is modified by the routine but restored on exit.</p>
<i>lda</i>	<p>INTEGER.</p> <p>The leading dimension of the array <i>a</i>. $lda \geq \max(1, k)$.</p>
<i>tau</i>	<p>REAL for <code>sorml2</code></p> <p>DOUBLE PRECISION for <code>dorml2</code></p> <p>COMPLEX for <code>cunml2</code></p> <p>COMPLEX*16 for <code>zunml2</code>.</p>

Array, **DIMENSION** (k).
 $\tau(i)$ must contain the scalar factor of the elementary reflector $H(i)$, as returned by `?gelqf`.

c **REAL** for `sorml2`
DOUBLE PRECISION for `dorml2`
COMPLEX for `cunml2`
COMPLEX*16 for `zunml2`.
Array, **DIMENSION** (ldc, n)
On entry, the m -by- n matrix C .

ldc **INTEGER**.
The leading dimension of the array c . $ldc \geq \max(1, m)$.

work **REAL** for `sorml2`
DOUBLE PRECISION for `dorml2`
COMPLEX for `cunml2`
COMPLEX*16 for `zunml2`.
Workspace array, **DIMENSION**
(n) if $side = 'L'$,
(m) if $side = 'R'$

Output Parameters

c On exit, c is overwritten by QC or $Q'C$ or CQ' or CQ .

info **INTEGER**.
= 0: successful exit
< 0: if $info = -i$, the i -th argument had an illegal value

?ormr2/?unmr2

Multiplies a general matrix by the orthogonal/unitary matrix from a RQ factorization determined by `?gerqf` (unblocked algorithm).

```
call sormr2 ( side, trans, m, n, k, a, lda, tau, c, ldc, work, info )
```

```
call dormr2 ( side, trans, m, n, k, a, lda, tau, c, ldc, work, info )
call cunmr2 ( side, trans, m, n, k, a, lda, tau, c, ldc, work, info )
call zunmr2 ( side, trans, m, n, k, a, lda, tau, c, ldc, work, info )
```

Discussion

The routine `?ormr2/?unmr2` overwrites the general real/complex m -by- n matrix C with

$Q * C$ if `side = 'L'` and `trans = 'N'`, or
 $Q' * C$ if `side = 'L'` and `trans = 'T'` (for real flavors) or
`trans = 'C'` (for complex flavors), or
 $C * Q$ if `side = 'R'` and `trans = 'N'`, or
 $C * Q'$ if `side = 'R'` and `trans = 'T'` (for real flavors) or
`trans = 'C'` (for complex flavors)

where Q is a real orthogonal or complex unitary matrix defined as the product of k elementary reflectors

$$Q = H(1) H(2) \dots H(k) \text{ or } Q = H(1)' H(2)' \dots H(k)'$$

as returned by `?gerqf`. Q is of order m if `side = 'L'` and of order n if `side = 'R'`.

Input Parameters

`side` CHARACTER*1.
= 'L': apply Q or Q' from the left
= 'R': apply Q or Q' from the right

`trans` CHARACTER*1.
= 'N': apply Q (No transpose)
= 'T': apply Q' (Transpose, for real flavors)
= 'C': apply Q' (Conjugate transpose, for complex flavors)

`m` INTEGER.
The number of rows of the matrix C . $m \geq 0$.

`n` INTEGER.
The number of columns of the matrix C . $n \geq 0$.

- k* **INTEGER.**
The number of elementary reflectors whose product defines the matrix Q .
If *side* = 'L', $m \geq k \geq 0$;
if *side* = 'R', $n \geq k \geq 0$.
- a* **REAL** for *sormr2*
DOUBLE PRECISION for *dormr2*
COMPLEX for *cunmr2*
COMPLEX*16 for *zunmr2*.
Array, **DIMENSION**
(*lda*, *m*) if *side* = 'L',
(*lda*, *n*) if *side* = 'R'
The *i*-th row must contain the vector which defines the elementary reflector $H(i)$, for $i = 1, 2, \dots, k$, as returned by *?gerqf* in the last *k* rows of its array argument *a*. The array *a* is modified by the routine but restored on exit.
- lda* **INTEGER.**
The leading dimension of the array *a*. $lda \geq \max(1, k)$.
- tau* **REAL** for *sormr2*
DOUBLE PRECISION for *dormr2*
COMPLEX for *cunmr2*
COMPLEX*16 for *zunmr2*.
Array, **DIMENSION** (*k*).
tau(*i*) must contain the scalar factor of the elementary reflector $H(i)$, as returned by *?gerqf*.
- c* **REAL** for *sormr2*
DOUBLE PRECISION for *dormr2*
COMPLEX for *cunmr2*
COMPLEX*16 for *zunmr2*.
Array, **DIMENSION** (*ldc*, *n*).
On entry, the *m*-by-*n* matrix *C*.
- ldc* **INTEGER.**
The leading dimension of the array *C*. $ldc \geq \max(1, m)$.

work REAL for `sormr2`
 DOUBLE PRECISION for `dormr2`
 COMPLEX for `cunmr2`
 COMPLEX*16 for `zunmr2`.
 Workspace array, DIMENSION
 (*n*) if *side* = 'L',
 (*m*) if *side* = 'R'

Output Parameters

c On exit, *c* is overwritten by *QC* or *Q'C* or *CQ'* or *CQ*.
info INTEGER.
 = 0: successful exit
 < 0: if *info* = -*i*, the *i*-th argument had an illegal value

?ormr3/?unmr3

Multiplies a general matrix by the orthogonal/unitary matrix from a RZ factorization determined by ?tzzrf (unblocked algorithm).

```
call sormr3 (side, trans, m, n, k, l, a, lda, tau, c, ldc, work, info)
call dormr3 (side, trans, m, n, k, l, a, lda, tau, c, ldc, work, info)
call cunmr3 (side, trans, m, n, k, l, a, lda, tau, c, ldc, work, info)
call zunmr3 (side, trans, m, n, k, l, a, lda, tau, c, ldc, work, info)
```

Discussion

The routine ?ormr3/?unmr3 overwrites the general real/complex *m*-by-*n* matrix *C* with

$Q * C$ if *side* = 'L' and *trans* = 'N', or
 $Q' * C$ if *side* = 'L' and *trans* = 'T' (for real flavors) or
trans = 'C' (for complex flavors), or

$C*Q$ if *side* = 'R' and *trans* = 'N', or
 $C*Q'$ if *side* = 'R' and *trans* = 'T' (for real flavors) or
trans = 'C' (for complex flavors)

where Q is a real orthogonal or complex unitary matrix defined as the product of k elementary reflectors

$$Q = H(1) H(2) \dots H(k)$$

as returned by `?tzzrf`. Q is of order m if *side* = 'L' and of order n if *side* = 'R'.

Input Parameters

side CHARACTER*1.
 = 'L': apply Q or Q' from the left
 = 'R': apply Q or Q' from the right

trans CHARACTER*1.
 = 'N': apply Q (No transpose)
 = 'T': apply Q' (Transpose, for real flavors)
 = 'C': apply Q' (Conjugate transpose, for complex flavors)

m INTEGER.
 The number of rows of the matrix C . $m \geq 0$.

n INTEGER.
 The number of columns of the matrix C . $n \geq 0$.

k INTEGER.
 The number of elementary reflectors whose product defines the matrix Q .
 If *side* = 'L', $m \geq k \geq 0$;
 if *side* = 'R', $n \geq k \geq 0$.

l INTEGER.
 The number of columns of the matrix A containing the meaningful part of the Householder reflectors.
 If *side* = 'L', $m \geq l \geq 0$,
 if *side* = 'R', $n \geq l \geq 0$.

<i>a</i>	<p>REAL for <code>sormr3</code> DOUBLE PRECISION for <code>dormr3</code> COMPLEX for <code>cunmr3</code> COMPLEX*16 for <code>zunmr3</code>.</p> <p>Array, DIMENSION (<i>lda</i>, <i>m</i>) if <i>side</i> = 'L', (<i>lda</i>, <i>n</i>) if <i>side</i> = 'R'</p> <p>The <i>i</i>-th row must contain the vector which defines the elementary reflector $H(i)$, for $i = 1, 2, \dots, k$, as returned by <code>?tzzrf</code> in the last <i>k</i> rows of its array argument <i>a</i>. The array <i>a</i> is modified by the routine but restored on exit.</p>
<i>lda</i>	<p>INTEGER.</p> <p>The leading dimension of the array <i>a</i>. $lda \geq \max(1, k)$.</p>
<i>tau</i>	<p>REAL for <code>sormr3</code> DOUBLE PRECISION for <code>dormr3</code> COMPLEX for <code>cunmr3</code> COMPLEX*16 for <code>zunmr3</code>.</p> <p>Array, DIMENSION (<i>k</i>).</p> <p><i>tau</i>(<i>i</i>) must contain the scalar factor of the elementary reflector $H(i)$, as returned by <code>?tzzrf</code>.</p>
<i>c</i>	<p>REAL for <code>sormr3</code> DOUBLE PRECISION for <code>dormr3</code> COMPLEX for <code>cunmr3</code> COMPLEX*16 for <code>zunmr3</code>.</p> <p>Array, DIMENSION (<i>ldc</i>, <i>n</i>).</p> <p>On entry, the <i>m</i>-by-<i>n</i> matrix <i>C</i>.</p>
<i>ldc</i>	<p>INTEGER.</p> <p>The leading dimension of the array <i>c</i>. $ldc \geq \max(1, m)$.</p>
<i>work</i>	<p>REAL for <code>sormr3</code> DOUBLE PRECISION for <code>dormr3</code> COMPLEX for <code>cunmr3</code> COMPLEX*16 for <code>zunmr3</code>.</p> <p>Workspace array, DIMENSION (<i>n</i>) if <i>side</i> = 'L', (<i>m</i>) if <i>side</i> = 'R'.</p>

Output Parameters

c On exit, *c* is overwritten by QC or $Q'C$ or CQ' or CQ .

info INTEGER.
 = 0: successful exit
 < 0: if *info* = -*i*, the *i*-th argument had an illegal value

?pbtf2

Computes the Cholesky factorization of a symmetric/ Hermitian positive definite band matrix (unblocked algorithm).

```
call spbtf2 ( uplo, n, kd, ab, ldab, info )
call dpbtf2 ( uplo, n, kd, ab, ldab, info )
call cpbtf2 ( uplo, n, kd, ab, ldab, info )
call zpbtf2 ( uplo, n, kd, ab, ldab, info )
```

Discussion

The routine computes the Cholesky factorization of a real symmetric or complex Hermitian positive definite band matrix A . The factorization has the form

$A = U' U$, if *uplo* = 'U', or

$A = L L'$, if *uplo* = 'L',

where U is an upper triangular matrix, U' is the transpose of U , and L is lower triangular.

This is the unblocked version of the algorithm, calling Level 2 BLAS.

Input Parameters

uplo CHARACTER*1.
 Specifies whether the upper or lower triangular part of the symmetric/Hermitian matrix A is stored:
 = 'U': Upper triangular
 = 'L': Lower triangular

<i>n</i>	<p>INTEGER.</p> <p>The order of the matrix <i>A</i>. $n \geq 0$.</p>
<i>kd</i>	<p>INTEGER.</p> <p>The number of super-diagonals of the matrix <i>A</i> if <i>uplo</i> = 'U', or the number of sub-diagonals if <i>uplo</i> = 'L'. $kd \geq 0$.</p>
<i>ab</i>	<p>REAL for <i>spbtf2</i> DOUBLE PRECISION for <i>dpbtf2</i> COMPLEX for <i>cpbtf2</i> COMPLEX*16 for <i>zpbtf2</i>.</p> <p>Array, DIMENSION (<i>ldab</i>, <i>n</i>).</p> <p>On entry, the upper or lower triangle of the symmetric/Hermitian band matrix <i>A</i>, stored in the first <i>kd</i>+1 rows of the array. The <i>j</i>-th column of <i>A</i> is stored in the <i>j</i>-th column of the array <i>ab</i> as follows: if <i>uplo</i> = 'U', $ab(kd+1+i-j, j) = A(i, j)$ for $\max(1, j-kd) \leq i \leq j$; if <i>uplo</i> = 'L', $ab(1+i-j, j) = A(i, j)$ for $j \leq i \leq \min(n, j+kd)$.</p>
<i>ldab</i>	<p>INTEGER.</p> <p>The leading dimension of the array <i>ab</i>. $ldab \geq kd+1$.</p>

Output Parameters

<i>ab</i>	<p>On exit, if <i>info</i> = 0, the triangular factor <i>U</i> or <i>L</i> from the Cholesky factorization $A = U' U$ or $A = L L'$ of the band matrix <i>A</i>, in the same storage format as <i>A</i>.</p>
<i>info</i>	<p>INTEGER.</p> <p>= 0: successful exit < 0: if <i>info</i> = -<i>k</i>, the <i>k</i>-th argument had an illegal value > 0: if <i>info</i> = <i>k</i>, the leading minor of order <i>k</i> is not positive definite, and the factorization could not be completed.</p>

?potf2

Computes the Cholesky factorization of a symmetric/Hermitian positive definite matrix (unblocked algorithm).

```
call spotf2 ( uplo, n, a, lda, info )
call dpotf2 ( uplo, n, a, lda, info )
call cpotf2 ( uplo, n, a, lda, info )
call zpotf2 ( uplo, n, a, lda, info )
```

Discussion

The routine ?potf2 computes the Cholesky factorization of a real symmetric or complex Hermitian positive definite matrix A . The factorization has the form

$A = U' U$, if `uplo = 'U'`, or

$A = L L'$, if `uplo = 'L'`,

where U is an upper triangular matrix and L is lower triangular.

This is the unblocked version of the algorithm, calling Level 2 BLAS.

Input Parameters

<code>uplo</code>	CHARACTER*1. Specifies whether the upper or lower triangular part of the symmetric/Hermitian matrix A is stored. = 'U': Upper triangular = 'L': Lower triangular
<code>n</code>	INTEGER. The order of the matrix A . $n \geq 0$.
<code>a</code>	REAL for <code>spotf2</code> DOUBLE PRECISION or dpotf2 COMPLEX for <code>cpotf2</code> COMPLEX*16 for <code>zpotf2</code> . Array, DIMENSION (<code>lda</code> , <code>n</code>). On entry, the symmetric/Hermitian matrix A .

If *uplo* = 'U', the leading *n*-by-*n* upper triangular part of *a* contains the upper triangular part of the matrix *A*, and the strictly lower triangular part of *a* is not referenced.

If *uplo* = 'L', the leading *n*-by-*n* lower triangular part of *a* contains the lower triangular part of the matrix *A*, and the strictly upper triangular part of *a* is not referenced.

lda INTEGER.

The leading dimension of the array *a*. $lda \geq \max(1, n)$.

Output Parameters

a On exit, if *info* = 0, the factor *U* or *L* from the Cholesky factorization $A = U' U$ or $A = L L'$.

info INTEGER.

= 0: successful exit

< 0: if *info* = -*k*, the *k*-th argument had an illegal value

> 0: if *info* = *k*, the leading minor of order *k* is not positive definite, and the factorization could not be completed.

?ptts2

Solves a tridiagonal system of the form $AX=B$ using the $L D L^H$ factorization computed by ?pttrf.

```
call sptts2 ( n, nrhs, d, e, b, ldb )
call dptts2 ( n, nrhs, d, e, b, ldb )
call cptts2 ( iuplo, n, nrhs, d, e, b, ldb )
call zptts2 ( iuplo, n, nrhs, d, e, b, ldb )
```

Discussion

The routine `?ptts2` solves a tridiagonal system of the form

$$A X = B$$

Real flavors `sptts2/dptts2` use the LDL' factorization of A computed by `sptrf/dptrf`, and complex flavors `cptts2/zptts2` use the $U'DU$ or LDL' factorization of A computed by `cpttrf/zpttrf`. D is a diagonal matrix specified in the vector d , U (or L) is a unit bidiagonal matrix whose superdiagonal (subdiagonal) is specified in the vector e , and X and B are n -by- $nrhs$ matrices.

Input Parameters

- `iuplo` **INTEGER**. Used with complex flavors only. Specifies the form of the factorization and whether the vector e is the superdiagonal of the upper bidiagonal factor U or the subdiagonal of the lower bidiagonal factor L .
 = 1: $A = U' D U$, e is the superdiagonal of U ;
 = 0: $A = L D L'$, e is the subdiagonal of L .
- `n` **INTEGER**.
 The order of the tridiagonal matrix A . $n \geq 0$.
- `nrhs` **INTEGER**.
 The number of right hand sides, that is, the number of columns of the matrix B . $nrhs \geq 0$.
- `d` **REAL** for `sptts2/cptts2`
DOUBLE PRECISION for `dptts2/zptts2`.
 Array, **DIMENSION** (n).
 The n diagonal elements of the diagonal matrix D from the factorization of A .
- `e` **REAL** for `sptts2`
DOUBLE PRECISION for `dptts2`
COMPLEX for `cptts2`
COMPLEX*16 for `zptts2`.
 Array, **DIMENSION** ($n-1$).
 Contains the ($n-1$) subdiagonal elements of the unit bidiagonal factor L from the LDL' factorization of A (for

real flavors, or for complex flavors when *iuplo* = 0). For complex flavors when *iuplo* = 1, *e* contains the (*n*-1) superdiagonal elements of the unit bidiagonal factor *U* from the factorization $A = U'DU$.

b REAL for *sptts2/cptts2*
DOUBLE PRECISION for *dptts2/zptts2*.
Array, DIMENSION (*ldb*, *nrhs*).
On entry, the right hand side vectors *B* for the system of linear equations.

ldb INTEGER.
The leading dimension of the array *B*. $ldb \geq \max(1, n)$.

Output Parameters

b On exit, the solution vectors, *X*.

?rscl

Multiplies a vector by the reciprocal of a real scalar.

```
call srscl ( n, sa, sx, incx )
call drscl ( n, sa, sx, incx )
call csrcsl ( n, sa, sx, incx )
call zdrscl ( n, sa, sx, incx )
```

Discussion

The routine *?rscl* multiplies an *n*-element real/complex vector *x* by the real scalar $1/a$. This is done without overflow or underflow as long as the final result x/a does not overflow or underflow.

Input Parameters

n INTEGER.
The number of components of the vector *x*.

sa REAL for `srscl/crscl`
DOUBLE PRECISION for `drscl/zdrscl`.
The scalar a which is used to divide each component of the vector x . **sa** must be ≥ 0 , or the subroutine will divide by zero.

sx REAL for `srscl`
DOUBLE PRECISION for `drscl`
COMPLEX for `crscl`
COMPLEX*16 for `zdrscl`.
Array, DIMENSION $(1+(n-1)*\text{abs}(\text{incx}))$.
The n -element vector x .

incx INTEGER.
The increment between successive values of the vector **sx**.
If $\text{incx} > 0$, $\text{sx}(1) = x(1)$ and
 $\text{sx}(1+(i-1)*\text{incx}) = x(i)$, $1 < i \leq n$.

Output Parameters

sx On exit, the result x/a .

?sygs2/?hegs2

Reduces a symmetric/Hermitian definite generalized eigenproblem to standard form, using the factorization results obtained from ?potrf (unblocked algorithm).

```
call ssygs2 ( itype, uplo, n, a, lda, b, ldb, info )
call dsygs2 ( itype, uplo, n, a, lda, b, ldb, info )
call chgs2 ( itype, uplo, n, a, lda, b, ldb, info )
call zhegs2 ( itype, uplo, n, a, lda, b, ldb, info )
```

Discussion

The routine `?sygs2/?hegs2` reduces a real symmetric-definite or a complex Hermitian-definite generalized eigenproblem to standard form.

If `itype = 1`, the problem is

$$Ax = \lambda Bx,$$

and `A` is overwritten by `inv(U)*A*inv(U)` or `inv(L)*A*inv(L)`.

If `itype = 2` or `3`, the problem is

$$ABx = \lambda x \text{ or } B Ax = \lambda x,$$

and `A` is overwritten by `UAU'` or `L'AL`. `B` must have been previously factorized as `U'U` or `LL'` by `?potrf`.

Input Parameters

<code>itype</code>	<p>INTEGER.</p> <p>= 1: compute <code>inv(U)*A*inv(U)</code> or <code>inv(L)*A*inv(L)</code>; = 2 or 3: compute <code>UAU'</code> or <code>L'AL</code>.</p>
<code>uplo</code>	<p>CHARACTER</p> <p>Specifies whether the upper or lower triangular part of the symmetric/Hermitian matrix <code>A</code> is stored, and how <code>B</code> has been factorized.</p> <p>= 'U': Upper triangular = 'L': Lower triangular</p>
<code>n</code>	<p>INTEGER.</p> <p>The order of the matrices <code>A</code> and <code>B</code>. $n \geq 0$.</p>
<code>a</code>	<p>REAL for <code>ssygs2</code> DOUBLE PRECISION for <code>dsygs2</code> COMPLEX for <code>chegs2</code> COMPLEX*16 for <code>zhegs2</code>.</p> <p>Array, DIMENSION (<code>lda</code>, <code>n</code>).</p> <p>On entry, the symmetric/Hermitian matrix <code>A</code>.</p> <p>If <code>uplo = 'U'</code>, the leading <code>n</code>-by-<code>n</code> upper triangular part of <code>a</code> contains the upper triangular part of the matrix <code>A</code>, and the strictly lower triangular part of <code>a</code> is not referenced.</p> <p>If <code>uplo = 'L'</code>, the leading <code>n</code>-by-<code>n</code> lower triangular part of <code>a</code> contains the lower triangular part of the matrix <code>A</code>, and the strictly upper triangular part of <code>a</code> is not referenced.</p>

lda **INTEGER.**
The leading dimension of the array *a*. $lda \geq \max(1,n)$.

b **REAL** for *ssygs2*
DOUBLE PRECISION for *dsygs2*
COMPLEX for *chegs2*
COMPLEX*16 for *zhegs2*.
Array, **DIMENSION** (*ldb*, *n*).
The triangular factor from the Cholesky factorization of *B* as returned by *?potrf*.

ldb **INTEGER.**
The leading dimension of the array *B*. $ldb \geq \max(1,n)$.

Output Parameters

a On exit, if *info* = 0, the transformed matrix, stored in the same format as *A*.

info **INTEGER.**
= 0: successful exit.
< 0: if *info* = *-i*, the *i*-th argument had an illegal value.

?sytd2/?hetd2

Reduces a symmetric/Hermitian matrix to real symmetric tridiagonal form by an orthogonal/unitary similarity transformation (unblocked algorithm).

```
call ssytd2 ( uplo, n, a, lda, d, e, tau, info )
call dsytd2 ( uplo, n, a, lda, d, e, tau, info )
call chetd2 ( uplo, n, a, lda, d, e, tau, info )
call zhetd2 ( uplo, n, a, lda, d, e, tau, info )
```

Discussion

The routine `?sytd2/?hetd2` reduces a real symmetric/complex Hermitian matrix A to real symmetric tridiagonal form T by an orthogonal/unitary similarity transformation: $Q' A Q = T$.

Input Parameters

<code>uplo</code>	<p>CHARACTER*1.</p> <p>Specifies whether the upper or lower triangular part of the symmetric/Hermitian matrix A is stored:</p> <p>= 'U': Upper triangular</p> <p>= 'L': Lower triangular</p>
<code>n</code>	<p>INTEGER.</p> <p>The order of the matrix A. $n \geq 0$.</p>
<code>a</code>	<p>REAL for <code>ssytd2</code></p> <p>DOUBLE PRECISION for <code>dsytd2</code></p> <p>COMPLEX for <code>chetd2</code></p> <p>COMPLEX*16 for <code>zhetd2</code>.</p> <p>Array, DIMENSION (lda, n).</p> <p>On entry, the symmetric/Hermitian matrix A.</p> <p>If <code>uplo</code> = 'U', the leading n-by-n upper triangular part of a contains the upper triangular part of the matrix A, and the strictly lower triangular part of a is not referenced.</p> <p>If <code>uplo</code> = 'L', the leading n-by-n lower triangular part of a contains the lower triangular part of the matrix A, and the strictly upper triangular part of a is not referenced.</p>
<code>lda</code>	<p>INTEGER.</p> <p>The leading dimension of the array a. $lda \geq \max(1, n)$.</p>

Output Parameters

<code>a</code>	<p>On exit, if <code>uplo</code> = 'U', the diagonal and first superdiagonal of a are overwritten by the corresponding elements of the tridiagonal matrix T, and the elements above the first superdiagonal, with the array <code>tau</code>, represent the orthogonal/unitary matrix Q as a product of elementary reflectors;</p>
----------------	---

if *uplo* = 'L', the diagonal and first subdiagonal of *a* are overwritten by the corresponding elements of the tridiagonal matrix *T*, and the elements below the first subdiagonal, with the array *tau*, represent the orthogonal/unitary matrix *Q* as a product of elementary reflectors.

- d* REAL for *ssytd2/chetd2*
 DOUBLE PRECISION for *dsytd2/zhetd2*.
 Array, DIMENSION (*n*).
 The diagonal elements of the tridiagonal matrix *T*:
 $d(i) = a(i,i)$.
- e* REAL for *ssytd2/chetd2*
 DOUBLE PRECISION for *dsytd2/zhetd2*.
 Array, DIMENSION (*n*-1).
 The off-diagonal elements of the tridiagonal matrix *T*:
 $e(i) = a(i,i+1)$ if *uplo* = 'U',
 $e(i) = a(i+1,i)$ if *uplo* = 'L'.
- tau* REAL for *ssytd2*
 DOUBLE PRECISION for *dsytd2*
 COMPLEX for *chetd2*
 COMPLEX*16 for *zhetd2*.
 Array, DIMENSION (*n*-1).
 The scalar factors of the elementary reflectors .
- info* INTEGER.
 = 0: successful exit
 < 0: if *info* = -*i*, the *i*-th argument had an illegal value.

?sytf2

Computes the factorization of a real/complex symmetric indefinite matrix, using the diagonal pivoting method (unblocked algorithm).

```
call ssytf2 ( uplo, n, a, lda, ipiv, info )
call dsytf2 ( uplo, n, a, lda, ipiv, info )
call csytf2 ( uplo, n, a, lda, ipiv, info )
call zsytf2 ( uplo, n, a, lda, ipiv, info )
```

Discussion

The routine `?sytf2` computes the factorization of a real/complex symmetric matrix A using the Bunch-Kaufman diagonal pivoting method:

$$A = U D U' \text{ or } A = L D L'$$

where U (or L) is a product of permutation and unit upper (lower) triangular matrices, U' is the transpose of U , and D is symmetric and block diagonal with 1-by-1 and 2-by-2 diagonal blocks.

This is the unblocked version of the algorithm, calling Level 2 BLAS.

Input Parameters

<code>uplo</code>	<code>CHARACTER*1</code> . Specifies whether the upper or lower triangular part of the symmetric matrix A is stored = 'U': Upper triangular = 'L': Lower triangular
<code>n</code>	<code>INTEGER</code> . The order of the matrix A . $n \geq 0$.
<code>a</code>	<code>REAL</code> for <code>ssytf2</code> <code>DOUBLE PRECISION</code> for <code>dsytf2</code> <code>COMPLEX</code> for <code>csytf2</code> <code>COMPLEX*16</code> for <code>zsytf2</code> . Array, <code>DIMENSION (lda, n)</code> .

On entry, the symmetric matrix A .

If $uplo = 'U'$, the leading n -by- n upper triangular part of a contains the upper triangular part of the matrix A , and the strictly lower triangular part of a is not referenced.

If $uplo = 'L'$, the leading n -by- n lower triangular part of a contains the lower triangular part of the matrix A , and the strictly upper triangular part of a is not referenced.

lda **INTEGER.**

The leading dimension of the array a . $lda \geq \max(1,n)$.

Output Parameters

a On exit, the block diagonal matrix D and the multipliers used to obtain the factor U or L .

$ipiv$ **INTEGER.**
Array, **DIMENSION** (n).

Details of the interchanges and the block structure of D
If $ipiv(k) > 0$, then rows and columns k and $ipiv(k)$ were interchanged and $D(k,k)$ is a 1-by-1 diagonal block.

If $uplo = 'U'$ and $ipiv(k) = ipiv(k-1) < 0$, then rows and columns $k-1$ and $-ipiv(k)$ were interchanged and $D(k-1:k, k-1:k)$ is a 2-by-2 diagonal block.

If $uplo = 'L'$ and $ipiv(k) = ipiv(k+1) < 0$, then rows and columns $k+1$ and $-ipiv(k)$ were interchanged and $D(k:k+1, k:k+1)$ is a 2-by-2 diagonal block.

$info$ **INTEGER.**

= 0: successful exit

< 0: if $info = -k$, the k -th argument had an illegal value

> 0: if $info = k$, $D(k,k)$ is exactly zero. The factorization has been completed, but the block diagonal matrix D is exactly singular, and division by zero will occur if it is used to solve a system of equations.

?hetf2

Computes the factorization of a complex Hermitian matrix, using the diagonal pivoting method (unblocked algorithm).

```
call chetf2 ( uplo, n, a, lda, ipiv, info )
call zhetf2 ( uplo, n, a, lda, ipiv, info )
```

Discussion

The routine computes the factorization of a complex Hermitian matrix A using the Bunch-Kaufman diagonal pivoting method:

$$A = U D U' \text{ or } A = L D L'$$

where U (or L) is a product of permutation and unit upper (lower) triangular matrices, U' is the conjugate transpose of U , and D is Hermitian and block diagonal with 1-by-1 and 2-by-2 diagonal blocks.

This is the unblocked version of the algorithm, calling Level 2 BLAS.

Input Parameters

<i>uplo</i>	CHARACTER*1. Specifies whether the upper or lower triangular part of the Hermitian matrix A is stored: = 'U': Upper triangular = 'L': Lower triangular
<i>n</i>	INTEGER. The order of the matrix A . $n \geq 0$.
<i>a</i>	COMPLEX for chetf2 COMPLEX*16 for zhetf2 . Array, DIMENSION (<i>lda</i> , <i>n</i>). On entry, the Hermitian matrix A . If <i>uplo</i> = 'U', the leading n -by- n upper triangular part of <i>a</i> contains the upper triangular part of the matrix A , and the strictly lower triangular part of <i>a</i> is not referenced.

If *uplo* = 'L', the leading *n*-by-*n* lower triangular part of *a* contains the lower triangular part of the matrix *A*, and the strictly upper triangular part of *a* is not referenced.

lda **INTEGER.**
The leading dimension of the array *a*. $lda \geq \max(1,n)$.

Output Parameters

a On exit, the block diagonal matrix *D* and the multipliers used to obtain the factor *U* or *L*.

ipiv **INTEGER.**
Array, **DIMENSION** (*n*).
Details of the interchanges and the block structure of *D*
If *ipiv*(*k*) > 0, then rows and columns *k* and *ipiv*(*k*) were interchanged and *D*(*k*,*k*) is a 1-by-1 diagonal block.
If *uplo* = 'U' and *ipiv*(*k*) = *ipiv*(*k*-1) < 0, then rows and columns *k*-1 and -*ipiv*(*k*) were interchanged and *D*(*k*-1:*k*,*k*-1:*k*) is a 2-by-2 diagonal block.
If *uplo* = 'L' and *ipiv*(*k*) = *ipiv*(*k*+1) < 0, then rows and columns *k*+1 and -*ipiv*(*k*) were interchanged and *D*(*k*:*k*+1,*k*:*k*+1) is a 2-by-2 diagonal block.

info **INTEGER.**
= 0: successful exit
< 0: if *info* = -*k*, the *k*-th argument had an illegal value
> 0: if *info* = *k*, *D*(*k*,*k*) is exactly zero. The factorization has been completed, but the block diagonal matrix *D* is exactly singular, and division by zero will occur if it is used to solve a system of equations.

?tgex2

Swaps adjacent diagonal blocks in an upper (quasi) triangular matrix pair by an orthogonal/unitary equivalence transformation.

```
call stgex2 ( wantq, wantz, n, a, lda, b, ldb, q, ldq, z,
             ldz, j1, n1, n2, work, lwork, info )
call dtgex2 ( wantq, wantz, n, a, lda, b, ldb, q, ldq, z,
             ldz, j1, n1, n2, work, lwork, info )
call ctgex2 ( wantq, wantz, n, a, lda, b, ldb, q, ldq, z,
             ldz, j1, info )
call ztgex2 ( wantq, wantz, n, a, lda, b, ldb, q, ldq, z,
             ldz, j1, info )
```

Discussion

The real routines `stgex2/dtgex2` swap adjacent diagonal blocks (A_{11} , B_{11}) and (A_{22} , B_{22}) of size 1-by-1 or 2-by-2 in an upper (quasi) triangular matrix pair (A, B) by an orthogonal equivalence transformation. (A, B) must be in generalized real Schur canonical form (as returned by `sgges/dgges`), that is, A is block upper triangular with 1-by-1 and 2-by-2 diagonal blocks. B is upper triangular.

The complex routines `ctgex2/ztgex2` swap adjacent diagonal 1-by-1 blocks (A_{11} , B_{11}) and (A_{22} , B_{22}) in an upper triangular matrix pair (A, B) by an unitary equivalence transformation. (A, B) must be in generalized Schur canonical form, that is, A and B are both upper triangular.

All routines optionally update the matrices Q and Z of generalized Schur vectors:

$$Q(\text{in}) * A(\text{in}) * Z(\text{in})' = Q(\text{out}) * A(\text{out}) * Z(\text{out})'$$

$$Q(\text{in}) * B(\text{in}) * Z(\text{in})' = Q(\text{out}) * B(\text{out}) * Z(\text{out})'$$

Input Parameters

<i>wantq</i>	LOGICAL. If <i>wantq</i> = <code>.TRUE.</code> : update the left transformation matrix <i>Q</i> ; If <i>wantq</i> = <code>.FALSE.</code> : do not update <i>Q</i> .
<i>wantz</i>	LOGICAL. If <i>wantz</i> = <code>.TRUE.</code> : update the right transformation matrix <i>Z</i> ; If <i>wantz</i> = <code>.FALSE.</code> : do not update <i>Z</i> .
<i>n</i>	INTEGER. The order of the matrices <i>A</i> and <i>B</i> . $n \geq 0$.
<i>a, b</i>	REAL for <code>stgex2</code> DOUBLE PRECISION for <code>dtgex2</code> COMPLEX for <code>ctgex2</code> COMPLEX*16 for <code>ztgex2</code> . Arrays, DIMENSION (<i>lda</i> , <i>n</i>) and (<i>ldb</i> , <i>n</i>), respectively. On entry, the matrices <i>A</i> and <i>B</i> in the pair (<i>A</i> , <i>B</i>).
<i>lda</i>	INTEGER. The leading dimension of the array <i>a</i> . $lda \geq \max(1, n)$.
<i>ldb</i>	INTEGER. The leading dimension of the array <i>b</i> . $ldb \geq \max(1, n)$.
<i>q, z</i>	REAL for <code>stgex2</code> DOUBLE PRECISION for <code>dtgex2</code> COMPLEX for <code>ctgex2</code> COMPLEX*16 for <code>ztgex2</code> . Arrays, DIMENSION (<i>ldq</i> , <i>n</i>) and (<i>ldz</i> , <i>n</i>), respectively. On entry, if <i>wantq</i> = <code>.TRUE.</code> , <i>q</i> contains the orthogonal/unitary matrix <i>Q</i> , and if <i>wantz</i> = <code>.TRUE.</code> , <i>z</i> contains the orthogonal/unitary matrix <i>Z</i> .
<i>ldq</i>	INTEGER. The leading dimension of the array <i>q</i> . $ldq \geq 1$. If <i>wantq</i> = <code>.TRUE.</code> , $ldq \geq n$.

<i>ldz</i>	<p>INTEGER.</p> <p>The leading dimension of the array <i>z</i>. $ldz \geq 1$. If <i>wantz</i> = <i>.TRUE.</i>, $ldz \geq n$.</p>
<i>j1</i>	<p>INTEGER.</p> <p>The index to the first block (A11, B11). $1 \leq j1 \leq n$.</p>
<i>n1</i>	<p>INTEGER. Used with real flavors only.</p> <p>The order of the first block (A11, B11). $n1 = 0, 1$ or 2.</p>
<i>n2</i>	<p>INTEGER. Used with real flavors only.</p> <p>The order of the second block (A22, B22). $n2 = 0, 1$ or 2.</p>
<i>work</i>	<p>REAL for <i>stgex2</i> DOUBLE PRECISION for <i>dtgex2</i>.</p> <p>Workspace array, DIMENSION (<i>lwork</i>). Used with real flavors only.</p>
<i>lwork</i>	<p>INTEGER.</p> <p>The dimension of the array <i>work</i>. $lwork \geq \max(n*(n2+n1), 2*(n2+n1)^2)$</p>

Output Parameters

<i>a</i>	On exit, the updated matrix <i>A</i> .
<i>b</i>	On exit, the updated matrix <i>B</i> .
<i>q</i>	On exit, the updated matrix <i>Q</i> . Not referenced if <i>wantq</i> = <i>.FALSE.</i>
<i>z</i>	On exit, the updated matrix <i>Z</i> . Not referenced if <i>wantz</i> = <i>.FALSE.</i>
<i>info</i>	<p>INTEGER.</p> <p>=0: Successful exit</p> <p>For <i>stgex2/dtgex2</i>: if <i>info</i> = 1, the transformed matrix (<i>A</i>, <i>B</i>) would be too far from generalized Schur form; the blocks are not swapped and (<i>A</i>, <i>B</i>) and (<i>Q</i>, <i>Z</i>) are unchanged. The problem of swapping is too ill-conditioned. If <i>info</i> = -16: <i>lwork</i> is too small. Appropriate value for <i>lwork</i> is returned in <i>work</i>(1).</p>

For `ctgex2/ztgex2`: if `info = 1`, the transformed matrix pair (A, B) would be too far from generalized Schur form; the problem is ill-conditioned. (A, B) may have been partially reordered, and `ilst` points to the first row of the current position of the block being moved.

?tgsy2

Solves the generalized Sylvester equation (unblocked algorithm).

```
call stgsy2 ( trans, ijob, m, n, a, lda, b, ldb, c, ldc,
             d, ldd, e, lde, f, ldf, scale, rdsum, rdscal,
             iwork, pq, info )
call dtgsy2 ( trans, ijob, m, n, a, lda, b, ldb, c, ldc,
             d, ldd, e, lde, f, ldf, scale, rdsum, rdscal,
             iwork, pq, info )
call ctgsy2 ( trans, ijob, m, n, a, lda, b, ldb, c, ldc,
             d, ldd, e, lde, f, ldf, scale, rdsum, rdscal,
             info )
call ztgsy2 ( trans, ijob, m, n, a, lda, b, ldb, c, ldc,
             d, ldd, e, lde, f, ldf, scale, rdsum, rdscal,
             info )
```

Discussion

The routine `?tgsy2` solves the generalized Sylvester equation:

$$\begin{aligned} AR - LB &= \text{scale} * C \\ DR - LE &= \text{scale} * F, \end{aligned} \quad (1)$$

using Level 1 and 2 BLAS, where R and L are unknown m -by- n matrices, (A, D) , (B, E) and (C, F) are given matrix pairs of size m -by- m , n -by- n and m -by- n , respectively.

For `stgsy2/dtgsy2`, pairs (A, D) and (B, E) must be in generalized Schur

canonical form, that is, A, B are upper quasi triangular and D, E are upper triangular. For `ctgsy2/ztgsy2`, matrices A, B, D and E are upper triangular (that is, (A, D) and (B, E) in generalized Schur form).

The solution (R, L) overwrites (C, F) . $0 \leq \text{scale} \leq 1$ is an output scaling factor chosen to avoid overflow.

In matrix notation, solving equation (1) corresponds to solve

$$Zx = \text{scale} * b,$$

where Z is defined as

$$Z = \begin{bmatrix} \text{kron}(I_n, A) & -\text{kron}(B', I_m) \\ \text{kron}(I_n, D) & -\text{kron}(E', I_m) \end{bmatrix} \quad (2)$$

Here I_k is the identity matrix of size k and X' is the transpose of X . $\text{kron}(X, Y)$ denotes the Kronecker product between the matrices X and Y .

If `trans = 'T'`, solve the transposed (conjugate transposed) system

$$Z'y = \text{scale} * b$$

for y , which is equivalent to solve for R and L in

$$\begin{aligned} A' R + D' L &= \text{scale} * C \\ R B' + L E' &= \text{scale} * (-F) \end{aligned} \quad (3)$$

This case is used to compute an estimate of $\text{Dif}[(A, D), (B, E)] = \text{sigma_min}(Z)$ using reverse communication with `?lacon`.

`?tgsy2` also (for $ijob \geq 1$) contributes to the computation in `?tgsyl` of an upper bound on the separation between two matrix pairs. Then the input $(A, D), (B, E)$ are sub-pencils of the matrix pair (two matrix pairs) in `?tgsyl`. See `?tgsyl` for details.

Input Parameters

- trans** CHARACTER
 If `trans = 'N'`, solve the generalized Sylvester equation (1);
 If `trans = 'T'`: solve the 'transposed' system (3).
- ijob** INTEGER.
 Specifies what kind of functionality is to be performed.
 If `ijob = 0`: solve (1) only.
 If `ijob = 1`: a contribution from this subsystem to a

Frobenius norm-based estimate of the separation between two matrix pairs is computed (look ahead strategy is used);
 If $ijob = 2$: a contribution from this subsystem to a Frobenius norm-based estimate of the separation between two matrix pairs is computed (?gecon on sub-systems is used).
 Not referenced if $trans = 'T'$.

m **INTEGER.**
 On entry, m specifies the order of A and D , and the row dimension of C , F , R and L .

n **INTEGER.**
 On entry, n specifies the order of B and E , and the column dimension of C , F , R and L .

a, b **REAL** for `stgsy2`
DOUBLE PRECISION for `dtgsy2`
COMPLEX for `ctgsy2`
COMPLEX*16 for `ztgsy2`.
 Arrays, **DIMENSION** (lda, m) and (ldb, n), respectively. On entry, a contains an upper (quasi) triangular matrix A and b contains an upper (quasi) triangular matrix B .

lda **INTEGER.**
 The leading dimension of the array a . $lda \geq \max(1, m)$.

ldb **INTEGER.**
 The leading dimension of the array b . $ldb \geq \max(1, n)$.

c, f **REAL** for `stgsy2`
DOUBLE PRECISION for `dtgsy2`
COMPLEX for `ctgsy2`
COMPLEX*16 for `ztgsy2`.
 Arrays, **DIMENSION** (ldc, n) and (ldf, n), respectively. On entry, c contains the right-hand-side of the first matrix equation in (1) and f contains the right-hand-side of the second matrix equation in (1).

<i>ldc</i>	INTEGER. The leading dimension of the array <i>c</i> . $ldc \geq \max(1, m)$.
<i>d, e</i>	REAL for <i>stgsy2</i> DOUBLE PRECISION for <i>dtgsy2</i> COMPLEX for <i>ctgsy2</i> COMPLEX*16 for <i>ztgsy2</i> . Arrays, DIMENSION (<i>ldd, m</i>) and (<i>lde, n</i>), respectively. On entry, <i>d</i> contains an upper triangular matrix <i>D</i> and <i>e</i> contains an upper triangular matrix <i>E</i> .
<i>ldd</i>	INTEGER. The leading dimension of the array <i>d</i> . $ldd \geq \max(1, m)$.
<i>lde</i>	INTEGER. The leading dimension of the array <i>e</i> . $lde \geq \max(1, n)$.
<i>ldf</i>	INTEGER. The leading dimension of the array <i>f</i> . $ldf \geq \max(1, m)$.
<i>rdsum</i>	REAL for <i>stgsy2/ctgsy2</i> DOUBLE PRECISION for <i>dtgsy2/ztgsy2</i> . On entry, the sum of squares of computed contributions to the Dif-estimate under computation by <i>?tgsy1</i> , where the scaling factor <i>rdscal</i> has been factored out.
<i>rdscal</i>	REAL for <i>stgsy2/ctgsy2</i> DOUBLE PRECISION for <i>dtgsy2/ztgsy2</i> . On entry, scaling factor used to prevent overflow in <i>rdsum</i> .
<i>iwork</i>	INTEGER. Used with real flavors only. Workspace array, DIMENSION (<i>m+n+2</i>).

Output Parameters

<i>c</i>	On exit, if <i>ijob</i> = 0, <i>c</i> has been overwritten by the solution <i>R</i> .
<i>f</i>	On exit, if <i>ijob</i> = 0, <i>f</i> has been overwritten by the solution <i>L</i> .

scale REAL for *stgsy2/ctgsy2*
DOUBLE PRECISION for *dtgsy2/ztgsy2*.
On exit, $0 \leq \textit{scale} \leq 1$. If $0 < \textit{scale} < 1$, the solutions *R* and *L* (*C* and *F* on entry) will hold the solutions to a slightly perturbed system, but the input matrices *A*, *B*, *D* and *E* have not been changed. If *scale* = 0, *R* and *L* will hold the solutions to the homogeneous system with *C* = *F* = 0. Normally *scale* = 1.

rdsum On exit, the corresponding sum of squares updated with the contributions from the current sub-system.
If *trans* = 'T', *rdsum* is not touched.
Note that *rdsum* only makes sense when *?tgsy2* is called by *?tgsyl*.

rdscal On exit, *rdscal* is updated with respect to the current contributions in *rdsum*.
If *trans* = 'T', *rdscal* is not touched.
Note that *rdscal* only makes sense when *?tgsy2* is called by *?tgsyl*.

pq INTEGER. Used with real flavors only.
On exit, the number of subsystems (of size 2-by-2, 4-by-4 and 8-by-8) solved by the routine *stgsy2/dtgsy2*.

info INTEGER.
On exit, if *info* is set to
=0: Successful exit
<0: If *info* = -*i*, the *i*-th argument had an illegal value.
>0: The matrix pairs (*A*, *D*) and (*B*, *E*) have common or very close eigenvalues.

?trti2

Computes the inverse of a triangular matrix (unblocked algorithm).

```
call strti2 ( uplo, diag, n, a, lda, info )
call dtrti2 ( uplo, diag, n, a, lda, info )
call ctrti2 ( uplo, diag, n, a, lda, info )
call ztrti2 ( uplo, diag, n, a, lda, info )
```

Discussion

The routine ?trti2 computes the inverse of a real/complex upper or lower triangular matrix.

This is the Level 2 BLAS version of the algorithm.

Input Parameters

<i>uplo</i>	<p>CHARACTER*1.</p> <p>Specifies whether the matrix <i>A</i> is upper or lower triangular.</p> <p>= 'U': Upper triangular</p> <p>= 'L': Lower triangular</p>
<i>diag</i>	<p>CHARACTER*1.</p> <p>Specifies whether or not the matrix <i>A</i> is unit triangular.</p> <p>= 'N': Non-unit triangular</p> <p>= 'U': Unit triangular</p>
<i>n</i>	<p>INTEGER.</p> <p>The order of the matrix <i>A</i>. $n \geq 0$.</p>
<i>a</i>	<p>REAL for strti2</p> <p>DOUBLE PRECISION for dtrti2</p> <p>COMPLEX for ctrti2</p> <p>COMPLEX*16 for ztrti2.</p> <p>Array, DIMENSION (<i>lda</i>, <i>n</i>).</p> <p>On entry, the triangular matrix <i>A</i>. If <i>uplo</i> = 'U', the leading <i>n</i>-by-<i>n</i> upper triangular part of the array <i>a</i></p>

contains the upper triangular matrix, and the strictly lower triangular part of *a* is not referenced. If *uplo* = 'L', the leading *n*-by-*n* lower triangular part of the array *a* contains the lower triangular matrix, and the strictly upper triangular part of *a* is not referenced. If *diag* = 'U', the diagonal elements of *a* are also not referenced and are assumed to be 1.

lda **INTEGER.**
The leading dimension of the array *a*. $lda \geq \max(1,n)$.

Output Parameters

a On exit, the (triangular) inverse of the original matrix, in the same storage format.

info **INTEGER.**
= 0: successful exit
< 0: if *info* = -*k*, the *k*-th argument had an illegal value

xerbla

*Error handling routine called by
LAPACK routines.*

```
call xerbla ( sname, info )
```

Discussion

The routine **xerbla** is an error handler for the LAPACK routines. It is called by a LAPACK routine if an input parameter has an invalid value. A message is printed and execution stops. Installers may consider modifying the **stop** statement in order to call system-specific exception-handling facilities.

Input Parameters

<i>sname</i>	CHARACTER*6 The name of the routine which called <i>xerbla</i> .
<i>info</i>	INTEGER. The position of the invalid parameter in the parameter list of the calling routine.

Vector Mathematical Functions

7

This chapter describes Vector Mathematical Functions Library (VML), which is designed to compute elementary functions on vector arguments. VML is an integral part of the Intel[®] MKL Kernel Library and the VML terminology is used here for simplicity in discussing this group of functions.

VML includes a set of highly optimized implementations of certain computationally expensive core mathematical functions (power, trigonometric, exponential, hyperbolic etc.) that operate on vectors.

Application programs that might significantly improve performance with VML include nonlinear programming software, integrals computation, and many others.

VML functions are divided into the following groups according to the operations they perform:

- [VML Mathematical Functions](#) compute values of elementary functions (such as sine, cosine, exponential, logarithm and so on) on vectors with unit increment indexing.
- [VML Pack/Unpack Functions](#) convert to and from vectors with positive increment indexing, vector indexing and mask indexing (see [Appendix A](#) for details on vector indexing methods).
- [VML Service Functions](#) allow the user to set /get the accuracy mode, and set/get the error code.

VML mathematical functions take an input vector as argument, compute values of the respective elementary function element-wise, and return the results in an output vector.

Data Types and Accuracy Modes

Mathematical and pack/unpack vector functions in VML have been implemented for vector arguments of single and double precision real data. Both Fortran- and C-interfaces to all functions, including VML service functions, are provided in the library. The differences in naming and calling the functions for Fortran- and C-interfaces are detailed in the [Function Naming Conventions](#) section below.

Each vector function from VML (for each data format) can work in two modes: High Accuracy (HA) and Low Accuracy (LA). For many functions, using the LA version will improve performance at the cost of accuracy. For some cases, the advantage of relaxing the accuracy improves performance very little so the same function is employed for both versions. Error behavior depends not only on whether the HA or LA version is chosen, but also depends on the processor on which the software runs. In addition, special value behavior may differ between the HA and LA versions of the functions. Any information on accuracy behavior can be found in the *VML Release Notes*.

Switching between the two modes (HA and LA) is accomplished by using `vm1SetMode(mode)` (see [Table 7-11](#)). The function `vm1GetMode()` will return the currently used mode. The High Accuracy mode is used by default.

Function Naming Conventions

Full names of all VML functions include only lowercase letters for Fortran-interface, whereas for C-interface names the lowercase letters are mixed with uppercase..



NOTE. *This naming convention is followed in the function descriptions in the manual. Actual function names in the library may differ slightly (with respect to lower- and uppercase usage) and will be sufficient to meet the requirements of the supported compilers.*

VML mathematical and pack/unpack function full names have the following structure:

```
v <p> <name> <mod>
```

The initial letter `v` is a prefix indicating that a function belongs to VML.

The `<p>` field is a precision prefix that indicates the data type:

```
s  REAL for Fortran–interface, or float for C–interface
d  DOUBLE PRECISION for Fortran–interface, or double for
   C–interface.
```

The `<name>` field indicates the function short name, with some of its letters in uppercase for C-interface (see [Table 7-2](#), [Table 7-9](#)).

The `<mod>` field (written in uppercase for C-interface) is present in pack/unpack functions only; it indicates the indexing method used:

```
i  indexing with positive increment
v  indexing with index vector
m  indexing with mask vector.
```

VML service function full names have the following structure:

```
vml <name>
```

where `vml` is a prefix indicating that a function belongs to VML, and `<name>` is the function short name, which includes some uppercase letters for C-interface (see [Table 7-10](#)).

To call VML functions from an application program, use conventional function calls. For example, the VML exponential function for single precision data can be called as

```
call vsexp ( n, a, y ) for Fortran–interface, or
vsExp ( n, a, y );   for C–interface.
```

Functions Interface

The interface to VML functions includes function full names and the arguments list.

The Fortran- and C-interface descriptions for different groups of VML functions are given below. Note that some functions (`Div`, `Pow`, and `Atan2`) have two input vectors `a` and `b` as their arguments, while `SinCos` function has two output vectors `y` and `z`.

VML Mathematical Functions:

Fortran:

```
call v<p><name>( n, a, y )  
call v<p><name>( n, a, b, y )  
call v<p><name>( n, a, y, z )
```

C:

```
v<p><name>( n, a, y );  
v<p><name>( n, a, b, y );  
v<p><name>( n, a, y, z );
```

Pack Functions:

Fortran:

```
call v<p>packi( n, a, inca, y )  
call v<p>packv( n, a, ia, y )  
call v<p>packm( n, a, ma, y )
```

C:

```
v<p>PackI( n, a, inca, y );  
v<p>PackV( n, a, ia, y );  
v<p>PackM( n, a, ma, y );
```

Unpack Functions:

Fortran:

```
call v<p>unpacki( n, a, y, incy )  
call v<p>unpackv( n, a, y, iy )  
call v<p>unpackm( n, a, y, my )
```

C:

```
v<p>UnpackI( n, a, y, incy );  
v<p>UnpackV( n, a, y, iy );  
v<p>UnpackM( n, a, y, my );
```

Service Functions:

Fortran:

```
oldmode = vmlsetmode( mode )  
mode    = vmlgetmode( )  
olderr  = vmlseterrstatus ( err )  
err     = vmlgeterrstatus( )  
olderr  = vmlclearerrstatus( )
```

```
oldcallback = vmlseterrorcallback( callback )
callback    = vmlgeterrorcallback( )
oldcallback = vmlclearerrorcallback( )
```

C:

```
oldmode = vmlSetMode( mode );
mode     = vmlGetMode( void);
olderr   = vmlSetErrStatus ( err );
err      = vmlGetErrStatus(void);
olderr   = vmlClearErrStatus(void);
oldcallback = vmlSetErrorCallBack(callback );
callback  = vmlGetErrorCallBack( void );
oldcallback = vmlClearErrorCallBack(void );
```

Input Parameters:

<i>n</i>	number of elements to be calculated
<i>a</i>	first input vector
<i>b</i>	second input vector
<i>inca</i>	vector increment for the input vector <i>a</i>
<i>ia</i>	index vector for the input vector <i>a</i>
<i>ma</i>	mask vector for the input vector <i>a</i>
<i>incy</i>	vector increment for the output vector <i>y</i>
<i>iy</i>	index vector for the output vector <i>y</i>
<i>my</i>	mask vector for the output vector <i>y</i>
<i>err</i>	error code
<i>mode</i>	VML mode
<i>callback</i>	address of the callback function

Output Parameters:

<i>y</i>	first output vector
<i>z</i>	second output vector
<i>err</i>	error code
<i>mode</i>	VML mode
<i>olderr</i>	former error code

oldmode former VML mode
oldcallback address of the former callback function

The data types of the parameters used in each function are specified in the respective function description section. All VML mathematical functions can perform in-place operations, which means that the same vector can be used as both input and output parameter. This holds true for functions with two input vectors as well, in which case one of them may be overwritten with the output vector. For functions with two output vectors, one of them may coincide with the input vector.

Vector Indexing Methods

Current VML mathematical functions work only with unit increment. Arrays with other increments, or more complicated indexing, can be accommodated by gathering the elements into a contiguous vector and then scattering them after the computation is complete.

Three following indexing methods are used to gather/scatter the vector elements in VML:

- positive increment
- index vector
- mask vector.

The indexing method used in a particular function is indicated by the indexing modifier (see the description of the `<mod>` field in [Function Naming Conventions](#)). For more information on indexing methods see [Vector Arguments in VML](#) in Appendix A.

Error Diagnostics

The VML library has its own error handler. The only difference for C- and Fortran- interfaces is that the Intel MKL error reporting routine `XERBLA` can be called after the Fortran- interface VML function encounters an error, and this routine gets information on `VML_STATUS_BADSIZE` and `VML_STATUS_BADMEM` input errors (see [Table 7-13](#)).

The VML error handler has the following properties:

- 1) The Error Status (`vmlErrStatus`) global variable is set after each VML function call. The possible values of this variable are shown in the [Table 7-13](#).
- 2) Depending on the VML mode, the error handler function invokes:
 - `errno` variable setting. The possible values are shown in the [Table 7-1](#)
 - writing error text information to the `stderr` stream
 - raising the appropriate exception on error, if necessary
 - calling the additional error handler callback function.

Table 7-1 Set Values of the `errno` Variable

Value of <code>errno</code>	Description
0	No errors are detected.
<code>EINVAL</code>	The array dimension is not positive.
<code>EACCES</code>	NULL pointer is passed.
<code>EDOM</code>	At least one of array values is out of a range of definition.
<code>ERANGE</code>	At least one of array values caused a singularity, overflow or underflow.

VML Mathematical Functions

This section describes VML functions which compute values of elementary mathematical functions on real vector arguments with unit increment.

Each function group is introduced by its short name, a brief description of its purpose, and the calling sequence for each type of data both for Fortran- and C-interfaces, as well as a description of the input/output arguments.

For all VML mathematical functions, the input range of parameters is equal to the mathematical range of definition in the set of defined values for the respective data type. Several VML functions, specifically `Div`, `Exp`, `Sinh`, `Cosh`, and `Pow`, can result in an overflow. For these functions, the respective input threshold values that mark off the precision overflow are specified in the function description section. Note that in these specifications, `FLT_MAX` denotes the maximum number representable in single precision data type, while `DBL_MAX` denotes the maximum number representable in double precision data type.

[Table 7-2](#) lists available mathematical functions and data types associated with them.

Table 7-2 VML Mathematical Functions

Function Short Name	Data Types	Description
Power and Root Functions		
Inv	s, d	Inversion of the vector elements
Div	s, d	Divide elements of one vector by elements of second vector
Sqrt	s, d	Square root of vector elements
InvSqrt	s, d	Inverse square root of vector elements
Cbrt	s, d	Cube root of vector elements
InvCbrt	s, d	Inverse cube root of vector elements
Pow	s, d	Each vector element raised to the specified power
Powx	s, d	Each vector element raised to the constant power

Table 6-2 VML Mathematical Functions (continued)

Function Short Name	Data Types	Description
Exponential and Logarithmic Functions		
<u>Exp</u>	s, d	Exponential of vector elements
<u>Ln</u>	s, d	Natural logarithm of vector elements
<u>Log10</u>	s, d	Denary logarithm of vector elements
Trigonometric Functions		
<u>Cos</u>	s, d	Cosine of vector elements
<u>Sin</u>	s, d	Sine of vector elements
<u>SinCos</u>	s, d	Sine and cosine of vector elements
<u>Tan</u>	s, d	Tangent of vector elements
<u>Acos</u>	s, d	Inverse cosine of vector elements
<u>Asin</u>	s, d	Inverse sine of vector elements
<u>Atan</u>	s, d	Inverse tangent of vector elements
<u>Atan2</u>	s, d	Four-quadrant inverse tangent of elements of two vectors
Hyperbolic Functions		
<u>Cosh</u>	s, d	Hyperbolic cosine of vector elements
<u>Sinh</u>	s, d	Hyperbolic sine of vector elements
<u>Tanh</u>	s, d	Hyperbolic tangent of vector elements
<u>Acosh</u>	s, d	Inverse hyperbolic cosine (nonnegative) of vector elements
<u>Asinh</u>	s, d	Inverse hyperbolic sine of vector elements
<u>Atanh</u>	s, d	Inverse hyperbolic tangent of vector elements
Special Functions		
<u>Erf</u>	s, d	Error function value of vector elements
<u>Erfc</u>	s, d	Complementary error function value of vector elements

Inv

Performs element by element inversion of the vector.

Fortran:

```
call vsinv( n, a, y )
call vdiv( n, a, y )
```

C:

```
vsInv( n, a, y );
vdInv( n, a, y );
```

Input Parameters

Fortran:

n `INTEGER, INTENT(IN)`. Specifies the number of elements to be calculated.

a `REAL, INTENT(IN)` for `vsinv`
`DOUBLE PRECISION, INTENT(IN)` for `vdiv`
Array, specifies the input vector *a*.

C:

n `int`. Specifies the number of elements to be calculated.

a `const float*` for `vsInv`
`const double*` for `vdInv`
Pointer to an array that contains the input vector *a*.

Output Parameters

Fortran:

y `REAL` for `vsinv`
`DOUBLE PRECISION` for `vdiv`
Array, specifies the output vector *y*.

C:

```
y          float*   for vsInv
           double*  for vdInv
           Pointer to an array that contains the output vector y.
```

Div

*Performs element by element division of vector **a** by vector **b**.*

Fortran:

```
call vsdiv( n, a, b, y )
call vddiv( n, a, b, y )
```

C:

```
vsDiv( n, a, b, y );
vdDiv( n, a, b, y );
```

Input Parameters

Fortran:

```
n          INTEGER, INTENT(IN). Specifies the number of elements
           to be calculated.
a, b       REAL, INTENT(IN)   for vsdiv
           DOUBLE PRECISION, INTENT(IN) for vddiv
           Arrays, specify the input vectors a and b.
```

C:

```
n          int. Specifies the number of elements to be calculated.
a, b       const float*   for vsDiv
           const double*  for vdDiv
           Pointers to arrays that contain the input vectors a and b.
```

Table 7-3 Precision Overflow Thresholds for Div Function

Data Type	Threshold Limitations on Input Parameters
single precision	$\text{abs}(a[i]) < \text{abs}(b[i]) * \text{FLT_MAX}$
double precision	$\text{abs}(a[i]) < \text{abs}(b[i]) * \text{DBL_MAX}$

Output Parameters

Fortran:

y **REAL** for **vsdiv**
 DOUBLE PRECISION for **vddiv**
 Array, specifies the output vector *y*.

C:

y **float*** for **vsDiv**
 double* for **vdDiv**
 Pointer to an array that contains the output vector *y*.

Sqrt

*Computes a square root
of vector elements.*

Fortran:

```
call vssqrt( n, a, y )
call vdsqrt( n, a, y )
```

C:

```
vsSqrt( n, a, y );
vdSqrt( n, a, y );
```

Input Parameters

Fortran:

n `INTEGER, INTENT(IN)`. Specifies the number of elements to be calculated.

a `REAL, INTENT(IN)` for `vssqrt`
`DOUBLE PRECISION, INTENT(IN)` for `vdsqrt`
Array, specifies the input vector *a*.

C:

n `int`. Specifies the number of elements to be calculated.

a `const float*` for `vsSqrt`
`const double*` for `vdSqrt`
Pointer to an array that contains the input vector *a*.

Output Parameters

Fortran:

y `REAL` for `vssqrt`
`DOUBLE PRECISION` for `vdsqrt`
Array, specifies the output vector *y*.

C:

y `float*` for `vsSqrt`
`double*` for `vdSqrt`
Pointer to an array that contains the output vector *y*.

InvSqrt

Computes an inverse square root of vector elements.

Fortran:

```
call vsinvsqrt( n, a, y )  
call vdinvsqrt( n, a, y )
```

C:

```
vsInvSqrt( n, a, y );  
vdInvSqrt( n, a, y );
```

Input Parameters

Fortran:

n **INTEGER, INTENT(IN)**. Specifies the number of elements to be calculated.

a **REAL, INTENT(IN)** for `vsinvsqrt`
DOUBLE PRECISION, INTENT(IN) for `vdinvsqrt`
Array, specifies the input vector *a*.

C:

n **int**. Specifies the number of elements to be calculated.

a **const float*** for `vsInvSqrt`
const double* for `vdInvSqrt`
Pointer to an array that contains the input vector *a*.

Output Parameters

Fortran:

y **REAL** for `vsinvsqrt`
DOUBLE PRECISION for `vdinvsqrt`
Array, specifies the output vector *y*.

C:

y **float*** for `vsInvSqrt`
double* for `vdInvSqrt`
Pointer to an array that contains the output vector *y*.

Cbrt

*Computes a cube root
of vector elements.*

Fortran:

```
call vsqrt( n, a, y )  
call vdsqrt( n, a, y )
```

C:

```
vsCbrt( n, a, y );  
vdCbrt( n, a, y );
```

Input Parameters

Fortran:

n **INTEGER, INTENT(IN)**. Specifies the number of elements to be calculated.

a **REAL, INTENT(IN)** for **vsqrt**
DOUBLE PRECISION, INTENT(IN) for **vdsqrt**
Array, specifies the input vector **a**.

C:

n **int**. Specifies the number of elements to be calculated.

a **const float*** for **vsCbrt**
const double* for **vdCbrt**
Pointer to an array that contains the input vector **a**.

Output Parameters

Fortran:

y **REAL** for **vsqrt**
DOUBLE PRECISION for **vdsqrt**
Array, specifies the output vector **y**.

C:

y **float*** for **vsCbrt**
double* for **vdCbrt**
Pointer to an array that contains the output vector **y**.

InvCbrt

Computes an inverse cube root of vector elements.

Fortran:

```
call vsinvcbrt( n, a, y )
call vdinvcbrt( n, a, y )
```

C:

```
vsInvCbrt( n, a, y );
vdInvCbrt( n, a, y );
```

Input Parameters

Fortran:

n `INTEGER, INTENT(IN)`. Specifies the number of elements to be calculated.

a `REAL, INTENT(IN)` for `vsinvcbrt`
`DOUBLE PRECISION, INTENT(IN)` for `vdinvcbrt`
Array, specifies the input vector *a*.

C:

n `int`. Specifies the number of elements to be calculated.

a `const float*` for `vsInvCbrt`
`const double*` for `vdInvCbrt`
Pointer to an array that contains the input vector *a*.

Output Parameters

Fortran:

y `REAL` for `vsinvcbrt`
`DOUBLE PRECISION` for `vdinvcbrt`
Array, specifies the output vector *y*.

C:

y `float*` for `vsInvCbrt`
`double*` for `vdInvCbrt`
Pointer to an array that contains the output vector *y*.

Pow

Computes a to the power b
for elements of two vectors.

Fortran:

```
call vspow( n, a, b, y )
call vdpow( n, a, b, y )
```

C:

```
vsPow( n, a, b, y );
vdPow( n, a, b, y );
```

Input Parameters

Fortran:

n **INTEGER, INTENT(IN)**. Specifies the number of elements to be calculated.

a, b **REAL, INTENT(IN)** for `vspow`
DOUBLE PRECISION, INTENT(IN) for `vdpow`
 Arrays, specify the input vectors a and b .

C:

n **int**. Specifies the number of elements to be calculated.

a, b **const float*** for `vsPow`
const double* for `vdPow`
 Pointers to arrays that contain the input vectors a and b .

Table 7-4 Precision Overflow Thresholds for Pow Function

Data Type	Threshold Limitations on Input Parameters
single precision	$\text{abs}(a[i]) < (\text{FLT_MAX})^{1/b[i]}$
double precision	$\text{abs}(a[i]) < (\text{DBL_MAX})^{1/b[i]}$

Output Parameters

Fortran:

`y` `REAL` for `vspow`
 `DOUBLE PRECISION` for `vdpow`
 Array, specifies the output vector `y`.

C:

`y` `float*` for `vsPow`
 `double*` for `vdPow`
 Pointer to an array that contains the output vector `y`.

Discussion

The function `Pow` has certain limitations on the input range of `a` and `b` parameters. Specifically, if `a[i]` is positive, then `b[i]` may be arbitrary. For negative or zero `a[i]`, the value of `b[i]` must be integer (either positive or negative).

Powx

*Raises each element of a vector
to the constant power.*

Fortran:

```
call vspowx( n, a, b, y )
call vdpowx( n, a, b, y )
```

C:

```
vsPowx( n, a, b, y );
vdPowx( n, a, b, y );
```

Input Parameters

Fortran:

`n` `INTEGER, INTENT(IN)`. Specifies the number of elements to be calculated.

`a, b` `REAL, INTENT(IN)` for `vspowx`

`DOUBLE PRECISION, INTENT(IN)` for `vdpowx`
 Array `a` specifies the input vector;
 scalar value `b` is the constant power.

C:

`n` `int`. Specifies the number of elements to be calculated.
`a` `const float*` for `vsPowx`
`const double*` for `vdPowx`
 Pointer to an array that contains the input vector `a`.
`b` `const float` for `vsPowx`
`const double` for `vdPowx`
 Constant value for power `b`.

Table 7-5 Precision Overflow Thresholds for `Powx` Function

Data Type	Threshold Limitations on Input Parameters
single precision	$\text{abs}(a[i]) < (\text{FLT_MAX})^{1/b}$
double precision	$\text{abs}(a[i]) < (\text{DBL_MAX})^{1/b}$

Output Parameters

Fortran:

`y` `REAL` for `vspowx`
`DOUBLE PRECISION` for `vdpowx`
 Array, specifies the output vector `y`.

C:

`y` `float*` for `vsPowx`
`double*` for `vdPowx`
 Pointer to an array that contains the output vector `y`.

Discussion

The function `Powx` has certain limitations on the input range of `a` and `b` parameters. Specifically, if `a[i]` is positive, then `b` may be arbitrary. For negative or zero `a[i]`, the value of `b` must be integer (either positive or negative).

Exp

*Computes an exponential
of vector elements.*

Fortran:

```
call vsexp( n, a, y )  
call vdexp( n, a, y )
```

C:

```
vsExp( n, a, y );  
vdExp( n, a, y );
```

Input Parameters

Fortran:

n **INTEGER, INTENT(IN)**. Specifies the number of elements to be calculated.

a **REAL, INTENT(IN)** for **vsexp**
DOUBLE PRECISION, INTENT(IN) for **vdexp**
Array, specifies the input vector **a**.

C:

n **int**. Specifies the number of elements to be calculated.

a **const float*** for **vsExp**
const double* for **vdExp**
Pointer to an array that contains the input vector **a**.

Table 7-6 Precision Overflow Thresholds for `Exp` Function

Data Type	Threshold Limitations on Input Parameters
single precision	$a[i] < \text{Ln}(\text{FLT_MAX})$
double precision	$a[i] < \text{Ln}(\text{DBL_MAX})$

Output Parameters

Fortran:

`y` `REAL` for `vsexp`
 `DOUBLE PRECISION` for `vdexp`
 Array, specifies the output vector `y`.

C:

`y` `float*` for `vsExp`
 `double*` for `vdExp`
 Pointer to an array that contains the output vector `y`.

Ln

*Computes natural logarithm
of vector elements.*

Fortran:

```
call vsln( n, a, y )
call vdln( n, a, y )
```

C:

```
vsLn( n, a, y );
vdLn( n, a, y );
```

Input Parameters

Fortran:

`n` `INTEGER, INTENT(IN)`. Specifies the number of elements
to be calculated.

`a` `REAL, INTENT(IN)` for `vsln`
 `DOUBLE PRECISION, INTENT(IN)` for `vdln`
 Array, specifies the input vector `a`.

C:

`n` `int`. Specifies the number of elements to be calculated.

`a` `const float*` for `vsLn`
 `const double*` for `vdLn`
 Pointer to an array that contains the input vector `a`.

Output Parameters

Fortran:

`y` `REAL` for `vsln`
 `DOUBLE PRECISION` for `vdln`
 Array, specifies the output vector `y`.

C:

`y` `float*` for `vsLn`
 `double*` for `vdLn`
 Pointer to an array that contains the output vector `y`.

Log10

*Computes denary logarithm
of vector elements.*

Fortran:

```
call vslog10( n, a, y )
call vdlog10( n, a, y )
```

C:

```
vsLog10( n, a, y );
vdLog10( n, a, y );
```

Input Parameters

Fortran:

n `INTEGER, INTENT(IN)`. Specifies the number of elements to be calculated.

a `REAL, INTENT(IN)` for `vslog10`
`DOUBLE PRECISION, INTENT(IN)` for `vdlog10`
Array, specifies the input vector *a*.

C:

n `int`. Specifies the number of elements to be calculated.

a `const float*` for `vsLog10`
`const double*` for `vdLog10`
Pointer to an array that contains the input vector *a*.

Output Parameters

Fortran:

y `REAL` for `vslog10`
`DOUBLE PRECISION` for `vdlog10`
Array, specifies the output vector *y*.

C:

y `float*` for `vsLog10`
`double*` for `vdLog10`
Pointer to an array that contains the output vector *y*.

Cos

Computes cosine of vector elements.

Fortran:

```
call vscos( n, a, y )  
call vdcos( n, a, y )
```

C:

```
vsCos( n, a, y );  
vdCos( n, a, y );
```

Input Parameters

Fortran:

n `INTEGER, INTENT(IN)`. Specifies the number of elements to be calculated.

a `REAL, INTENT(IN)` for `vscos`
`DOUBLE PRECISION, INTENT(IN)` for `vdcos`
Array, specifies the input vector *a*.

C:

n `int`. Specifies the number of elements to be calculated.

a `const float*` for `vsCos`
`const double*` for `vdCos`
Pointer to an array that contains the input vector *a*.

Output Parameters

Fortran:

y `REAL` for `vscos`
`DOUBLE PRECISION` for `vdcos`
Array, specifies the output vector *y*.

C:

y `float*` for `vsCos`
`double*` for `vdCos`
Pointer to an array that contains the output vector *y*.

Sin

Computes sine of vector elements.

Fortran:

```
call vssin( n, a, y )  
call vdsin( n, a, y )
```

C:

```
vsSin( n, a, y );  
vdSin( n, a, y );
```

Input Parameters

Fortran:

n `INTEGER, INTENT(IN)`. Specifies the number of elements to be calculated.

a `REAL, INTENT(IN)` for `vssin`
`DOUBLE PRECISION, INTENT(IN)` for `vdSin`
Array, specifies the input vector *a*.

C:

n `int`. Specifies the number of elements to be calculated.

a `const float*` for `vsSin`
`const double*` for `vdSin`
Pointer to an array that contains the input vector *a*.

Output Parameters

Fortran:

y `REAL` for `vssin`
`DOUBLE PRECISION` for `vdSin`
Array, specifies the output vector *y*.

C:

y `float*` for `vsSin`
`double*` for `vdSin`
Pointer to an array that contains the output vector *y*.

SinCos

Computes sine and cosine of vector elements.

Fortran:

```
call vssincos( n, a, y, z )
call vdsincos( n, a, y, z )
```

C:

```
vsSinCos( n, a, y, z );
vdSinCos( n, a, y, z );
```

Input Parameters

Fortran:

n `INTEGER, INTENT(IN)`. Specifies the number of elements to be calculated.

a `REAL, INTENT(IN)` for `vssincos`
`DOUBLE PRECISION, INTENT(IN)` for `vdsincos`
 Array, specifies the input vector *a*.

C:

n `int`. Specifies the number of elements to be calculated.

a `const float*` for `vsSinCos`
`const double*` for `vdSinCos`
 Pointer to an array that contains the input vector *a*.

Output Parameters

Fortran:

y, z `REAL` for `vssincos`
`DOUBLE PRECISION` for `vdsincos`
 Arrays, specify the output vectors *y* (for sine values) and *z* (for cosine values).

C:

y, z `float*` for `vsSinCos`
`double*` for `vdSinCos`
 Pointers to arrays that contain the output vectors *y* (for sine values) and *z* (for cosine values).

Tan

Computes tangent of vector elements.

Fortran:

```
call vstan( n, a, y )
call vdtan( n, a, y )
```

C:

```
vsTan( n, a, y );
vdTan( n, a, y );
```

Input Parameters

Fortran:

n **INTEGER, INTENT(IN)**. Specifies the number of elements to be calculated.

a **REAL, INTENT(IN)** for **vstan**
DOUBLE PRECISION, INTENT(IN) for **vdtan**
Array, specifies the input vector **a**.

C:

n **int**. Specifies the number of elements to be calculated.

a **const float*** for **vsTan**
const double* for **vdTan**
Pointer to an array that contains the input vector **a**.

Output Parameters

Fortran:

y **REAL** for **vstan**
DOUBLE PRECISION for **vdtan**
Array, specifies the output vector **y**.

C:

y **float*** for **vsTan**
double* for **vdTan**
Pointer to an array that contains the output vector **y**.

Acos

*Computes inverse cosine
of vector elements.*

Fortran:

```
call vsacos( n, a, y )  
call vdacos( n, a, y )
```

C:

```
vsAcos( n, a, y );  
vdAcos( n, a, y );
```

Input Parameters

Fortran:

n **INTEGER, INTENT(IN)**. Specifies the number of elements to be calculated.

a **REAL, INTENT(IN)** for **vsacos**
DOUBLE PRECISION, INTENT(IN) for **vdacos**
Array, specifies the input vector **a**.

C:

n **int**. Specifies the number of elements to be calculated.

a **const float*** for **vsAcos**
const double* for **vdAcos**
Pointer to an array that contains the input vector **a**.

Output Parameters

Fortran:

y **REAL** for **vsacos**
DOUBLE PRECISION for **vdacos**
Array, specifies the output vector **y**.

C:

y **float*** for **vsAcos**
double* for **vdAcos**
Pointer to an array that contains the output vector **y**.

Asin

*Computes inverse sine
of vector elements.*

Fortran:

```
call vsasin( n, a, y )  
call vdasin( n, a, y )
```

C:

```
vsAsin( n, a, y );  
vdAsin( n, a, y );
```

Input Parameters

Fortran:

n **INTEGER, INTENT(IN)**. Specifies the number of elements to be calculated.

a **REAL, INTENT(IN)** for *vsasin*
DOUBLE PRECISION, INTENT(IN) for *vdasin*
Array, specifies the input vector *a*.

C:

n **int**. Specifies the number of elements to be calculated.

a **const float*** for *vsAsin*
const double* for *vdAsin*
Pointer to an array that contains the input vector *a*.

Output Parameters

Fortran:

y **REAL** for *vsasin*
DOUBLE PRECISION for *vdasin*
Array, specifies the output vector *y*.

C:

`y` `float*` for `vsAsin`
 `double*` for `vdAsin`
 Pointer to an array that contains the output vector `y`.

Atan

*Computes inverse tangent
of vector elements.*

Fortran:

```
call vsatan( n, a, y )
call vdatan( n, a, y )
```

C:

```
vsAtan( n, a, y );
vdAtan( n, a, y );
```

Input Parameters

Fortran:

`n` `INTEGER, INTENT(IN)`. Specifies the number of elements to be calculated.

`a` `REAL, INTENT(IN)` for `vsatan`
 `DOUBLE PRECISION, INTENT(IN)` for `vdatan`
 Array, specifies the input vector `a`.

C:

`n` `int`. Specifies the number of elements to be calculated.

`a` `const float*` for `vsAtan`
 `const double*` for `vdAtan`
 Pointer to an array that contains the input vector `a`.

Output Parameters

Fortran:

y `REAL` for `vsatan`
 `DOUBLE PRECISION` for `vdatan`
 Array, specifies the output vector *y*.

C:

y `float*` for `vsAtan`
 `double*` for `vdAtan`
 Pointer to an array that contains the output vector *y*.

Atan2

*Computes four-quadrant inverse
 tangent of elements of two vectors.*

Fortran:

```
call vsatan2( n, a, b, y )
call vdatan2( n, a, b, y )
```

C:

```
vsAtan2( n, a, b, y );
vdAtan2( n, a, b, y );
```

Input Parameters

Fortran:

n `INTEGER, INTENT(IN)`. Specifies the number of elements
 to be calculated.

a, b `REAL, INTENT(IN)` for `vsatan2`
 `DOUBLE PRECISION, INTENT(IN)` for `vdatan2`
 Arrays, specify the input vectors *a* and *b*.

C:

n `int`. Specifies the number of elements to be calculated.

a, b `const float*` for `vsAtan2`
 `const double*` for `vdAtan2`
 Pointers to arrays that contain the input vectors *a* and *b*.

Output Parameters

Fortran:

y **REAL** for `vsatan2`
 DOUBLE PRECISION for `vdatan2`
Array, specifies the output vector *y*.

C:

y **float*** for `vsAtan2`
 double* for `vdAtan2`
Pointer to an array that contains the output vector *y*.

The elements of the output vector *y* are computed as the four-quadrant arctangent of $a[i] / b[i]$.

Cosh

*Computes hyperbolic cosine
of vector elements.*

Fortran:

```
call vscosh( n, a, y )  
call vdcosh( n, a, y )
```

C:

```
vsCosh( n, a, y );  
vdCosh( n, a, y );
```

Input Parameters

Fortran:

n **INTEGER, INTENT(IN)**. Specifies the number of elements
to be calculated.

`a` `REAL, INTENT(IN)` for `vscosh`
 `DOUBLE PRECISION, INTENT(IN)` for `vdcosh`
 Array, specifies the input vector `a`.

C:

`n` `int`. Specifies the number of elements to be calculated.

`a` `const float*` for `vsCosh`
 `const double*` for `vdCosh`
 Pointer to an array that contains the input vector `a`.

Table 7-7 Precision Overflow Thresholds for Cosh Function

Data Type	Threshold Limitations on Input Parameters
single precision	$-\text{Ln}(\text{FLT_MAX}) - \text{Ln}2 < a[i] < \text{Ln}(\text{FLT_MAX}) + \text{Ln}2$
double precision	$-\text{Ln}(\text{DBL_MAX}) - \text{Ln}2 < a[i] < \text{Ln}(\text{DBL_MAX}) + \text{Ln}2$

Output Parameters

Fortran:

`y` `REAL` for `vscosh`
 `DOUBLE PRECISION` for `vdcosh`
 Array, specifies the output vector `y`.

C:

`y` `float*` for `vsCosh`
 `double*` for `vdCosh`
 Pointer to an array that contains the output vector `y`.

Sinh

*Computes hyperbolic sine
of vector elements.*

Fortran:

```
call vssinh( n, a, y )  
call vdsinh( n, a, y )
```

C:

```
vsSinh( n, a, y );  
vdSinh( n, a, y );
```

Input Parameters

Fortran:

n `INTEGER, INTENT(IN)`. Specifies the number of elements to be calculated.

a `REAL, INTENT(IN)` for `vssinh`
`DOUBLE PRECISION, INTENT(IN)` for `vdsinh`
Array, specifies the input vector *a*.

C:

n `int`. Specifies the number of elements to be calculated.

a `const float*` for `vsSinh`
`const double*` for `vdSinh`
Pointer to an array that contains the input vector *a*.

Table 7-8 Precision Overflow Thresholds for `sinh` Function

Data Type	Threshold Limitations on Input Parameters
single precision	$-\text{Ln}(\text{FLT_MAX}) - \text{Ln}2 < a[i] < \text{Ln}(\text{FLT_MAX}) + \text{Ln}2$
double precision	$-\text{Ln}(\text{DBL_MAX}) - \text{Ln}2 < a[i] < \text{Ln}(\text{DBL_MAX}) + \text{Ln}2$

Output Parameters

Fortran:

`y` `REAL` for `vssinh`
 `DOUBLE PRECISION` for `vdsinh`
 Array, specifies the output vector `y`.

C:

`y` `float*` for `vsSinh`
 `double*` for `vdSinh`
 Pointer to an array that contains the output vector `y`.

Tanh

*Computes hyperbolic tangent
of vector elements.*

Fortran:

```
call vstanh( n, a, y )
call vdtanh( n, a, y )
```

C:

```
vsTanh( n, a, y );
vdTanh( n, a, y );
```

Input Parameters

Fortran:

`n` `INTEGER, INTENT(IN)`. Specifies the number of elements
to be calculated.

`a` `REAL, INTENT(IN)` for `vstanh`
 `DOUBLE PRECISION, INTENT(IN)` for `vdTanh`
 Array, specifies the input vector `a`.

C:

`n` `int`. Specifies the number of elements to be calculated.

`a` `const float*` for `vsTanh`
 `const double*` for `vdTanh`
 Pointer to an array that contains the input vector `a`.

Output Parameters

Fortran:

`y` `REAL` for `vstanh`
 `DOUBLE PRECISION` for `vdTanh`
 Array, specifies the output vector `y`.

C:

`y` `float*` for `vsTanh`
 `double*` for `vdTanh`
 Pointer to an array that contains the output vector `y`.

Acosh

*Computes inverse hyperbolic cosine
 (nonnegative) of vector elements.*

Fortran:

```
call vsacosh( n, a, y )
call vdacosh( n, a, y )
```

C:

```
vsAcosh( n, a, y );
vdAcosh( n, a, y );
```

Input Parameters

Fortran:

n `INTEGER, INTENT(IN)`. Specifies the number of elements to be calculated.

a `REAL, INTENT(IN)` for `vsacosh`
`DOUBLE PRECISION, INTENT(IN)` for `vdacosh`
Array, specifies the input vector *a*.

C:

n `int`. Specifies the number of elements to be calculated.

a `const float*` for `vsAcosh`
`const double*` for `vdAcosh`
Pointer to an array that contains the input vector *a*.

Output Parameters

Fortran:

y `REAL` for `vsacosh`
`DOUBLE PRECISION` for `vdacosh`
Array, specifies the output vector *y*.

C:

y `float*` for `vsAcosh`
`double*` for `vdAcosh`
Pointer to an array that contains the output vector *y*.

Asinh

Computes inverse hyperbolic sine of vector elements.

Fortran:

```
call vsasinh( n, a, y )  
call vdasinh( n, a, y )
```

C:

```
vsAsinh( n, a, y );  
vdAsinh( n, a, y );
```

Input Parameters

Fortran:

n `INTEGER, INTENT(IN)`. Specifies the number of elements to be calculated.

a `REAL, INTENT(IN)` for `vsasinh`
`DOUBLE PRECISION, INTENT(IN)` for `vdasinh`
Array, specifies the input vector *a*.

C:

n `int`. Specifies the number of elements to be calculated.

a `const float*` for `vsAsinh`
`const double*` for `vdAsinh`
Pointer to an array that contains the input vector *a*.

Output Parameters

Fortran:

y `REAL` for `vsasinh`
`DOUBLE PRECISION` for `vdasinh`
Array, specifies the output vector *y*.

C:

y `float*` for `vsAsinh`
`double*` for `vdAsinh`
Pointer to an array that contains the output vector *y*.

Atanh

Computes inverse hyperbolic tangent of vector elements.

Fortran:

```
call vsatanh( n, a, y )
call vdatanh( n, a, y )
```


C:

```
vsAtanh( n, a, y );
vdAtanh( n, a, y );
```

Input Parameters

Fortran:

n **INTEGER, INTENT(IN)**. Specifies the number of elements to be calculated.

a **REAL, INTENT(IN)** for **vsatanh**
DOUBLE PRECISION, INTENT(IN) for **vdatanh**
 Array, specifies the input vector **a**.

C:

n **int**. Specifies the number of elements to be calculated.

a **const float*** for **vsAtanh**
const double* for **vdAtanh**
 Pointer to an array that contains the input vector **a**.

Output Parameters

Fortran:

y **REAL** for **vsatanh**
DOUBLE PRECISION for **vdatanh**
 Array, specifies the output vector **y**.

C:

y **float*** for **vsAtanh**
double* for **vdAtanh**
 Pointer to an array that contains the output vector **y**.

Erf

Computes the error function value of vector elements.

Fortran:

```
call vserf( n, a, y )
call vderf( n, a, y )
```

C:

```
vsErf( n, a, y );
vdErf( n, a, y );
```

Input Parameters

Fortran:

n **INTEGER, INTENT(IN)**. Specifies the number of elements to be calculated.

a **REAL, INTENT(IN)** for **vserf**
DOUBLE PRECISION, INTENT(IN) for **vderf**
 Array, specifies the input vector **a**.

C:

n **int**. Specifies the number of elements to be calculated.

a **const float*** for **vsErf**
const double* for **vdErf**
 Pointer to an array that contains the input vector **a**.

Output Parameters

Fortran:

y **REAL** for **vserf**
DOUBLE PRECISION for **vderf**
 Array, specifies the output vector **y**.

C:

y **float*** for **vsErf**
double* for **vdErf**
 Pointer to an array that contains the output vector **y**.

Discussion

The function **Erf** computes the error function values for elements of the input vector **a** and writes them to the output vector **y**.

The error function is defined as given by:

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$$

Erfc

Computes the complementary error function value of vector elements.

Fortran:

```
call vserfc( n, a, y )
call vderfc( n, a, y )
```

C:

```
vsErfc( n, a, y );
vdErfc( n, a, y );
```

Input Parameters

Fortran:

n **INTEGER, INTENT(IN)**. Specifies the number of elements to be calculated.

a **REAL, INTENT(IN)** for **vserfc**
DOUBLE PRECISION, INTENT(IN) for **vderfc**
Array, specifies the input vector **a**.

C:

n **int**. Specifies the number of elements to be calculated.

a **const float*** for **vsErfc**
const double* for **vdErfc**
Pointer to an array that contains the input vector **a**.

Output Parameters

Fortran:

y **REAL** for **vserfc**
DOUBLE PRECISION for **vderfc**
Array, specifies the output vector **y**.

C:

`y` `float*` for `vsErfc`
 `double*` for `vdErfc`
 Pointer to an array that contains the output vector `y`.

Discussion

The function `Erfc` computes the complementary error function values for elements of the input vector `a` and writes them to the output vector `y`.

The complementary error function is defined as given by:

$$\operatorname{erfc}(x) = 1 - \operatorname{erf}(x)$$

or, in other words,

$$\operatorname{erfc}(x) = \frac{2}{\sqrt{\pi}} \int_x^{\infty} e^{-t^2} dt .$$

VML Pack/Unpack Functions

This section describes VML functions which convert vectors with unit increment to and from vectors with positive increment indexing, vector indexing and mask indexing (see [Appendix A](#) for details on vector indexing methods).

[Table 7-9](#) lists available VML Pack/Unpack functions, together with data types and indexing methods associated with them.

Table 7-9 VML Pack/Unpack Functions

Function Short Name	Data Types	Indexing Methods	Description
Pack	s, d	I,V,M	Gathers elements of arrays, indexed by different methods.
Unpack	s, d	I,V,M	Scatters vector elements to arrays with different indexing.

Pack

*Copies elements of an array
with specified indexing to
a vector with unit increment.*

Fortran:

```
call vspacki( n, a, inca, y )
call vspackv( n, a, ia, y )
call vspackm( n, a, ma, y )
call vdpacki( n, a, inca, y )
call vdpackv( n, a, ia, y )
call vdpackm( n, a, ma, y )
```

C:

```

vsPackI( n, a, inca, y );
vsPackV( n, a, ia, y );
vsPackM( n, a, ma, y );
vdPackI( n, a, inca, y );
vdPackV( n, a, ia, y );
vdPackM( n, a, ma, y );

```

Input Parameters

Fortran:

n **INTEGER, INTENT(IN)**. Specifies the number of elements to be calculated.

a **REAL, INTENT(IN)** for *vspacki*, *vspackv*, *vspackm*
DOUBLE PRECISION, INTENT(IN) for *vdpacki*,
vdpackv, *vdpackm*
 Array, **DIMENSION** at least $(1 + (n-1)*inca)$ for *vspacki*,
 at least $\max(n, \max(ia[j]))$, $j=0, \dots, n-1$, for *vspackv*,
 at least *n* for *vspackm*,
 Specifies the input vector *a*.

inca **INTEGER, INTENT(IN)** for *vspacki*, *vdpacki*.
 Specifies the increment for the elements of *a*.

ia **INTEGER, INTENT(IN)** for *vspackv*, *vdpackv*.
 Array, **DIMENSION** at least *n*
 Specifies the index vector for the elements of *a*.

ma **INTEGER, INTENT(IN)** for *vspackm*, *vdpackm*.
 Array, **DIMENSION** at least *n*
 Specifies the mask vector for the elements of *a*.

C:

n **int**. Specifies the number of elements to be calculated

a **const float*** for *vsPackI*, *vsPackV*, *vsPackM*
const double* for *vdPackI*, *vdPackV*, *vdPackM*
 Specifies the pointer to an array that contains the input vector *a*.
 Size of the array must be:

at least $(1 + (n-1)*inca)$ for `vsPackI`,
 at least $\max(n, \max(ia[j]), j=0, \dots, n-1)$, for `vsPackV`,
 at least n for `vsPackM`.

`inca` `int` for `vsPackI`, `vdPackI`.

Specifies the increment for the elements of `a`.

`ia` `const int*` for `vsPackV`, `vdPackV`. Specifies the pointer to
 an array of size at least n that contains the index vector
 for the elements of `a`.

`ma` `const int*` for `vsPackM`, `vdPackM`. Specifies the pointer to
 an array of size at least n that contains the mask vector
 for the elements of `a`.

Output Parameters

Fortran:

`y` `REAL` for `vspacki`, `vspackv`, `vspackm`
`DOUBLE PRECISION` for `vdpacki`, `vdpackv`, `vdpackm`
 Array, `DIMENSION` at least n , specifies the output vector `y`.

C:

`y` `float*` for `vsPackI`, `vsPackV`, `vsPackM`
`double*` for `vdPackI`, `vdPackV`, `vdPackM`
 Specifies the pointer to an array of size at least n that contains
 the output vector `y`.

Unpack

*Copies elements of a vector with unit increment
 to an array with specified indexing.*

Fortran:

```
call vsunpacki( n, a, y, incy )
call vsunpackv( n, a, y, iy )
call vsunpackm( n, a, y, my )
```

```

call vdunpacki( n, a, y, incy )
call vdunpackv( n, a, y, iy )
call vdunpackm( n, a, y, my )

```

C:

```

vsUnpackI( n, a, y, incy );
vsUnpackV( n, a, y, iy );
vsUnpackM( n, a, y, my );
vdUnpackI( n, a, y, incy );
vdUnpackV( n, a, y, iy );
vdUnpackM( n, a, y, my );

```

Input Parameters

Fortran:

n **INTEGER, INTENT(IN)**. Specifies the number of elements to be calculated.

a **REAL, INTENT(IN)** for **vsunpacki**, **vsunpackv**, **vsunpackm**
DOUBLE PRECISION, INTENT(IN) for **vdunpacki**, **vdunpackv**, **vdunpackm**.
 Array, **DIMENSION** at least *n*, specifies the input vector **a**.

incy **INTEGER, INTENT(IN)** for **vsunpacki**, **vdunpacki**.
 Specifies the increment for the elements of **y**.

iy **INTEGER, INTENT(IN)** for **vsunpackv**, **vdunpackv**.
 Array, **DIMENSION** at least *n*, specifies the index vector for the elements of **y**.

my **INTEGER, INTENT(IN)** for **vsunpackm**, **vdunpackm**.
 Array, **DIMENSION** at least *n*, specifies the mask vector for the elements of **y**.

C:

n **int**. Specifies the number of elements to be calculated.

a **const float*** for **vsUnpackI**, **vsUnpackV**, **vsUnpackM**
const double* for **vdUnpackI**, **vdUnpackV**, **vdUnpackM**
 Specifies the pointer to an array of size at least *n* that contains the input vector **a**.

`incy` `int` for `vsUnpackI`, `vdUnpackI`.
 Specifies the increment for the elements of `y`.

`iy` `const int*` for `vsUnpackV`, `vdUnpackV`. Specifies the pointer to an array of size at least `n` that contains the index vector for the elements of `a`.

`my` `const int*` for `vsUnpackM`, `vdUnpackM`. Specifies the pointer to an array of size at least `n` that contains the mask vector for the elements of `a`.

Output Parameters

Fortran:

`y` `REAL` for `vsunpacki`, `vsunpackv`, `vsunpackm`
`DOUBLE PRECISION` for `vdunpacki`, `vdunpackv`,
`vdunpackm`.
 Array, `DIMENSION`
 at least $(1 + (n-1)*incy)$ for `vsunpacki`,
 at least $\max(n, \max(iy[j]), j=0, \dots, n-1)$, for `vsunpackv`,
 at least `n` for `vsunpackm`
 Specifies the output vector `y`.

C:

`y` `float*` for `vsUnpackI`, `vsUnpackV`, `vsUnpackM`
`double*` for `vdUnpackI`, `vdUnpackV`, `vdUnpackM`
 Specifies the pointer to an array that contains the output vector `y`.
 Size of the array must be:
 at least $(1 + (n-1)*incy)$ for `vsUnPackI`,
 at least $\max(n, \max(ia[j]), j=0, \dots, n-1)$, for `vsUnPackV`,
 at least `n` for `vsUnPackM`.

VML Service Functions

This section describes VML functions which allow the user to set /get the accuracy mode, and set/get the error code. All these functions are available both in Fortran- and C- interfaces.

[Table 7-10](#) lists available VML Service functions and their short description.

Table 7-10 VML Service Functions

Function Short Name	Description
SetMode	Sets the VML mode
GetMode	Gets the VML mode
SetErrStatus	Sets the VML error status
GetErrStatus	Gets the VML error status
ClearErrStatus	Clears the VML error status
SetErrorCallBack	Sets the additional error handler callback function
GetErrorCallBack	Gets the additional error handler callback function
ClearErrorCallBack	Deletes the additional error handler callback function

SetMode

Sets the new mode for VML functions according to `mode` parameter and stores the previous VML mode to `oldmode`.

Fortran:

```
oldmode = vmlsetmode( mode )
```

C:

```
oldmode = vmlSetMode( mode );
```

Input Parameters

Fortran:

mode `INTEGER, INTENT(IN)`. Specifies the VML mode to be set.

C:

mode `int`. Specifies the VML mode to be set.

Output Parameters

Fortran:

oldmode `INTEGER`. Specifies the former VML mode.

C:

oldmode `int`. Specifies the former VML mode.

Discussion

The *mode* parameter is designed to control accuracy, FPU and error handling options. [Table 7-11](#) lists values of the *mode* parameter. All other possible values of the *mode* parameter may be obtained from these values by using bitwise OR (|) operation to combine one value for accuracy, one for FPU, and one for error control options. The default value of the *mode* parameter is `VML_HA | VML_ERRMODE_DEFAULT`. Thus, the current FPU control word (FPU precision and the rounding method) is used by default.

If any VML mathematical function requires different FPU precision, or rounding method, it changes these options automatically and then restores the former values. The *mode* parameter enables you to minimize switching the internal FPU mode inside each VML mathematical function that works with similar precision and accuracy settings. To accomplish this, set the *mode* parameter to `VML_FLOAT_CONSISTENT` for single precision functions, or to `VML_DOUBLE_CONSISTENT` for double precision functions. These values of the *mode* parameter are the optimal choice for the respective function groups, as they are required for most of the VML mathematical functions. After the execution is over, set the *mode* to `VML_RESTORE` if you need to restore the previous FPU mode.

Table 7-11 Values of the *mode* Parameter

Value of <i>mode</i>	Description
Accuracy Control	
VML_HA	High accuracy versions of VML functions will be used
VML_LA	Low accuracy versions of VML functions will be used
Additional FPU Mode Control	
VML_FLOAT_CONSISTENT	The optimal FPU mode (control word) for single precision functions is set, and the previous FPU mode is saved
VML_DOUBLE_CONSISTENT	The optimal FPU mode (control word) for double precision functions is set, and the previous FPU mode is saved
VML_RESTORE	The previously saved FPU mode is restored
Error Mode Control	
VML_ERRMODE_IGNORE	No action is set for computation errors
VML_ERRMODE_ERRNO	On error, the <code>errno</code> variable is set
VML_ERRMODE_STDERR	On error, the error text information is written to <code>stderr</code>
VML_ERRMODE_EXCEPT	On error, an exception is raised
VML_ERRMODE_CALLBACK	On error, an additional error handler function is called
VML_ERRMODE_DEFAULT	On error, the <code>errno</code> variable is set, an exception is raised, and an additional error handler function is called

Examples

Several examples of calling the function `vmlSetMode()` with different values of the *mode* parameter are given below:

Fortran:

```
oldmode = vmlsetmode( VML_LA )
call vmlsetmode( IOR(VML_LA, IOR(VML_FLOAT_CONSISTENT,
                               VML_ERRMODE_IGNORE )))
call vmlsetmode( VML_RESTORE)
```

C:

```
vmlSetMode( VML_LA );  
vmlSetMode( VML_LA | VML_FLOAT_CONSISTENT | VML_ERRMODE_IGNORE );  
vmlSetMode( VML_RESTORE );
```

GetMode

Gets the VML mode.

Fortran:

```
mod = vmlgetmode()
```

C:

```
mod = vmlGetMode( void );
```

Output Parameters

Fortran:

mod **INTEGER**. Specifies the packed *mode* parameter.

C:

mod **int**. Specifies the packed *mode* parameter.

Discussion

The function `vmlGetMode()` returns the VML *mode* parameter which controls accuracy, FPU and error handling options. The *mod* variable value is some combination of the values listed in the [Table 7-11](#). You can obtain some of these values using the respective mask from the [Table 7-12](#), for example:

Fortran:

```
mod = vmlgetmode()  
accm = IAND(mod, VML_ACCURACY_MASK)  
fpum = IAND(mod, VML_FPUMODE_MASK)  
errm = IAND(mod, VML_ERRMODE_MASK)
```

C:

```

accm = vmlGetMode(void )& VML_ACCURACY_MASK;
fpum = vmlGetMode(void )& VML_FPUMODE _MASK;
errm = vmlGetMode(void )& VML_ERRMODE _MASK;

```

Table 7-12 Values of Mask for the *mode* Parameter

Value of mask	Description
VML_ACCURACY_MASK	Specifies mask for accuracy <i>mode</i> selection.
VML_FPUMODE_MASK	Specifies mask for FPU <i>mode</i> selection.
VML_ERRMODE_MASK	Specifies mask for error <i>mode</i> selection.

SetErrStatus

Sets the new VML error status according to *err* and stores the previous VML error status to *olderr*.

Fortran:

```
olderr = vmlseterrstatus( err )
```

C:

```
olderr = vmlSetErrStatus( err );
```

Input Parameters

Fortran:

err INTEGER, INTENT(IN). Specifies the VML error status to be set.

C:

err int. Specifies the VML error status to be set.

Output Parameters

Fortran:

olderr **INTEGER**. Specifies the former VML error status.

C:

olderr **int**. Specifies the former VML error status.

[Table 7-13](#) lists possible values of the *err* parameter.

Table 7-13 **Values of the VML Error Status**

Error Status	Description
VML_STATUS_OK	The execution was completed successfully.
VML_STATUS_BADSIZE	The array dimension is not positive.
VML_STATUS_BADMEM	NULL pointer is passed.
VML_STATUS_ERRDOM	At least one of array values is out of a range of definition.
VML_STATUS_SING	At least one of array values caused a singularity.
VML_STATUS_OVERFLOW	An overflow has happened during the calculation process.
VML_STATUS_UNDERFLOW	An underflow has happened during the calculation process.

Examples:

```
vmlSetErrStatus( VML_STATUS_OK );
vmlSetErrStatus( VML_STATUS_ERRDOM );
vmlSetErrStatus( VML_STATUS_UNDERFLOW );
```

GetErrStatus

Gets the VML error status.

Fortran:

```
err = vmlgeterrstatus( )
```

C:

```
err = vmlGetErrStatus( void );
```

Output Parameters

Fortran:

`err` **INTEGER**. Specifies the VML error status.

C:

`err` **int**. Specifies the VML error status.

ClearErrStatus

*Sets the VML error status to **VML_STATUS_OK** and stores the previous VML error status to `olderr`.*

Fortran:

```
olderr = vmlclearerrstatus( )
```

C:

```
olderr = vmlClearErrStatus( void );
```

Output Parameters

Fortran:

`olderr` **INTEGER**. Specifies the former VML error status.

C:

olderr *int*. Specifies the former VML error status.

SetErrorCallback

Sets the additional error handler callback function and gets the old callback function.

Fortran:

```
oldcallback = vmlseterrorcallback( callback )
```

C:

```
oldcallback = vmlSetErrorCallBack( callback );
```

Input Parameters

Fortran:

```
callback      Address of the callback function.  

                 The callback function has the following format:  

INTEGER FUNCTION ERRFUNC(par)  

  TYPE (ERROR_STRUCTURE) par  

  ! ...  

  ! user error processing  

  ! ...  

  ERRFUNC = 0  

  ! if ERRFUNC = 0 - standard VML error  

  handler  

  ! is called after the callback  

  ! if ERRFUNC != 0 - standard VML error  

  handler  

  ! is not called  

END
```

The passed error structure is defined as follows:

```

TYPE ERROR_STRUCTURE
SEQUENCE
INTEGER*4 ICODE
INTEGER*4 IINDEX
REAL*8 DBA1
REAL*8 DBA2
REAL*8 DBR1
REAL*8 DBR2
CHARACTER(64) CFUNCNAME
INTEGER*4 IFUNCNAMELEN
END TYPE ERROR_STRUCTURE

```

C:

callback

Pointer to the callback function.

The callback function has the following format:

```

static int __stdcall
MyHandler(DefVmlErrorContext*
pContext)
{
    /* Handler body */
};

```

The passed error structure is defined as follows:

```

typedef struct _DefVmlErrorContext
{
    int iCode;          /* Error status value */
    int iIndex;        /* Index for bad array
                        element, or bad array
                        dimension, or bad
                        array pointer */
    double dbA1;       /* Error argument 1 */
    double dbA2;       /* Error argument 2 */
    double dbR1;       /* Error result 1 */
    double dbR2;       /* Error result 2 */
    char cFuncName[64]; /* Function name */
    int iFuncNameLen;  /* Length of function
                        name*/
} DefVmlErrorContext;

```

Output Parameters

Fortran:

oldcallback Address of the former callback function.

C:

oldcallback Pointer to the former callback function.

Discussion

The callback function is called on each VML mathematical function error if `VML_ERRMODE_CALLBACK` error mode is set (see [Table 7-11](#)).

Use the `vmlSetErrorCallBack()` function if you need to define your own callback function instead of default empty callback function.

The input structure for a callback function contains the following information

about the encountered error:

- the input value which caused an error
- location (array index) of this value
- the computed result value
- error code
- name of the function in which the error occurred.

You can insert your own error processing into the callback function. This may include correcting the passed result values in order to pass them back and resume computation. The standard error handler is called after the callback function only if it returns 0.

GetErrorCallBack

Gets the additional error handler callback function.

Fortran:

```
fun = vmlgeterrorcallback( )
```

C:

```
fun = vmlGetErrorCallBack( void );
```

Output Parameters

Fortran:

fun Address of the callback function.

C:

fun Pointer to the callback function.

ClearErrorCallBack

Deletes the additional error handler callback function and retrieves the former callback function.

Fortran:

```
oldcallback = vmlclearerrorcallback( )
```

C:

```
oldcallback = vmlClearErrorCallBack( void );
```

Output Parameters

Fortran:

oldcallback **INTEGER**. Address of the former callback function.

C:

oldcallback **int**. Pointer to the former callback function.

Vector Generators of Statistical Distributions

8

This chapter describes the part of Intel[®] MKL which is known as Vector Statistical Library (VSL) and is designed for the purpose of generating vectors of pseudorandom numbers.

VSL provides a set of pseudorandom number generator subroutines implementing basic continuous and discrete distributions. To speed up performance, all these subroutines were developed using the calls to the highly optimized *Basic Random Number Generators* (BRNGs) and the library of vector mathematical functions (VML, see [chapter 7](#)).

All VSL subroutines can be classified into three major categories:

- Pseudorandom number generators for different types of statistical distributions, for example, uniform, normal (Gaussian), binomial, etc. Detailed description of the generators can be found in [Pseudorandom Generators](#) section.
- Basic subroutines to handle random number streams: create, initialize, delete, copy, get the index of a basic generator. The description of these subroutines can be found in [Service Subroutines](#) section.
- Registration subroutines for basic pseudorandom generators and subroutines that obtain properties of the registered generators (see [Advanced Service Subroutines](#) section).

The last two categories will be referred to as service subroutines.

Conventions

In this discussion, a Random Number Generator (RNG) means a number-theoretic deterministic algorithm that generates number sequences, which can be interpreted as random samplings from a universal set with a given probability distribution function. Since random numbers are generated by a deterministic algorithm, they cannot be truly random and should be referred to as pseudorandom. The respective generators should be also called pseudorandom. However, in this chapter no specific differentiation is made between random and pseudorandom numbers, as well as between random and pseudorandom generators unless the context requires otherwise. Likewise, the terms *random number* and *variate*, *statistical distribution* and *probability distribution*, are not distinguished here either.

The choice of a number-theoretic algorithm A and initial conditions I identifies a unique sequence of random numbers, which is called a random stream. The pair $\langle A, I \rangle$ is referred to as the random stream state. In VSL a stream is identified by a *stream descriptor* represented as `TYPE (VSL_STREAM_STATE)` structure in FORTRAN interface, and `VSLStreamStatePtr` pointer in C interface.

All generators of nonuniform distributions, both discrete and continuous, are built on the basis of the uniform distribution generators, called Basic Random Number Generators (BRNGs). The pseudorandom numbers with nonuniform distribution are obtained through an appropriate transformation of the uniformly distributed pseudorandom numbers. The most common transformation techniques include the inverse Cumulative Distribution Function (CDF), acceptance/rejection method, and mixtures. For certain types of distribution, several generation methods are implemented.

VSL subroutines for pseudorandom number generation accept the stream descriptor and the distribution parameters as input and write the result in a vector of pseudorandom numbers with a given distribution. For a given statistical distribution, several generation methods can be used, which may differ in efficiency for particular ranges of input parameters. Consequently, the most efficient generators often use different methods for different ranges. To establish the generation method to be used in the subroutine, you

should specify the input parameter called the method number. Description of methods available for each generator can be found in [Pseudorandom Generators](#) section.

In the discussion that follow, the terms *multiprocessor system*, *computational node*, and *processor* refer to any configuration of the system with shared or distributed memory, or combination of the two. Specifically, a computational node, or a processor, refers to a computational unit capable of performing independent parallel computations (this may be either a physical processor, a cluster node, or a logical parallel process).

Mathematical Notation

The following notation is used throughout the text:

N	The set of natural numbers $N = \{1, 2, 3, \dots\}$.
Z	The set of integers $Z = \{\dots, -3, -2, -1, 0, 1, 2, 3, \dots\}$.
R	The set of real numbers.
$\lfloor a \rfloor$	The floor of a (the largest integer less than or equal to a).
\oplus or xor	Bitwise exclusive OR.
C_{α}^k or $\binom{\alpha}{k}$	Binomial coefficient or combination ($\alpha \in R$, $\alpha \geq 0$; $k \in N \cup \{0\}$). $C_{\alpha}^0 = 1$. For $\alpha \geq k$ binomial coefficient is defined as $C_{\alpha}^k = \frac{\alpha(\alpha-1) \dots (\alpha-k+1)}{k!}$. If $\alpha < k$, then $C_{\alpha}^k = 0$.
$\Phi(x)$	Cumulative Gaussian distribution function $\Phi(x) = \int_{-\infty}^x \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{y^2}{2}\right) dy$, defined over $-\infty < x < +\infty$. $\Phi(-\infty) = 0$, $\Phi(+\infty) = 1$.
LCG(a, c, m)	Linear Congruential Generator $x_{n+1} = (ax_n + c) \bmod m$, where a is called the <i>multiplier</i> , c is called the <i>increment</i> and m is called the <i>modulus</i> of the generator.

- MCG(a,m) Multiplicative Congruential Generator $x_{n+1} = (ax_n) \bmod m$ is a special case of Linear Congruential Generator, where the increment c is taken to be 0.
- GFSR(p,q) Generalized Feedback Shift Register Generator

$$x_n = x_{n-p} \oplus x_{n-q}.$$

Naming Conventions

The names of all VSL functions in FORTRAN are lowercase; names in C may contain both lowercase and uppercase letters.



NOTE. *This naming convention is followed in the function descriptions in the manual. Actual function names in the library may differ slightly (with respect to lower- and uppercase usage) and will be sufficient to meet the requirements of the supported compilers.*

The names of generator subroutines have the following structure:

v <type of result>rng<distribution> for FORTRAN-interface
 v <type of result>Rng<distribution> for C-interface

where v is the prefix of a VSL vector function, and the field <type of result> is either s , d , or i and specifies one of the following types:

s	REAL for FORTRAN-interface float for C-interface
d	DOUBLE PRECISION for FORTRAN-interface double for C-interface
i	INTEGER for FORTRAN-interface int for C-interface

Prefixes s and d apply to continuous distributions only, prefix i applies only to discrete case. The prefix rng indicates that the subroutine is a pseudorandom generator, and the <distribution> field specifies the type of statistical distribution.

Names of service subroutines follow the template below:

```
vsl<name> ,
```

where `vsl` is the prefix of a VSL service function. The field `<name>` contains a short function name. For a more detailed description of service subroutines refer to [Service Subroutines](#) and [Advanced Service Subroutines](#) sections.

Prototype of each generator subroutine implementing a given type of random number distribution fits the following structure:

```
<function name>( method, stream, n, r, [<distribution parameters>] ),
```

where

- `method` is the number specifying the method of generation. A detailed description of this parameter can be found in [Pseudorandom Generators](#) section.
- `stream` defines the random stream descriptor and must have a nonzero value. Random streams and their usage are discussed further in [Random Streams](#) and [Service Subroutines](#).
- `n` defines the number of pseudorandom values to be generated. If `n` is less than or equal to zero, no values are generated. Furthermore, if `n` is negative, an error condition is set.
- `r` defines the destination array for the generated numbers. The dimension of the array must be large enough to store at least `n` pseudorandom numbers.

Additional parameters included into `<distribution parameters>` field are individual for each generator subroutine and are described in detail in [Pseudorandom Generators](#) section.

To invoke a pseudorandom generator, use a call to the respective VSL subroutine. For example, to obtain a vector `r`, composed of `n` independent and identically distributed pseudorandom numbers with normal (Gaussian) distribution, that have the mean value `a` and standard deviation `sigma`, write the following:

for FORTRAN-interface

```
call vsrnggaussian( method, stream, n, r, a, sigma )
```

for C-interface

```
vsRngGaussian( method, stream, n, r, a, sigma )
```

Basic Pseudorandom Generators

Basic Random Number Generators (BRNGs) are the major and widely spread tool to obtain uniformly distributed pseudorandom numbers.

VSL provides a number of basic generators that differ in speed and quality: the 32-bit multiplicative congruential generator $MCG(1132489760, 2^{31} - 1)$ [[L'Ecuyer99](#)], the 32-bit generalized feedback shift register generator $GFSR(250, 103)$ [[Kirkpatrick81](#)], and the combined multiple recursive generator $MRG-32k3a$ [[L'Ecuyer99a](#)], as well as the 59-bit multiplicative congruential generator $MCG(13^{13}, 2^{59})$ and Wichmann-Hill generator (in fact, this is a set of 273 basic generators) from NAG Numerical Libraries [[NAG](#)]. Essentially, applicability of a basic generator to a given computational task is very difficult to estimate. To ensure more reliable results, basic generators are usually tested in a series of statistical tests prior to actual computation. Comparative performance analysis of the generators and testing results can be found in [VSLNotes](#).

Users may want to design and use their own basic generators. VSL provides means of registration of such user-designed generators through the steps described in [Advanced Service Subroutines](#) section.

For some basic generators, VSL provides two methods of creating independent random streams in multiprocessor computations, which are the leapfrog method and the block-splitting method. The properties of the generators designed for parallel computations are discussed in detail in [[Coddington94](#)].

For a more detailed description of the generator properties and testing results refer to [VSLNotes](#).

Random Streams

Several random streams may be used in one application for a number of reasons.

First, it may be necessary to supply random data to different computational nodes of a multiprocessor system. In this case, the following options are available:

- use an individual basic generator for each computational node, so that each random stream is filled from a different basic generator;

- use one basic generator for all computational nodes and generate several independent random streams using the leapfrog method or the block-splitting method;
- use combination of the two approaches, when one basic generator is used to generate independent streams for all nodes and each of the nodes in turn uses its own generator.

Another reason is related to the fact that many Monte Carlo simulations require additional randomization. A simple illustration is the necessity to assign random streams to different elements of the model or to run variance reduction methods [[Bratley87](#)].

In either case, the correlation between different random streams can affect reliability of the final result.

Data Types

Stream State. Random numbers can be generated by portions using the notion of a *stream state*, which is a structure created after a call to the stream creating subroutine. A stream state descriptor is used to access the structure:

FORTRAN

```
TYPE VSL_STREAM_STATE
    INTEGER*4 descriptor1
    INTEGER*4 descriptor2
END TYPE VSL_STREAM_STATE
```

C

```
typedef (void*) VSLStreamStatePtr;
```

See [Advanced Service Subroutines](#) for the format of the stream state structure for user-designed generators.

Service Subroutines

Stream handling comprises subroutines for creating, deleting, or copying the streams and getting the index of a basic generator.

[Table 8-1](#) lists all available service subroutines

Table 8-1 Service Subroutines

Subroutine	Short Description
NewStream	Creates and initializes a random stream.
NewStreamEx	Creates and initializes a random stream for the generators with multiple initial conditions.
DeleteStream	Deletes previously created stream.
CopyStream	Copies a stream to another stream.
CopyStreamState	Creates a copy of a random stream state.
LeapfrogStream	Initializes the stream of <i>k</i> -th computational node in a <i>nstreams</i> -node cluster by the leapfrog method.
SkipAheadStream	Initializes the stream by the block-splitting method.
GetStreamStateBrng	Obtains the index of the basic generator responsible for the generation of a given random stream.
GetNumRegBrng	Obtains the number of currently registered basic generators.



NOTE. *In the above table, the `vs1` prefix in the function names is omitted. In the function reference this prefix is always used in function prototypes and code examples.*

Most of the generator-based work comprises three basic steps:

1. Creating and initializing a stream ([NewStream](#), [NewStreamEx](#), [CopyStream](#), [CopyStreamState](#), [LeapfrogStream](#), [SkipAheadStream](#)).
2. Generating pseudorandom numbers with given distribution, see [Pseudorandom Generators](#).
3. Deleting the stream ([DeleteStream](#)).

Note that you can concurrently create multiple streams and obtain pseudorandom data from one or several generators by using the stream state. You must use the `DeleteStream` function to delete all the streams afterwards.

NewStream

Creates and initializes a random stream.

Fortran:

```
call vslnewstream( stream, brng, seed )
```

C:

```
vslNewStream( stream, brng, seed )
```

Discussion

For a basic generator with number `brng`, this function creates a new stream and initializes it with a 32-bit seed. The function is also applicable for generators with multiple initial conditions. See [VSLNotes](#) for a more detailed description of stream initialization for different basic generators.

Input Parameters

FORTRAN:

`brng` `INTEGER, INTENT(IN)`. Index of the basic generator to initialize the stream.

`seed` `INTEGER, INTENT(IN)`. Initial condition of the stream.

C:

`brng` `int`. Index of the basic generator to initialize the stream.

`seed` `unsigned int`. Initial condition of the stream.

Output Parameters

FORTRAN:

stream TYPE(VSL_STREAM_STATE),
INTENT(OUT). Stream state descriptor.

C:

stream VSLStreamStatePtr*. Pointer to the stream
state structure.

NewStreamEx

*Creates and initializes a random stream
for generators with multiple initial
conditions.*

Fortran:

```
call vslnewstreamex( stream, brng, n, params )
```

C:

```
vslNewStreamEx( stream, brng, n, params )
```

Discussion

This function provides an advanced tool to set the initial conditions for a basic generator if its input arguments imply several initialization parameters. This subroutine should not be used unless it is specially necessary. Whenever possible, use `vslNewStream`, which is analogous to `vslNewStreamEx` except that it takes only one 32-bit initial condition. In particular, `vslNewStreamEx` may be used to initialize the seed tables in Generalized Feedback Shift Register Generators (GFSRs). A more detailed description of this issue can be found in [VSLNotes](#).

Input Parameters

FORTRAN:

<i>brng</i>	<code>INTEGER, INTENT(IN)</code> . Index of the basic generator to initialize the stream.
<i>n</i>	<code>INTEGER, INTENT(IN)</code> . Number of initial conditions contained in <i>params</i> .
<i>params</i>	<code>INTEGER, INTENT(IN)</code> . Array of initial conditions necessary for the basic generator <i>brng</i> to initialize the stream.

C:

<i>brng</i>	<code>int</code> . Index of the basic generator to initialize the stream.
<i>n</i>	<code>int</code> . Number of initial conditions contained in <i>params</i> .
<i>params</i>	<code>const unsigned int[]</code> . Array of initial conditions necessary for the basic generator <i>brng</i> to initialize the stream.

Output Parameters

FORTRAN:

<i>stream</i>	<code>TYPE(VSL_STREAM_STATE), INTENT(OUT)</code> . Stream state descriptor.
---------------	---

C:

<i>stream</i>	<code>VSLStreamStatePtr*</code> . Pointer to the stream state structure.
---------------	--

DeleteStream

Deletes a random stream.

Fortran:

```
call vsldeletestream( stream )
```

C:

```
vs1DeleteStream( stream )
```

Discussion

This function deletes the random stream created by one of the initialization functions.

Input/Output Parameters

FORTRAN:

stream `TYPE(VSL_STREAM_STATE),`
 `INTENT(INOUT)`. Descriptor of the stream to
 be deleted; must have non-zero value.

C:

stream `VSLStreamStatePtr*`. Pointer to the stream
 structure; must have non-zero value.
 After the stream is successfully deleted, the
 stream pointer is set to `NULL`.

CopyStream

Creates a copy of a random stream.

Fortran:

```
call vslcopystream( newstream, srcstream )
```

C:

```
vs1CopyStream( newstream, srcstream )
```


Discussion

The function creates an exact copy of *srcstream* and stores its descriptor to *newstream*.

Input Parameters

FORTRAN:

srcstream `TYPE(VSL_STREAM_STATE),`
 `INTENT(IN)`. Descriptor of the stream to be copied.

C:

srcstream `VSLStreamStatePtr`. Pointer to the stream state structure to be copied.

Output Parameters

FORTRAN:

newstream `TYPE(VSL_STREAM_STATE),`
 `INTENT(OUT)`. Descriptor of the stream copy.

C:

newstream `VSLStreamStatePtr*`. Pointer to the copy of the stream state structure.

CopyStreamState

Creates a copy of a random stream state.

Fortran:

```
call vslcopystreamstate( deststream, srcstream )
```

C:

```
vslCopyStreamState( deststream, srcstream )
```

Discussion

The function copies a stream state from *srcstream* to the existing *deststream* stream. Both streams should be generated by the same basic generator. An error message is generated when the index of the BRNG that produced *deststream* stream differs from the index of the BRNG that generated *srcstream* stream.

Unlike [CopyStream](#) function, which creates a new stream and copies both the stream state and other data from *srcstream*, the function [CopyStreamState](#) copies only *srcstream* stream state data to the generated *deststream* stream.

Input Parameters

FORTRAN:

<i>srcstream</i>	<code>TYPE(VSL_STREAM_STATE), INTENT(IN)</code> . Descriptor of the stream with the state to be copied.
------------------	---

C:

<i>srcstream</i>	<code>VSLStreamStatePtr</code> . Pointer to the stream state structure from which the stream state is copied.
------------------	---

Output Parameters

FORTRAN:

<i>deststream</i>	<code>TYPE(VSL_STREAM_STATE), INTENT(IN)</code> . Descriptor of the destination stream where the state of <i>srcstream</i> stream is copied.
-------------------	--

C:

`deststream`

`VSLStreamStatePtr`. Pointer to the stream state structure where the stream state is copied.

LeapfrogStream

Initializes stream of k -th computational node in $nstreams$ -node cluster using the leapfrog method.

Fortran:

```
call vslleapfrogstream( stream, k, nstreams )
```

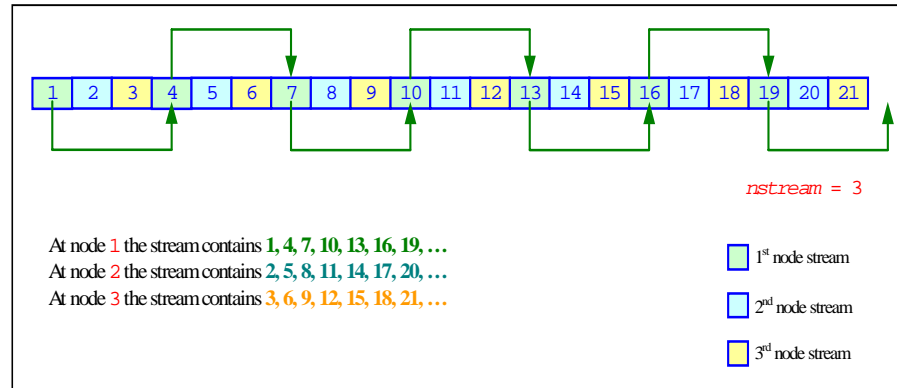
C:

```
vslLeapfrogstream( stream, k, nstreams )
```

Discussion

The function uses the leapfrog method (see [Figure 8-1](#)) to generate an independent random stream for the computational node k , $0 \leq k < nstreams$, where $nstreams$ is the largest number of computational nodes used.

Figure 8-1 Leapfrog Method



The following code examples illustrate the initialization of three independent streams using the leapfrog method:

Example 8-1 FORTRAN Code for Leapfrog Method

```

...
type(VSL_STREAM_STATE)stream1
type(VSL_STREAM_STATE)stream2
type(VSL_STREAM_STATE)stream3

! Creating 3 identical streams
call vslnewstream(stream1, VSL_BRNG_MCG31, 174)
call vslcopystream(stream2, stream1)
call vslcopystream(stream3, stream1)

! Leapfrogging the streams
call vslleapfrogstream(stream1, 0, 3)
call vslleapfrogstream(stream2, 1, 3)
call vslleapfrogstream(stream3, 2, 3)

! Generating random numbers
...
! Deleting the streams
call vsldeletestream(stream1)
call vsldeletestream(stream2)
call vsldeletestream(stream3)
...
    
```

Example 8-2 C Code for Leapfrog Method

```
...
VSLStreamStatePtr stream1;
VSLStreamStatePtr stream2;
VSLStreamStatePtr stream3;

/* Creating 3 identical streams */
vslNewStream(&stream1, VSL_BRNG_MCG31, 174);
vslCopyStream(&stream2, stream1);
vslCopyStream(&stream3, stream1);

/* Leapfrogging the streams */
vslLeapfrogStream(stream1, 0, 3);
vslLeapfrogStream(stream2, 1, 3);
vslLeapfrogStream(stream3, 2, 3);

/* Generating random numbers */
...
/* Deleting the streams */
vslDeleteStream(&stream1);
vslDeleteStream(&stream2);
vslDeleteStream(&stream3);
...
```

Input Parameters

FORTRAN:

<i>stream</i>	<code>TYPE(VSL_STREAM_STATE),</code> <code>INTENT(IN)</code> . Descriptor of the stream to which the leapfrog method is applied.
<i>k</i>	<code>INTEGER, INTENT(IN)</code> . Index of the computational node, or stream number.
<i>nstreams</i>	<code>INTEGER, INTENT(IN)</code> . Largest number of computational nodes, or number of independent streams.

C:

<i>stream</i>	<code>VSLStreamStatePtr</code> . Pointer to the stream state structure to which the leapfrog method is applied.
<i>k</i>	<code>int</code> . Index of the computational node, or stream number.
<i>nstreams</i>	<code>int</code> . Largest number of computational nodes, or number of independent streams.

SkipAheadStream

Initializes a stream using the block-splitting method.

Fortran:


```
call vslskipaheadstream( stream, nskip )
```

C:

```
vslSkipAheadStream( stream, nskip )
```

Discussion

This function initializes an independent random stream of a given computational node through the block-splitting method (see [Figure 8-2](#)). The maximum number of computational nodes is unlimited. The largest number of elements skipped in a given stream is *nskip*.

Figure 8-2 Block-Splitting Method

The following code examples illustrate how to initialize three independent streams using the `SkipAheadStream` function:

Example 8-3 FORTRAN Code for Block-Splitting Method

```

...
TYPE(VSL_STREAM_STATE)stream1
TYPE(VSL_STREAM_STATE)stream2
TYPE(VSL_STREAM_STATE)stream3

! Creating the 1st stream
call vslnewstream(stream1, VSL_BRNG_MCG31, 174)

! Skipping ahead by 7 elements the 2nd stream
call vslcopystream(stream2, stream1);
call vslskipaheadstream(stream2, 7);

! Skipping ahead by 7 elements the 3rd stream
call vslcopystream(stream3, stream2);
call vslskipaheadstream(stream3, 7);

! Generating random numbers
...
! Deleting the streams
call vsldeletestream(stream1)
call vsldeletestream(stream2)
call vsldeletestream(stream3)
...

```

Example 8-4 C Code for Block-Splitting Method

```

VSLStreamStatePtr stream1;
VSLStreamStatePtr stream2;
VSLStreamStatePtr stream3;

/* Creating the 1st stream */
vslNewStream(&stream1, VSL_BRNG_MCG31, 174);

/* Skipping ahead by 7 elements the 2nd stream */
vslCopyStream(&stream2, stream1);
vslSkipAheadStream(stream2, 7);

/* Skipping ahead by 7 elements the 3rd stream */
vslCopyStream(&stream3, stream2);
vslSkipAheadStream(stream3, 7);

/* Generating random numbers */
...
/* Deleting the streams */
vslDeleteStream(&stream1);
vslDeleteStream(&stream2);
vslDeleteStream(&stream3);
...

```

Input Parameters

FORTRAN:

<i>stream</i>	TYPE(VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream to which the block-splitting method is applied.
<i>nskip</i>	INTEGER, INTENT(IN). Number of skipped elements.

C:

<i>stream</i>	VSLStreamStatePtr. Pointer to the stream state structure to which the block-splitting method is applied.
<i>nskip</i>	int. Number of skipped elements.

GetStreamStateBrng

Returns index of a basic generator used for generation of a given random stream.

Fortran:

```
brng = vslgetstreamstatebrng( stream )
```

C:

```
brng = vslGetStreamStateBrng( stream )
```

Discussion

This function retrieves the index of a basic generator used for generation of a given random stream.

Input Parameters

FORTRAN:

stream `TYPE(VSL_STREAM_STATE), INTENT(IN)`.
Descriptor of the stream state.

C:

stream `VSLStreamStatePtr`. Pointer to the stream
state structure.

Output Parameters

FORTRAN:

brng `INTEGER`. Index of the basic generator assigned
for the generation of *stream* ; negative in case
of an error.

C:

*brng**int*. Index of the basic generator assigned for the generation of *stream* ; negative in case of an error.

GetNumRegBrng

Obtains the number of currently registered basic generators.

Fortran:*nregbrng* = vslgetnumregbrngs()**C:***nregbrng* = vslGetNumRegBrngs(void)

Discussion

This function obtains the number of currently registered basic generators. Whenever the user registers a user-defined basic generator, the number of registered basic generators is incremented. The maximum number of basic generators that can be registered is determined by `VSL_MAX_REG_BRNGS` parameter.

Output Parameters

FORTRAN:

*nregbrngs***INTEGER**. The number of basic generators registered at the moment of the function call.

C:

*nregbrngs***int**. The number of basic generators registered at the moment of the function call.

Pseudorandom Generators

This section contains description of VSL subroutines for generating random numbers with different types of distribution. Each function group is introduced by the type of underlying distribution and contains a short description of its functionality, as well as specifications of the call sequence for both FORTRAN and C-interface and the explanation of input and output parameters.

[Table 8-2](#) and [Table 8-3](#) list the pseudorandom number generator subroutines, together with used data types and output distributions.

Table 8-2 Continuous Distribution Generators

Type of Distribution	Data Types	Description
Uniform	s, d	Uniform continuous distribution on the interval (a,b) .
Gaussian	s, d	Normal (Gaussian) distribution.
Exponential	s, d	Exponential distribution.
Laplace	s, d	Laplace distribution (double exponential distribution).
Weibull	s, d	Weibull distribution.
Cauchy	s, d	Cauchy distribution.
Rayleigh	s, d	Rayleigh distribution.
Lognormal	s, d	Lognormal distribution.
Gumbel	s, d	Gumbel (extreme value) distribution.

Table 8-3 Discrete Distribution Generators

Type of Distribution	Data Types	Description
Uniform	i	Uniform discrete distribution on the interval $[a,b)$.
UniformBits	i	Generator of integer random values with uniform bit distribution.
Bernoulli	i	Bernoulli distribution.
Geometric	i	Geometric distribution.
Binomial	i	Binomial distribution.
Hypergeometric	i	Hypergeometric distribution.

Table 8-3 Discrete Distribution Generators (continued)

Type of Distribution	Data Types	Description
Poisson	i	Poisson distribution.
NegBinomial	i	Negative binomial distribution, or Pascal distribution.

Continuous Distributions

This section describes routines for generating pseudorandom numbers with continuous distribution.

Uniform

Generates pseudorandom numbers with uniform distribution.

Fortran:

```
call vsrnguniform( method, stream, n, r, a, b )
call vdrnguniform( method, stream, n, r, a, b )
```

C:

```
vsRngUniform( method, stream, n, r, a, b )
vdRngUniform( method, stream, n, r, a, b )
```

Discussion

This function generates pseudorandom numbers uniformly distributed over the interval (a, b) , where a, b are the left and right bounds of the interval, respectively, and $a, b \in R; a < b$.

The probability density function is given by:

$$f_{a, b}(x) = \begin{cases} \frac{1}{b-a}, & x \in (a, b) \\ 0, & x \notin (a, b) \end{cases}, \quad -\infty < x < +\infty.$$

The cumulative distribution function is as follows:

$$F_{a, b}(x) = \begin{cases} 0, & x < a \\ \frac{x-a}{b-a}, & a \leq x < b, \quad -\infty < x < +\infty. \\ 1, & x \geq b \end{cases}$$

Input Parameters

FORTRAN:

<i>method</i>	INTEGER, INTENT(IN). Generation method; dummy and set to 0 in case of uniform distribution.
<i>stream</i>	TYPE (VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream state structure.
<i>n</i>	INTEGER, INTENT(IN). Number of random values to be generated.
<i>a</i>	REAL, INTENT(IN) for vsrnguniform. DOUBLE PRECISION, INTENT(IN) for vdrnguniform. Left bound <i>a</i> .
<i>b</i>	REAL, INTENT(IN) for vsrnguniform. DOUBLE PRECISION, INTENT(IN) for vdrnguniform. Right bound <i>b</i> .

C:

<i>method</i>	int. Generation method; dummy and set to 0 in case of uniform distribution.
<i>stream</i>	VSLStreamStatePtr. Pointer to the stream state structure.

<i>n</i>	<code>int</code> . Number of random values to be generated.
<i>a</i>	<code>float</code> for <code>vsRngUniform</code> . <code>double</code> for <code>vdRngUniform</code> . Left bound <i>a</i> .
<i>b</i>	<code>float</code> for <code>vsRngUniform</code> . <code>double</code> for <code>vdRngUniform</code> . Right bound <i>b</i> .

Output Parameters

FORTRAN:

<i>r</i>	<code>REAL, INTENT(OUT)</code> for <code>vsrnguniform</code> . <code>DOUBLE PRECISION, INTENT(OUT)</code> for <code>vdRngUniform</code> . Vector of <i>n</i> pseudorandom numbers uniformly distributed over the interval (<i>a</i> , <i>b</i>).
----------	--

C:

<i>r</i>	<code>float*</code> for <code>vsRngUniform</code> . <code>double*</code> for <code>vdRngUniform</code> . Vector of <i>n</i> pseudorandom numbers uniformly distributed over the interval (<i>a</i> , <i>b</i>).
----------	---

Gaussian

Generates normally distributed pseudorandom numbers.

Fortran:

```
call vsrnggaussian( method, stream, n, r, a, sigma )
```

```
call vdrnggaussian( method, stream, n, r, a, sigma )
```

C:

```
vsRngGaussian( method, stream, n, r, a, sigma )
vdRngGaussian( method, stream, n, r, a, sigma )
```

Discussion

This function generates pseudorandom numbers with normal (Gaussian) distribution with mean value a and standard deviation σ , where

$a, \sigma \in R; \sigma > 0$.

The probability density function is given by:

$$f_{a, \sigma}(x) = \frac{1}{\sqrt{2\pi\sigma}} \exp\left(-\frac{(x-a)^2}{2\sigma^2}\right), \quad -\infty < x < +\infty.$$

The cumulative distribution function is as follows:

$$F_{a, \sigma}(x) = \int_{-\infty}^x \frac{1}{\sqrt{2\pi\sigma}} \exp\left(-\frac{(y-a)^2}{2\sigma^2}\right) dy, \quad -\infty < x < +\infty.$$

The cumulative distribution function $F_{a, \sigma}(x)$ can be expressed in terms of standard normal distribution $\Phi(x)$ as

$$F_{a, \sigma}(x) = \Phi((x-a)/\sigma).$$

Input Parameters

FORTTRAN:

<i>method</i>	INTEGER, INTENT(IN). Generation method.
<i>stream</i>	TYPE (VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream state structure.
<i>n</i>	INTEGER, INTENT(IN). Number of random values to be generated.

a REAL, INTENT(IN) for `vsrnggaussian`.
 DOUBLE PRECISION, INTENT(IN) for
`vdrnggaussian`.

Mean value *a*.

sigma REAL, INTENT(IN) for `vsrnggaussian`.
 DOUBLE PRECISION, INTENT(IN) for
`vdrnggaussian`.

Standard deviation σ .

C:

method int. Generation method.

stream VSLStreamStatePtr. Pointer to the stream
 state structure.

n int. Number of random values to be
 generated.

a float for `vsRngGaussian`.
 double for `vdRngGaussian`.

Mean value *a*.

sigma float for `vsRngGaussian`.
 double for `vdRngGaussian`.

Standard deviation σ .

Output Parameters

FORTRAN:

r REAL, INTENT(OUT) for `vsrnggaussian`.
 DOUBLE PRECISION, INTENT(OUT) for
`vdrnggaussian`.

Vector of *n* normally distributed
 pseudorandom numbers.

C:

`r` `float*` for `vsRngGaussian`.
`double*` for `vdRngGaussian`.
 Vector of `n` normally distributed
 pseudorandom numbers.

Exponential

*Generates exponentially distributed
 pseudorandom numbers.*

Fortran:

call `vsrngexponential(method, stream, n, r, a, beta)`
 call `vdRngexponential(method, stream, n, r, a, beta)`

C:

`vsRngExponential(method, stream, n, r, a, beta)`
`vdRngExponential(method, stream, n, r, a, beta)`

Discussion

This function generates pseudorandom numbers with exponential distribution that has the displacement a and scalefactor β , where $a, \beta \in R; \beta > 0$.

The probability density function is given by:

$$f_{a, \beta}(x) = \begin{cases} \frac{1}{\beta} \exp(-(x-a)/\beta), & x \geq a \\ 0, & x < a \end{cases}, \quad -\infty < x < +\infty.$$

The cumulative distribution function is as follows:

$$F_{a, \beta}(x) = \begin{cases} 1 - \exp(-(x-a)/\beta), & x \geq a \\ 0, & x < a \end{cases}, \quad -\infty < x < +\infty.$$

Input Parameters

FORTRAN:

<i>method</i>	INTEGER, INTENT(IN). Generation method.
<i>stream</i>	TYPE (VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream state structure.
<i>n</i>	INTEGER, INTENT(IN). Number of random values to be generated.
<i>a</i>	REAL, INTENT(IN) for vsrngexponential. DOUBLE PRECISION, INTENT(IN) for vdrngexponential. Displacement <i>a</i> .
<i>beta</i>	REAL, INTENT(IN) for vsrngexponential. DOUBLE PRECISION, INTENT(IN) for vdrngexponential. Scalefactor β .

C:

<i>method</i>	int. Generation method.
<i>stream</i>	VSLStreamStatePtr. Pointer to the stream state structure.
<i>n</i>	int. Number of random values to be generated.
<i>a</i>	float for vsRngExponential. double for vdRngExponential. Displacement <i>a</i> .

beta `float` for `vsRngExponential`.
 `double` for `vdRngExponential`.
 Scalefactor β .

Output Parameters

FORTRAN:

r `REAL, INTENT(OUT)` for
 `vsrngexponential`.
 `DOUBLE PRECISION, INTENT(OUT)` for
 `vdrngexponential`.
 Vector of *n* exponentially distributed
 pseudorandom numbers.

C:

r `float*` for `vsRngExponential`.
 `double*` for `vdRngExponential`.
 Vector of *n* exponentially distributed
 pseudorandom numbers.

Laplace

*Generates pseudorandom numbers with
 Laplace distribution.*

Fortran:

```
call vsrnglaplace( method, stream, n, r, a, beta )
call vdrnglaplace( method, stream, n, r, a, beta )
```

C:

```
vsRngLaplace( method, stream, n, r, a, beta )
vdRngLaplace( method, stream, n, r, a, beta )
```

Discussion

This function generates pseudorandom numbers with Laplace distribution with mean value (or average) a and scalefactor β , where

$a, \beta \in \mathbf{R}; \beta > 0$. The scalefactor value determines the standard deviation as

$$\sigma = \beta\sqrt{2}.$$

The probability density function is given by:

$$f_{a, \beta}(x) = \frac{1}{\sqrt{2}\beta} \exp\left(-\frac{|x-a|}{\beta}\right), \quad -\infty < x < +\infty.$$

The cumulative distribution function is as follows:

$$F_{a, \beta}(x) = \begin{cases} \frac{1}{2} \exp\left(-\frac{|x-a|}{\beta}\right), & x < a \\ 1 - \frac{1}{2} \exp\left(-\frac{|x-a|}{\beta}\right), & x \geq a \end{cases}, \quad -\infty < x < +\infty.$$

Input Parameters

FORTRAN:

<i>method</i>	INTEGER, INTENT(IN). Generation method.
<i>stream</i>	TYPE (VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream state structure.
<i>n</i>	INTEGER, INTENT(IN). Number of random values to be generated.
<i>a</i>	REAL, INTENT(IN) for <code>vsrnglaplace</code> . DOUBLE PRECISION, INTENT(IN) for <code>vdrnglaplace</code> . Mean value <i>a</i> .

beta REAL, INTENT(IN) for `vsrnglaplace`.
DOUBLE PRECISION, INTENT(IN) for `vdrnglaplace`.
Scalefactor β .

C:

method int. Generation method.
stream VSLStreamStatePtr. Pointer to the stream state descriptor.

n int. Number of random values to be generated.

a float for `vsRngLaplace`.
double for `vdRngLaplace`.
Mean value *a*.

beta float for `vsRngLaplace`.
double for `vdRngLaplace`.
Scalefactor β .

Output Parameters

FORTRAN:

r REAL, INTENT(OUT) for `vsrnglaplace`.
DOUBLE PRECISION, INTENT(OUT) for `vdrnglaplace`.
Vector of *n* Laplace distributed pseudorandom numbers.

C:

r float* for `vsRngLaplace`.
double* for `vdRngLaplace`.
Vector of *n* Laplace distributed pseudorandom numbers.

Weibull

Generates Weibull distributed pseudorandom numbers.

Fortran:

```
call vsrngweibull( method, stream, n, r, alpha, a, beta )
call vdrngweibull( method, stream, n, r, alpha, a, beta )
```

C:

```
vsRngWeibull( method, stream, n, r, alpha, a, beta )
vdRngWeibull( method, stream, n, r, alpha, a, beta )
```

Discussion

This function generates Weibull distributed pseudorandom numbers with displacement a , scalefactor β , and shape α , where $\alpha, \beta, a \in R$; $\alpha > 0$; $\beta > 0$.

The probability density function is given by:

$$f_{a, \alpha, \beta}(x) = \begin{cases} \frac{\alpha}{\beta^\alpha} (x-a)^{\alpha-1} \exp\left(-\left(\frac{x-a}{\beta}\right)^\alpha\right), & x \geq a \\ 0, & x < a \end{cases}$$

The cumulative distribution function is as follows:

$$F_{a, \alpha, \beta}(x) = \begin{cases} 1 - \exp\left(-\left(\frac{x-a}{\beta}\right)^\alpha\right), & x \geq a \\ 0, & x < a \end{cases}, \quad -\infty < x < +\infty.$$

Input Parameters

FORTRAN:

<i>method</i>	INTEGER, INTENT(IN). Generation method.
<i>stream</i>	TYPE (VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream state structure.
<i>n</i>	INTEGER, INTENT(IN). Number of random values to be generated.
<i>alpha</i>	REAL, INTENT(IN) for vsrngweibull. DOUBLE PRECISION, INTENT(IN) for vdrngweibull. Shape α .
<i>a</i>	REAL, INTENT(IN) for vsrngweibull. DOUBLE PRECISION, INTENT(IN) for vdrngweibull. Displacement <i>a</i> .
<i>beta</i>	REAL, INTENT(IN) for vsrngweibull. DOUBLE PRECISION, INTENT(IN) for vdrngweibull. Scalefactor β .

C:

<i>method</i>	int. Generation method.
<i>stream</i>	VSLStreamStatePtr. Pointer to the stream state structure.
<i>n</i>	int. Number of random values to be generated.
<i>alpha</i>	float for vsRngWeibull. double for vdRngWeibull. Shape α .

`a` float for `vsRngWeibull`.
double for `vdRngWeibull`.
Displacement `a`.

`beta` float for `vsRngWeibull`.
double for `vdRngWeibull`.
Scalefactor β .

Output Parameters

FORTRAN:

`r` REAL, INTENT(OUT) for `vsrngweibull`.
DOUBLE PRECISION, INTENT(OUT) for
`vdRngWeibull`.
Vector of `n` Weibull distributed pseudorandom
numbers.

C:

`r` float* for `vsRngWeibull`.
double* for `vdRngWeibull`.
Vector of `n` Weibull distributed pseudorandom
numbers.

Cauchy

*Generates Cauchy distributed
pseudorandom values.*

Fortran:

```
call vsrngcauchy( method, stream, n, r, a, beta )
call vdrngcauchy( method, stream, n, r, a, beta )
```


C:

```
vsRngCauchy( method, stream, n, r, a, beta )
vdRngCauchy( method, stream, n, r, a, beta )
```

Discussion

This function generates Cauchy distributed pseudorandom numbers with displacement a and scalefactor β , where $a, \beta \in R; \beta > 0$.

The probability density function is given by:

$$f_{a, \beta}(x) = \frac{1}{\pi\beta\left(1 + \left(\frac{x-a}{\beta}\right)^2\right)}, \quad -\infty < x < +\infty.$$

The cumulative distribution function is as follows:

$$F_{a, \beta}(x) = \frac{1}{2} + \frac{1}{\pi} \arctan\left(\frac{x-a}{\beta}\right), \quad -\infty < x < +\infty.$$

Input Parameters

FORTRAN:

<i>method</i>	INTEGER, INTENT(IN). Generation method.
<i>stream</i>	TYPE (VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream state structure.
<i>n</i>	INTEGER, INTENT(IN). Number of random values to be generated.
<i>a</i>	REAL, INTENT(IN) for vsrngcauchy. DOUBLE PRECISION, INTENT(IN) for vdRngCauchy. Displacement a .

<i>beta</i>	REAL, INTENT(IN) for <i>vsrngcauchy</i> . DOUBLE PRECISION, INTENT(IN) for <i>vdrngcauchy</i> . Scalefactor β .
C:	
<i>method</i>	int. Generation method.
<i>stream</i>	VSLStreamStatePtr. Pointer to the stream state structure.
<i>n</i>	int. Number of random values to be generated.
<i>a</i>	float for <i>vsRngCauchy</i> . double for <i>vdRngCauchy</i> . Displacement <i>a</i> .
<i>beta</i>	float for <i>vsRngCauchy</i> . double for <i>vdRngCauchy</i> . Scalefactor β .

Output Parameters

FORTRAN:

<i>r</i>	REAL, INTENT(OUT) for <i>vsrngcauchy</i> . DOUBLE PRECISION, INTENT(OUT) for <i>vdrngcauchy</i> . Vector of <i>n</i> Cauchy distributed pseudorandom numbers.
----------	---

C:

<i>r</i>	float* for <i>vsRngCauchy</i> . double* for <i>vdRngCauchy</i> . Vector of <i>n</i> Cauchy distributed pseudorandom numbers.
----------	--

Rayleigh

Generates Rayleigh distributed pseudorandom values.

Fortran:

```
call vsrngrayleigh( method, stream, n, r, a, beta )
call vdrnggrayleigh( method, stream, n, r, a, beta )
```

C:

```
vsRngRayleigh( method, stream, n, r, a, beta )
vdRngRayleigh( method, stream, n, r, a, beta )
```

Discussion

This function generates Rayleigh distributed pseudorandom numbers with displacement a and scalefactor β , where $a, \beta \in R; \beta > 0$.

Rayleigh distribution is a special case of Weibull distribution, where the shape parameter $\alpha = 2$.

The probability density function is given by:

$$f_{a, \beta}(x) = \begin{cases} \frac{2(x-a)}{\beta^2} \exp\left(-\frac{(x-a)^2}{\beta^2}\right), & x \geq a \\ 0, & x < a \end{cases}, \quad -\infty < x < +\infty.$$

The cumulative distribution function is as follows:

$$F_{a, \beta}(x) = \begin{cases} 1 - \exp\left(-\frac{(x-a)^2}{\beta^2}\right), & x \geq a \\ 0, & x < a \end{cases}, \quad -\infty < x < +\infty.$$

Input Parameters

FORTRAN:

<i>method</i>	INTEGER, INTENT(IN). Generation method.
<i>stream</i>	TYPE (VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream state structure.
<i>n</i>	INTEGER, INTENT(IN). Number of random values to be generated.
<i>a</i>	REAL, INTENT(IN) for vsrngrayleigh. DOUBLE PRECISION, INTENT(IN) for vdrnggrayleigh. Displacement <i>a</i> .
<i>beta</i>	REAL, INTENT(IN) for vsrngrayleigh. DOUBLE PRECISION, INTENT(IN) for vdrnggrayleigh. Scalefactor β .

C:

<i>method</i>	int. Generation method.
<i>stream</i>	VSLStreamStatePtr. Pointer to the stream state structure.
<i>n</i>	int. Number of random values to be generated.
<i>a</i>	float for vsRngRayleigh. double for vdRngRayleigh. Displacement <i>a</i> .
<i>beta</i>	float for vsRngRayleigh. double for vdRngRayleigh. Scalefactor β .

Output Parameters

FORTRAN:

r REAL, INTENT(OUT) for vsrngrayleigh.
DOUBLE PRECISION, INTENT(OUT) for
vdrnggrayleigh.
Vector of *n* Rayleigh distributed
pseudorandom numbers.

C:

r float* for vsRngRayleigh.
double* for vdRngRayleigh.
Vector of *n* Rayleigh distributed
pseudorandom numbers.

Lognormal

*Generates lognormally distributed
pseudorandom numbers.*

Fortran:

```
call vsrnglognormal( method, stream, n, r, a, sigma, b,  
beta )  
call vdrnglognormal( method, stream, n, r, a, sigma, b,  
beta )
```

C:

```
vsRngLognormal( method, stream, n, r, a, sigma, b, beta )  
vdRngLognormal( method, stream, n, r, a, sigma, b, beta )
```

Discussion

This function generates lognormally distributed pseudorandom numbers with average of distribution a and standard deviation σ of subject normal distribution, displacement b , and scalefactor β , where

$$a, \sigma, b, \beta \in R; \sigma > 0; \beta > 0.$$

The probability density function is given by:

$$f_{a, \sigma, b, \beta}(x) = \begin{cases} \frac{1}{\sigma(x-b)\sqrt{2\pi}} \exp\left(-\frac{[\ln((x-b)/\beta) - a]^2}{2\sigma^2}\right), & x > b \\ 0, & x \leq b \end{cases}$$

The cumulative distribution function is as follows:

$$F_{a, \sigma, b, \beta}(x) = \begin{cases} \Phi((\ln((x-b)/\beta) - a)/\sigma), & x > b \\ 0, & x \leq b \end{cases}$$

Input Parameters

FORTRAN:

<i>method</i>	INTEGER, INTENT(IN). Generation method.
<i>stream</i>	TYPE (VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream state structure.
<i>n</i>	INTEGER, INTENT(IN). Number of random values to be generated.
<i>a</i>	REAL, INTENT(IN) for vsrnglognormal. DOUBLE PRECISION, INTENT(IN) for vdrnglognormal. Average a of the subject normal distribution.

<i>sigma</i>	<p>REAL, INTENT(IN) for vsrnglognormal. DOUBLE PRECISION, INTENT(IN) for vdrnglognormal.</p> <p>Standard deviation σ of the subject normal distribution.</p>
<i>b</i>	<p>REAL, INTENT(IN) for vsrnglognormal. DOUBLE PRECISION, INTENT(IN) for vdrnglognormal.</p> <p>Displacement <i>b</i>.</p>
<i>beta</i>	<p>REAL, INTENT(IN) for vsrnglognormal. DOUBLE PRECISION, INTENT(IN) for vdrnglognormal.</p> <p>Scalefactor value β.</p>
C:	
<i>method</i>	<i>int</i> . Generation method.
<i>stream</i>	VSLStreamStatePtr. Pointer to the stream state structure.
<i>n</i>	<i>int</i> . Number of random values to be generated.
<i>a</i>	<p>float for vsRngLognormal. double for vdRngLognormal.</p> <p>Average <i>a</i> of the subject normal distribution.</p>
<i>sigma</i>	<p>float for vsRngLognormal. double for vdRngLognormal.</p> <p>Standard deviation σ of the subject normal distribution.</p>
<i>b</i>	<p>float for vsRngLognormal. double for vdRngLognormal.</p> <p>Displacement <i>b</i>.</p>

beta float for vsRngLognormal.
double for vdRngLognormal.
Scalefactor value β .

Output Parameters

FORTRAN:

r REAL, INTENT(OUT) for vsrnglognormal.
DOUBLE PRECISION, INTENT(OUT) for
vdrnglognormal.
Vector of *n* lognormally distributed
pseudorandom numbers.

C:

r float* for vsRngLognormal.
double* for vdRngLognormal.
Vector of *n* lognormally distributed
pseudorandom numbers.

Gumbel

*Generates Gumbel distributed
pseudorandom values.*

Fortran:

call vsrnggumbel(*method, stream, n, r, a, beta*)
call vdrnggumbel(*method, stream, n, r, a, beta*)

C:

vsRngGumbel(*method, stream, n, r, a, beta*)
vdRngGumbel(*method, stream, n, r, a, beta*)

Discussion

This function generates Gumbel distributed pseudorandom numbers with displacement a and scalefactor β , where $a, \beta \in R; \beta > 0$.

The probability density function is given by:

$$f_{a, \beta}(x) = \frac{1}{\beta} \exp\left(\frac{x-a}{\beta}\right) \exp(-\exp((x-a)/\beta)), \quad -\infty < x < +\infty$$

The cumulative distribution function is as follows:

$$F_{a, \beta}(x) = 1 - \exp(-\exp((x-a)/\beta)), \quad -\infty < x < +\infty.$$

Input Parameters

FORTRAN:

<i>method</i>	INTEGER, INTENT(IN). Generation method.
<i>stream</i>	TYPE (VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream state structure.
<i>n</i>	INTEGER, INTENT(IN). Number of random values to be generated.
<i>a</i>	REAL, INTENT(IN) for vsrnggumbel. DOUBLE PRECISION, INTENT(IN) for vdrnggumbel. Displacement a .
<i>beta</i>	REAL, INTENT(IN) for vsrnggumbel. DOUBLE PRECISION, INTENT(IN) for vdrnggumbel. Scalefactor β .

C:

<i>method</i>	int. Generation method.
---------------	-------------------------

<i>stream</i>	<code>VSLStreamStatePtr</code> . Pointer to the stream state structure.
<i>n</i>	<code>int</code> . Number of random values to be generated.
<i>a</i>	<code>float</code> for <code>vsRngGumbel</code> . <code>double</code> for <code>vdRngGumbel</code> . Displacement <i>a</i> .
<i>beta</i>	<code>float</code> for <code>vsRngGumbel</code> . <code>double</code> for <code>vdRngGumbel</code> . Scalefactor β .

Output Parameters

FORTRAN:

<i>r</i>	<code>REAL, INTENT(OUT)</code> for <code>vsrnggumbel</code> . <code>DOUBLE PRECISION, INTENT(OUT)</code> for <code>vdrrnggumbel</code> . Vector of <i>n</i> pseudorandom values with Gumbel distribution.
----------	---

C:

<i>r</i>	<code>float*</code> for <code>vsRngGumbel</code> . <code>double*</code> for <code>vdRngGumbel</code> . Vector of <i>n</i> pseudorandom values with Gumbel distribution.
----------	---

Discrete Distributions

This section describes routines for generating pseudorandom numbers with discrete distribution.

<i>stream</i>	<code>TYPE (VSL_STREAM_STATE), INTENT(IN)</code> . Descriptor of the stream state structure.
<i>n</i>	<code>INTEGER, INTENT(IN)</code> . Number of random values to be generated.
<i>a</i>	<code>INTEGER, INTENT(IN)</code> . Left interval bound <i>a</i> .
<i>b</i>	<code>INTEGER, INTENT(IN)</code> . Right interval bound <i>b</i> .
C:	
<i>method</i>	<code>int</code> . Generation method.
<i>stream</i>	<code>VSLStreamStatePtr</code> . Pointer to the stream state structure.
<i>n</i>	<code>int</code> . Number of random values to be generated.
<i>a</i>	<code>int</code> . Left interval bound <i>a</i> .
<i>b</i>	<code>int</code> . Right interval bound <i>b</i> .

Output Parameters

FORTRAN:

<i>r</i>	<code>INTEGER, INTENT(OUT)</code> . Vector of <i>n</i> pseudorandom values uniformly distributed over the interval [<i>a</i> , <i>b</i>).
----------	---

C:

<i>r</i>	<code>int*</code> . Vector of <i>n</i> pseudorandom values uniformly distributed over the interval [<i>a</i> , <i>b</i>).
----------	---

UniformBits

Generates integer random values with uniform bit distribution.

Fortran:

```
call virnguniformbits( method, stream, n, r )
```

C:

```
viRngUniformBits( method, stream, n, r )
```

Discussion

This function generates integer random values with uniform bit distribution. The generators of uniformly distributed numbers can be represented as recurrence relations over integer values in modular arithmetic. Apparently, each integer can be treated as a vector of several bits. In a truly random generator, these bits are random, while in pseudorandom generators this randomness can be violated. For example, a well known drawback of linear congruential generators is that lower bits are less random than higher bits (for example, see [Knuth81]). For this reason, care should be taken when using this function. Typically, in a 32-bit LCG only 24 higher bits of an integer value can be considered truly random. See [VSLNotes](#) for details.

Input Parameters

FORTTRAN:

<i>method</i>	INTEGER, INTENT(IN). Generation method. A dummy argument in <code>virnguniformbits</code> . Should be zero.
<i>stream</i>	TYPE (VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream state structure.
<i>n</i>	INTEGER, INTENT(IN). Number of random values to be generated.

C:

method *int*. Generation method. A dummy argument in *viRngUniformBits*. Should be zero.

stream *VSLStreamStatePtr*. Pointer to the stream state structure.

n *int*. Number of random values to be generated.

Output Parameters

FORTRAN:

r *INTEGER, INTENT(OUT)*. Vector of *n* pseudorandom integer numbers. If the *stream* was generated by a 64 or a 128-bit generator, each integer value is represented by two or four elements of *r* respectively. The number of bytes occupied by each integer is contained in the field *wordsize* of the structure *VSL_BRNG_PROPERTIES*. The total number of bits that are actually used to store the value are contained in the field *nbits* of the same structure. See [Advanced Service Subroutines](#) for a more detailed discussion of *VSL_BRNG_PROPERTIES*.

C:

`r` `unsigned int*`. Vector of `n` pseudorandom integer numbers. If the `stream` was generated by a 64 or a 128-bit generator, each integer value is represented by two or four elements of `r` respectively. The number of bytes occupied by each integer is contained in the field `WordSize` of the structure `VSLBrngProperties`. The total number of bits that are actually used to store the value are contained in the field `NBits` of the same structure. See [Advanced Service Subroutines](#) for a more detailed discussion of `VSLBrngProperties`.

Bernoulli

Generates Bernoulli distributed pseudorandom values.

Fortran:

```
call virngbernoulli( method, stream, n, r, p )
```

C:

```
viRngBernoulli( method, stream, n, r, p )
```

Discussion

This function generates Bernoulli distributed pseudorandom numbers with probability p of a single trial success, where

$$p \in R; 0 \leq p \leq 1 .$$

A variate is called Bernoulli distributed, if after a trial it is equal to 1 with probability of success p , and to 0 with probability $1-p$.

The probability distribution is given by:

$$P(X = 1) = p,$$

$$P(X = 0) = 1 - p.$$

The cumulative distribution function is as follows:

$$F_p(x) = \begin{cases} 0, & x < 0 \\ 1 - p, & 0 \leq x < 1 \\ 1, & x \geq 1 \end{cases}, x \in \mathbb{R}.$$

Input Parameters

FORTRAN:

<i>method</i>	INTEGER, INTENT(IN). Generation method.
<i>stream</i>	TYPE (VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream state structure.
<i>n</i>	INTEGER, INTENT(IN). Number of random values to be generated.
<i>p</i>	DOUBLE PRECISION, INTENT(IN). Success probability <i>p</i> of a trial.

C:

<i>method</i>	int. Generation method.
<i>stream</i>	VSLStreamStatePtr. Pointer to the stream state structure.
<i>n</i>	int. Number of random values to be generated.
<i>p</i>	double. Success probability <i>p</i> of a trial.

Output Parameters

FORTRAN:

r `INTEGER, INTENT(OUT)`. Vector of n Bernoulli distributed pseudorandom values.

C:

r `int*`. Vector of n Bernoulli distributed pseudorandom values.

Geometric

Generates geometrically distributed pseudorandom values.

Fortran:

`call virnggeometric(method, stream, n, r, p)`

C:

`virNgGeometric(method, stream, n, r, p)`

Discussion

This function generates geometrically distributed pseudorandom numbers with probability p of a single trial success, where $p \in \mathbf{R}$; $0 < p < 1$.

A geometrically distributed variate represents the number of independent Bernoulli trials preceding the first success. The probability of a single Bernoulli trial success is p .

The probability distribution is given by:

$$P(X = k) = p \cdot (1 - p)^k, \quad k \in \{0, 1, 2, \dots\}.$$

The cumulative distribution function is as follows:

$$F_p(x) = \begin{cases} 0, & x < 0 \\ 1 - (1 - p)^{\lfloor x + 1 \rfloor}, & x \geq 0 \end{cases}, \quad x \in \mathbf{R}.$$

Input Parameters

FORTRAN:

<i>method</i>	INTEGER, INTENT(IN). Generation method.
<i>stream</i>	TYPE (VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream state structure.
<i>n</i>	INTEGER, INTENT(IN). Number of random values to be generated.
<i>p</i>	DOUBLE PRECISION, INTENT(IN). Success probability <i>p</i> of a trial.

C:

<i>method</i>	int. Generation method.
<i>stream</i>	VSLStreamStatePtr. Pointer to the stream state structure.
<i>n</i>	int. Number of random values to be generated.
<i>p</i>	double. Success probability <i>p</i> of a trial.

Output Parameters

FORTRAN:

<i>r</i>	INTEGER, INTENT(OUT). Vector of <i>n</i> geometrically distributed pseudorandom values.
----------	---

C:

<i>r</i>	int*. Vector of <i>n</i> geometrically distributed pseudorandom values.
----------	---

Binomial

Generates binomially distributed pseudorandom numbers.

Fortran:

```
call virngbinomial( method, stream, n, r, ntrial, p )
```

C:

```
viRngBinomial( method, stream, n, r, ntrial, p )
```

Discussion

This function generates binomially distributed pseudorandom numbers with number of independent Bernoulli trials m , and with probability p of a single trial success, where $p \in R$; $0 \leq p \leq 1$, $m \in N$.

A binomially distributed variate represents the number of successes in m independent Bernoulli trials with probability of a single trial success p .

The probability distribution is given by:

$$P(X = k) = C_m^k p^k (1-p)^{m-k}, \quad k \in \{0, 1, \dots, m\}.$$

The cumulative distribution function is as follows:

$$F_{m,p}(x) = \begin{cases} 0, & x < 0 \\ \sum_{k=0}^{\lfloor x \rfloor} C_m^k p^k (1-p)^{m-k}, & 0 \leq x < m, \quad x \in R. \\ 1, & x \geq m \end{cases}$$

Input Parameters

FORTTRAN:

<i>method</i>	INTEGER, INTENT(IN). Generation method.
<i>stream</i>	TYPE (VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream state structure.

<i>n</i>	<code>INTEGER, INTENT(IN)</code> . Number of random values to be generated.
<i>ntrial</i>	<code>INTEGER, INTENT(IN)</code> . Number of independent trials <i>m</i> .
<i>p</i>	<code>DOUBLE PRECISION, INTENT(IN)</code> . Success probability <i>p</i> of a single trial.
C:	
<i>method</i>	<code>int</code> . Generation method.
<i>stream</i>	<code>VSLStreamStatePtr</code> . Pointer to the stream state structure.
<i>n</i>	<code>int</code> . Number of random values to be generated.
<i>ntrial</i>	<code>int</code> . Number of independent trials <i>m</i> .
<i>p</i>	<code>double</code> . Success probability <i>p</i> of a single trial.

Output Parameters

FORTRAN:

<i>r</i>	<code>INTEGER, INTENT(OUT)</code> . Vector of <i>n</i> binomially distributed pseudorandom values.
----------	--

C:

<i>r</i>	<code>int*</code> . Vector of <i>n</i> binomially distributed pseudorandom values.
----------	--

Hypergeometric

Generates hypergeometrically distributed pseudorandom values.

Fortran:

```
call virnghypergeometric( method, stream, n, r, l, s, m )
```

C:

```
viRngHypergeometric( method, stream, n, r, l, s, m )
```

Discussion

This function generates hypergeometrically distributed pseudorandom values with lot size l , size of sampling s , and number of marked elements in the lot m , where $l, m, s \in N \cup \{0\}$; $l \geq \max(s, m)$.

Consider a lot of l elements comprising m “marked” and $l-m$ “unmarked” elements. A trial sampling without replacement of exactly s elements from this lot helps to define the hypergeometric distribution, which is the probability that the group of s elements contains exactly k marked elements.

The probability distribution is given by:

$$P(X = k) = \frac{C_m^k C_{l-m}^{s-k}}{C_l^s}, \quad k \in \{\max(0, s + m - l), \dots, \min(s, m)\}.$$

The cumulative distribution function is as follows:

$$F_{l, s, m}(x) = \begin{cases} 0, & x < \max(0, s + m - l) \\ \sum_{k=\max(0, s+m-l)}^{\lfloor x \rfloor} \frac{C_m^k C_{l-m}^{s-k}}{C_l^s}, & \max(0, s + m - l) \leq x \leq \min(s, m) \\ 1, & x > \min(s, m) \end{cases}$$

Input Parameters

FORTRAN:

<i>method</i>	INTEGER, INTENT(IN). Generation method.
<i>stream</i>	TYPE (VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream state structure.
<i>n</i>	INTEGER, INTENT(IN). Number of random values to be generated.
<i>l</i>	INTEGER, INTENT(IN). Lot size <i>l</i> .
<i>s</i>	INTEGER, INTENT(IN). Size of sampling without replacement <i>s</i> .
<i>m</i>	INTEGER, INTENT(IN). Number of marked elements <i>m</i> .

C:

<i>method</i>	int. Generation method.
<i>stream</i>	VSLStreamStatePtr. Pointer to the stream state structure.
<i>n</i>	int. Number of random values to be generated.
<i>l</i>	int. Lot size <i>l</i> .
<i>s</i>	int. Size of sampling without replacement <i>s</i> .
<i>m</i>	int. Number of marked elements <i>m</i> .

Output Parameters

FORTRAN:

<i>r</i>	INTEGER, INTENT(OUT). Vector of <i>n</i> hypergeometrically distributed pseudorandom values.
----------	--

C:

`r` `int*`. Vector of `n` hypergeometrically distributed pseudorandom values.

Poisson

Generates Poisson distributed pseudorandom values.

Fortran:

`call virngpoisson(method, stream, n, r, lambda)`

C:

`viRngPoisson(method, stream, n, r, lambda)`

Discussion

This function generates Poisson distributed pseudorandom numbers with distribution parameter λ , where $\lambda \in R$; $\lambda > 0$.

The probability distribution is given by:

$$P(X = k) = \frac{\lambda^k e^{-\lambda}}{k!}, \quad k \in \{0, 1, 2, \dots\}.$$

The cumulative distribution function is as follows:

$$F_{\lambda}(x) = \begin{cases} \sum_{k=0}^{\lfloor x \rfloor} \frac{\lambda^k e^{-\lambda}}{k!}, & x \geq 0 \\ 0, & x < 0 \end{cases}, \quad x \in R.$$

Input Parameters

FORTRAN:

<i>method</i>	INTEGER, INTENT(IN). Generation method.
<i>stream</i>	TYPE (VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream state structure.
<i>n</i>	INTEGER, INTENT(IN). Number of random values to be generated.
<i>lambda</i>	DOUBLE PRECISION, INTENT(IN). Distribution parameter λ .

C:

<i>method</i>	int. Generation method.
<i>stream</i>	VSLStreamStatePtr. Pointer to the stream state structure.
<i>n</i>	int. Number of random values to be generated.
<i>lambda</i>	double. Distribution parameter λ .

Output Parameters

FORTRAN:

<i>r</i>	INTEGER, INTENT(OUT). Vector of <i>n</i> Poisson distributed pseudorandom values.
----------	---

C:

<i>r</i>	int*. Vector of <i>n</i> Poisson distributed values.
----------	--

NegBinomial

Generates pseudorandom numbers with negative binomial distribution.

Fortran:

```
call virngnegbinomial( method, stream, n, r, a, p )
```

C:

```
viRngNegBinomial( method, stream, n, r, a, p )
```

Discussion

This function generates pseudorandom numbers with negative binomial distribution and distribution parameters a and p , where $p, a \in \mathbf{R}$; $0 < p < 1$; $a > 0$.

If the first distribution parameter $a \in \mathbf{N}$, this distribution is the same as Pascal distribution. If $a \in \mathbf{N}$, the distribution can be interpreted as the expected time of a -th success in a sequence of Bernoulli trials, when the probability of success is p .

The probability distribution is given by:

$$P(X = k) = C_{a+k-1}^k p^a (1-p)^k, \quad k \in \{0, 1, 2, \dots\}.$$

The cumulative distribution function is as follows:

$$F_{a, p}(x) = \begin{cases} \sum_{k=0}^{\lfloor x \rfloor} C_{a+k-1}^k p^a (1-p)^k, & x \geq 0 \\ 0, & x < 0 \end{cases}, \quad x \in \mathbf{R}.$$

Input Parameters

FORTRAN:

<i>method</i>	INTEGER, INTENT(IN). Generation method.
<i>stream</i>	TYPE (VSL_STREAM_STATE), INTENT(IN). Descriptor of the stream state structure.
<i>n</i>	INTEGER, INTENT(IN). Number of random values to be generated.
<i>a</i>	DOUBLE PRECISION, INTENT(IN). The first distribution parameter <i>a</i> .
<i>p</i>	DOUBLE PRECISION, INTENT(IN). The second distribution parameter <i>p</i> .

C:

<i>method</i>	int. Generation method.
<i>stream</i>	VSLStreamStatePtr. Pointer to the stream state structure.
<i>n</i>	int. Number of random values to be generated.
<i>a</i>	double. The first distribution parameter <i>a</i> .
<i>p</i>	double. The second distribution parameter <i>p</i> .

Output Parameters

FORTRAN:

<i>r</i>	INTEGER, INTENT(OUT). Vector of <i>n</i> pseudorandom values with negative binomial distribution.
----------	---

C:

<i>r</i>	int*. Vector of <i>n</i> pseudorandom values with negative binomial distribution.
----------	---

Advanced Service Subroutines

This section describes service subroutines for registering a basic generator and obtaining properties of the previously registered basic generators.

Table 8-4 Advanced Service Subroutines

Subroutine	Short Description
RegisterBrng	Registers a user-designed basic generator.
GetBrngProperties	Returns the structure with properties of the basic generator with a given number.

Data types

The subroutines of this section refer to a structure defining the properties of the basic generator:

Example 8-5 Fortran Version

```

TYPE VSL_BRNG_PROPERTIES
    INTEGER streamstatesize
    INTEGER nseeds
    INTEGER includeszero
    INTEGER wordsize
    INTEGER nbits
    INTEGER initstream
    INTEGER sbrng
    INTEGER dbrng
    INTEGER ibrng
END TYPE VSL_BRNG_PROPERTIES

```

Example 8-6 C Version

```

typedef struct _VSLBRngProperties {
    int          StreamStateSize;
    int          NSeeds;
    int          IncludesZero;
    int          WordSize;
    int          NBits;
    InitStreamPtr  InitStream;
    sBRngPtr      sBRng;
    dBRngPtr      dBRng;
    iBRngPtr      iBRng;
} VSLBRngProperties;

```

Example 8-7 Pointers to Functions

```

typedef int (*InitStreamPtr)( int method, void * stream, int n,
                             const unsigned int params[] );
typedef void (*sBRngPtr)( void * stream, int n, float r[],
                          float a, float b );
typedef void (*dBRngPtr)( void * stream, int n, double r[],
                          double a, double b );
typedef void (*iBRngPtr)( void * stream, int n,
                          unsigned int r[] );

```

Table 8-5 Field Descriptions

Field	Short Description
FORTRAN: <code>streamstatesize</code> C: <code>StreamStateSize</code>	The size, in bytes, of the stream state structure for a given basic generator.
FORTRAN: <code>nseeds</code> C: <code>NSeeds</code>	The number of 32-bit initial conditions (seeds) necessary to initialize the stream state structure for a given basic generator.
FORTRAN: <code>includeszero</code> C: <code>IncludesZero</code>	Flag value indicating whether the generator can produce a pseudorandom 0 ¹ .
FORTRAN: <code>wordsize</code> C: <code>WordSize</code>	Machine word size, in bytes, used in integer-value computations. Possible values: 4, 8, and 16 for 32, 64, and 128-bit generators, respectively.
FORTRAN: <code>nbits</code> C: <code>NBits</code>	The number of bits required to represent a pseudorandom value in integer arithmetic. Note that, for instance, 48-bit pseudorandom values are stored to 64-bit (8 byte) memory locations. In this case, <code>WordSize</code> is equal to 8 (number of bytes used to store the pseudorandom value), while <code>NBits</code> contains the actual number of bits occupied by the value (in this example, 48).

Table 8-5 Field Descriptions (continued)

Field	Short Description
FORTRAN: <code>initstream</code> C: <code>InitStream</code>	Contains the pointer to the initialization subroutine of a given basic generator.
FORTRAN: <code>sbrng</code> C: <code>sBRng</code>	Contains the pointer to the basic generator of single precision real numbers uniformly distributed over the interval (a,b) (<code>REAL</code> in FORTRAN and <code>float</code> in C).
FORTRAN: <code>dbrng</code> C: <code>dBRng</code>	Contains the pointer to the basic generator of double precision real numbers uniformly distributed over the interval (a,b) (<code>DOUBLE PRECISION</code> in FORTRAN and <code>double</code> in C).
FORTRAN: <code>ibrng</code> C: <code>iBRng</code>	Contains the pointer to the basic generator of integer numbers with uniform bit distribution ² (<code>INTEGER</code> in FORTRAN and <code>unsigned int</code> in C).

1. Certain types of generators, for example, generalized feedback shift registers can potentially generate a pseudorandom 0. On the other hand, generators like multiplicative congruential generators never generate such a number. In most cases this information is irrelevant because the probability of generating a zero value is extremely small. However, in certain non-uniform distribution generators the possibility for a basic generator to produce a pseudorandom zero may lead to generation of an infinitely large number (overflow). Even though the software handles overflows correctly, so that they may be interpreted as $+\infty$ and $-\infty$, the user has to be careful and verify the final results. If an infinitely large number may affect the computation, the user should either remove such numbers from the generated vector, or use safe generators, which do not produce pseudorandom 0.
2. A specific generator that permits operations over single bits and bit groups of pseudorandom numbers.

RegisterBrng

Registers user-defined basic generator.

Fortran:

```
brng = vslregisterbrng( properties )
```

C:

```
brng = vslRegisterBrng( properties )
```

Discussion

An example of a registration procedure can be found in the respective directory `vs1\examples`.

Input Parameters

FORTRAN:

<i>properties</i>	<code>TYPE (VSL_BRNG_PROPERTIES), INTENT (IN)</code> . Structure containing properties of the basic generator to be registered.
-------------------	---

C:

<i>properties</i>	<code>VSLBrngProperties*</code> . Structure containing properties of the basic generator to be registered.
-------------------	--

Output Parameters

FORTRAN:

<i>brng</i>	<code>INTEGER</code> . The number (index) of the registered basic generator; used for identification. Negative values indicate the registration error.
-------------	--

C:

brng **int**. The number (index) of the registered basic generator; used for identification. Negative values indicate the registration error.

GetBrngProperties

Returns structure with properties of a given basic generator.

Fortran:

```
call vslgetbrngproperties( brng, properties )
```

C:

```
call vslGetBrngProperties( brng, properties )
```

Input Parameters

FORTRAN:

brng **INTEGER, INTENT(IN)**. Number (index) of the registered basic generator.

C:

brng **int**. Number (index) of the registered basic generator.

Output Parameters

FORTRAN:

properties **TYPE (VSL_BRNG_PROPERTIES), INTENT(OUT)**. Structure containing properties of the generator with number *brng*.

C:

properties

*VSLBrngProperties**. Structure containing properties of the generator with number *brng*.

Formats for User-Designed Generators

To register a user-designed basic generator using *RegisterBrng* function, you need to pass the pointer *iBrng* to the integer-value implementation of the generator; the pointers *sBrng* and *dBrng* to the generator implementations for single and double precision values, respectively; and pass the pointer *InitStream* to the stream initialization subroutine. This section contains recommendations on defining such functions with input and output arguments. An example of the registration procedure for a user-designed generator can be found in the respective directory *vsl\examples*.

InitStream

FORTRAN:

```

INTEGER FUNCTION mybrnginitstream( method, stream, n, params )
    INTEGER, INTENT (IN) :: method
    TYPE(MYSTREAM_STATE), INTENT (INOUT):: stream
    INTEGER, INTENT (IN) :: n
    INTEGER, INTENT (IN) :: params
! Initialize the stream
    ...
END SUBROUTINE mybrnginitstream

```

C:

```

int MyBrngInitStream( int method, VSLStreamStatePtr stream,
    int n, const unsigned int params[] )
{
    /* Initialize the stream */
    ...
}

```



```
} /* MyBrngInitStream */
```

Discussion

The initialization subroutine of a user-designed generator must initialize *stream* according to the specified initialization *method*, initial conditions *params* and the argument *n*. The value of *method* determines the initialization method to be used.

- If *method* is equal to 0, the initialization is by the standard generation method, which must be supported by all basic generators. In this case the function assumes that the *stream* structure was not previously initialized. The value of *n* is used as the actual number of 32-bit values passed as initial conditions through *params*. Note, that the situation when the actual number of initial conditions passed to the function is not sufficient to initialize the generator is not an error. Whenever it occurs, the basic generator must initialize the missing conditions using default settings.
- If *method* is equal to 1, the generation is by the leapfrog method, where *n* specifies the number of computational nodes (independent streams). Here the function assumes that the *stream* was previously initialized by the standard generation method. In this case *params* contains only one element, which identifies the computational node. If the generator does not support the leapfrog method, the function must return the error code `VSL_ERROR_LEAPFROG_UNSUPPORTED`.
- If *method* is equal to 2, the generation is by the block-splitting method. Same as above, the *stream* is assumed to be previously initialized by the standard generation method; *params* is not used, *n* identifies the number of skipped elements. If the generator does not support the block-splitting method, the function must return the error code `VSL_ERROR_SKIPAHEAD_UNSUPPORTED`.

For a more detailed description of the leapfrog and the block-splitting methods, refer to the description of `LeapfrogStream` and `SkipAheadStream`, respectively.

Stream state structure is individual for every generator. However, each structure has a number of fields that are the same for all the generators:

FORTRAN:

```
type(mystream_state)
    INTEGER*4   reserved1
    INTEGER*4   reserved2
    INTEGER*4   reserved3
    INTEGER*4   reserved4
    [ fields specific for the given generator ]
end type mystream_state
```

C:

```
typedef struct
{
    uint64 Reserved1;
    uint64 Reserved2;
    [ fields specific for the given generator ]
} MyStreamState
```

The fields *Reserved1* and *Reserved2* are reserved for private needs only, and must not be modified by the user. When including specific fields into the structure, follow the rules below:

- The fields must fully describe the current state of the generator. For example, the state of a linear congruential generator can be identified by only one initial condition;
- If the generator can use both the leapfrog and the block-splitting methods, additional fields should be introduced to identify the independent streams. For example, in $LCG(a, c, m)$, apart from the initial conditions, two more fields should be specified: the value of the multiplier a^k and the value of the increment $(a^k - 1)c / (a - 1)$.

For a more detailed discussion, refer to [[Knuth81](#)], and [[Gentle98](#)]. An example of the registration procedure can be found in the respective directory `vs1\examples`.

iBRng

FORTRAN:

```
SUBROUTINE imybrng( stream, n, r )
```

```
TYPE(MYSTREAM_STATE), INTENT(INOUT):: stream
INTEGER, INTENT(IN)    :: n
INTEGER, DIMENSION(*), INTENT(OUT) :: r
! Generating integer random numbers
! Pay attention to word size needed to
! store one random number
DO i = 1, n
    R(I) = ...
END DO
! Update stream state
END SUBROUTINE imybrng
```

C:

```
void iMyBrng( VSLStreamStatePtr stream, int n,
             unsigned int r[] )
{
    int    i;    /* Loop variable */
    /* Generating integer random numbers */
    /* Pay attention to word size needed to
       store only random number */
    for( i = 0; i < n; i++ )
    {
        r[i] = ...
    }
    /* Update stream state */
    ...
} /* iMyBrng */
```



NOTE. When using 64 and 128-bit generators, consider digit capacity to store the numbers to the pseudorandom vector r correctly. For example, storing one 64-bit value requires two elements of r , the first to store the lower 32 bits and the second to store the higher 32 bits. Similarly, use 4 elements of r to store a 128-bit value.

sBRng

FORTRAN:

```

SUBROUTINE smybrng( stream, n, r, a, b )
    TYPE(MYSTREAM_STATE), INTENT(INOUT):: stream
    INTEGER, INTENT(IN)    :: n
    REAL, DIMENSION(n), INTENT(OUT)  :: r
    REAL, INTENT(IN)      :: a
    REAL, INTENT(IN)      :: b
! Generating real (a,b) random numbers
    DO i = 1, n
        R(I) = ...
    END DO
! Update stream state
END SUBROUTINE smybrng

```

C:

```

void sMyBrng( VSLStreamStatePtr stream, int n, float r[],
             float a, float b )
{
    int    i;    /* Loop variable */
    /* Generating float (a,b) random numbers */
    for ( i = 0; i < n; i++ )
    {
        r[i] = ...
    }
    /* Update stream state */
}

```

```

...
} /* sMyBrng */

```

dBRng

FORTRAN:

```

SUBROUTINE dmybrng( stream, n, r, a, b )
  TYPE(MYSTREAM_STATE), INTENT(INOUT) :: stream
  INTEGER, INTENT(IN)      :: n
  DOUBLE PRECISION, DIMENSION(n), INTENT(OUT)  :: r
  REAL, INTENT(IN)      :: a
  REAL, INTENT(IN)      :: b
! Generating double precision (a,b) random numbers
  DO i = 1, n
    R(I) = ...
  END DO
! Update stream state
...
END SUBROUTINE dmybrng

```

C:

```

void dMyBrng( VSLStreamStatePtr stream, int n, double r[],
             double a, double b )
{
  int      i;      /* Loop variable */
  /* Generating double (a,b) random numbers */
  for ( i = 0; i < n; i++ )
  {
    r[i] = ...
  }
  /* Update stream state */
  ...
} /* dMyBrng */

```

Advanced DFT Functions

9

The Fast Fourier Transform (FFT) algorithm that calculates the Discrete Fourier Transform (DFT) is an indispensable tool in a vast number of fields. This chapter describes the set of new DFT functions implemented in Intel® MKL, which present a uniform and easy-to-use Applications Programmer Interface (API) providing fast computation of DFT via FFT.

The Discrete Fourier Transform function library of Intel MKL provides one-dimensional, two-dimensional, and multi-dimensional (up to the order of 7) routines.

Both Fortran- and C-interfaces exist for all transform functions.

Although Intel MKL still supports the FFT interface described in chapter 3 of this manual, users are encouraged to migrate to the new advanced DFT functions in their application programs. Unlike the older FFT routines, the DFT routines support transform lengths of other than powers of 2 mixed radix.

The full list of DFT functions implemented in Intel MKL is given in the table below:

Table 9-1 DFT Functions in Intel MKL

Function Name	Operation
Descriptor Manipulation Functions	
<u>DftiCreateDescriptor</u>	Allocates memory for the descriptor data structure and instantiates it with default configuration settings.
<u>DftiCommitDescriptor</u>	Performs all initialization that facilitates the actual DFT computation.
<u>DftiCopyDescriptor</u>	Copies an existing descriptor.
<u>DftiFreeDescriptor</u>	Frees memory allocated for a descriptor.

Table 9-1 DFT Functions in Intel MKL (continued)

Function Name	Operation
DFT Computation Functions	
DftiComputeForward	Computes the forward DFT.
DftiComputeBackward	Computes the backward DFT.
Descriptor Configuration Functions	
DftiSetValue	Sets one particular configuration parameter with the specified configuration value.
DftiGetValue	Gets the configuration value of one particular configuration parameter.
Status Checking Functions	
DftiErrorClass	Checks if the status reflects an error of a predefined class.
DftiErrorMessage	Generates an error message.

Description of DFT functions is followed by discussion of configuration settings (see [Configuration Settings](#)) and various configuration parameters used.

Computing DFT

DFT functions described later in this chapter are implemented in Fortran and C interface. Fortran stands for Fortran 95. DFT interface relies critically on many modern features offered in Fortran 95 that have no counterpart in Fortran 77.



NOTE. *Following the explicit function interface in Fortran, data array must be defined as one-dimensional for any transformation type.*

The materials presented in this chapter assume the availability of native complex types in C as they are specified in C9X.

Before the presenting every function, a couple of usage examples are given.

For most common situations, we expect a DFT computation can be effected by four function calls. Here are the examples of two one-dimensional computations.

Example 9-1 One-dimensional DFT (Fortran-interface)

```
! Fortran example.
! 1D complex to complex, and real to conjugate even
Use MKL_DFTI
Complex :: X(32)
Real :: Y(34)
type(DFTI_DESCRIPTOR), POINTER :: My_Desc1_Handle, My_Desc2_Handle
Integer :: Status
...put input data into X(1),...,X(32); Y(1),...,Y(32)

! Perform a complex to complex transform
Status = DftiCreateDescriptor( My_Desc1_Handle, DFTI_SINGLE,
    DFTI_COMPLEX, 1, 32 )
Status = DftiCommitDescriptor( My_Desc1_Handle )
Status = DftiComputeForward( My_Desc1_Handle, X )
Status = DftiFreeDescriptor(My_Desc1_Handle)
! result is given by {X(1),X(2),...,X(32)}

! Perform a real to complex conjugate even transform
Status = DftiCreateDescriptor(My_Desc2_Handle, DFTI_SINGLE,
    DFTI_REAL, 1, 32)
Status = DftiCommitDescriptor(My_Desc2_Handle)
Status = DftiComputeForward(My_Desc2_Handle, Y)
Status = DftiFreeDescriptor(My_Desc2_Handle)
! result is given by {Y(1)+iY(2), Y(3)+iY(4), ..., Y(33)+iY(34),
! Y(31)-iY(32), Y(29)-iY(30), ..., Y(3)-iY(4)}.
```

Example 9-2 One-dimensional DFT (C-interface)

```

/* C example, float _Complex is defined in C9X */
#include "mkl_dfti.h"
float _Complex x[32];
float y[34];
DFTI_DESCRIPTOR *my_desc1_handle, *my_desc2_handle;
/* .... or alternatively
DFTI_DESCRIPTOR_HANDLE my_desc1_handle, my_desc2_handle; */

long status;
...put input data into x[0],...,x[31]; y[0],...,y[31]
status = DftiCreateDescriptor( &my_desc1_handle, DFTI_SINGLE,
    DFTI_COMPLEX, 1, 32);
status = DftiCommitDescriptor( my_desc1_handle );
status = DftiComputeForward( my_desc1_handle, x);
status = DftiFreeDescriptor(&my_desc1_handle);
/* result is x[0], ..., x[31] */
status = DftiCreateDescriptor( &my_desc2_handle, DFTI_SINGLE,
    DFTI_REAL, 1, 32);
status = DftiCommitDescriptor( my_desc2_handle);
status = DftiComputeForward( my_desc2_handle, y);
status = DftiFreeDescriptor(&my_desc2_handle);
/* y[0]+iy[1], ..., y[32]+iy[33], y[30]-iy[31], ..., y[2]-iy[3] */

```

The following is an example of two simple two-dimensional transforms. Notice that the data and result parameters in computation functions are all declared as assumed-size rank-1 array `DIMENSION(0:*)`. Therefore two-dimensional array must be transformed to one-dimensional array by `EQUIVALENCE` statement or other facilities of Fortran.

Example 9-3 Two-dimensional DFT (Fortran-interface)

```
! Fortran example.
! 2D complex to complex, and real to conjugate even
Use MKL_DFTI
Complex :: X_2D(32,100)
Real :: Y_2D(34, 102)
Complex :: X(3200)
Real :: Y(3468)
Equivalence (X_2D, X)
Equivalence (Y_2D, Y)
type(DFTI_DESCRIPTOR), POINTER :: My_Desc1_Handle, My_Desc2_Handle
Integer :: Status, L(2)
...put input data into X_2D(j,k), Y_2D(j,k), 1<=j=32,1<=k=100
...set L(1) = 32, L(2) = 100
...the transform is a 32-by-100

! Perform a complex to complex transform
Status = DftiCreateDescriptor( My_Desc1_Handle, DFTI_SINGLE,
    DFTI_COMPLEX, 2, L)
Status = DftiCommitDescriptor( My_Desc1_Handle)
Status = DftiComputeForward( My_Desc1_Handle, X)
Status = DftiFreeDescriptor(My_Desc1_Handle)
! result is given by X_2D(j,k), 1<=j<=32, 1<=k<=100

! Perform a real to complex conjugate even transform
Status = DftiCreateDescriptor( My_Desc2_Handle, DFTI_SINGLE,
    DFTI_REAL, 2, L)
Status = DftiCommitDescriptor( My_Desc2_Handle)
Status = DftiComputeForward( My_Desc2_Handle, Y)
Status = DftiFreeDescriptor(My_Desc2_Handle)
! result is given by the complex value z(j,k) 1<=j<=32; 1<=k<=100 where
! z(j,k) = Y_2D(2j-1,k) + iY_2D(2j,k) 1<=j<=17; 1<=k<=100
! z(j,k) = Y_2D(2(34-j)-1,k) - iY_2D(2(34-j),k) 18<=j<=32; 1<=k<=100
```

Example 9-4 Two-dimensional DFT (C-interface)

```

/* C example */
#include "mkl_dfti.h"
float _Complex x[32][100];
float y[34][102];
DFTI_DESCRIPTOR_HANDLE my_desc1_handle, my_desc2_handle;
/* or alternatively
DFTI_DESCRIPTOR *my_desc1_handle, *my_desc2_handle; */
long status, l[2];
...put input data into x[j][k] 0<=j<=31, 0<=k<=99
...put input data into y[j][k] 0<=j<=31, 0<=k<=99
l[0] = 32; l[1] = 100;
status = DftiCreateDescriptor( &my_desc1_handle, DFTI_SINGLE,
                             DFTI_COMPLEX, 2, 1);
status = DftiCommitDescriptor( my_desc1_handle);
status = DftiComputeForward( my_desc1_handle, x);
status = DftiFreeDescriptor(&my_desc1_handle);
/* result is the complex value x[j][k], 0<=j<=31, 0<=k<=99 */
status = DftiCreateDescriptor( &my_desc2_handle, DFTI_SINGLE,
                             DFTI_REAL, 2, 1);
status = DftiCommitDescriptor( my_desc2_handle);
status = DftiComputeForward( my_desc2_handle, y);
status = DftiFreeDescriptor(&my_desc2_handle);
/* result is the complex value z(j,k) 0<=j<=31; 0<=k<=99
/* z(j,k) = y[2j][k] + iy[2j+1][k] 0<=j<=16; 0<=k<=99 */
/* z(j,k) = y[2(32-j)][k] - iy[2(32-j)+1][k] 17<=j<=31; 1<=k<=100 */

```

The record of type `DFTI_DESCRIPTOR`, when created, contains information about the length and domain of the DFT to be computed. Moreover, it contains the setting of a rather large number of configuration parameters. The examples above use the default settings for all of these parameters, which include, for example, the following:

- the DFT to be computed does not have a scale factor;

- there is only one set of data to be transformed;
- the data is stored contiguously in memory;
- the forward transform is defined to be the formula using $e^{-i2\pi jk/n}$ rather than $e^{+i2\pi jk/n}$;
- complex data is stored in the native complex data type;
- the computed result overwrites (in place) the input data; etc.

Should any one of these many default settings be inappropriate, they can be changed one-at-a-time through the function [DftiSetValue](#). For example, to preserve the input data after the DFT computation, the configuration of the F should be changed to "not in place" from the default choice of "in place." The code below illustrates how this can be done:

Example 9-5 Changing Default Settings (Fortran)

```
! Fortran example
! 1D complex to complex, not in place
Use MKL_DFTI
Complex :: X_in(32), X_out(32)
type(DFTI_DESCRIPTOR), POINTER :: My_Desc_Handle
Integer :: Status
...put input data into X_in(j), 1<=j<=32
Status = DftiCreateDescriptor( My_Desc_Handle, DFTI_SINGLE,
    DFTI_COMPLEX, 1, 32)
Status = DftiSetValue( My_Desc_Handle, DFTI_PLACEMENT, DFTI_NOT_INPLACE)
Status = DftiCommitDescriptor( My_Desc_Handle)
Status = DftiComputeForward( My_Desc_Handle, X_in, X_out)
Status = DftiFreeDescriptor( My_Desc_Handle)
! result is X_out(1),X_out(2),...,X_out(32)
```

Example 9-6 Changing Default Settings (C)

```
/* C example */
#include "mkl_dfti.h"
float _Complex x_in[32], x_out[32];
DFTI_DESCRIPTOR_HANDLE my_desc_handle;
/* or alternatively
DFTI_DESCRIPTOR *my_desc_handle; */
long status;
...put input data into x_in[j], 0 <= j < 32
status = DftiCreateDescriptor( &my_desc_handle, DFTI_SINGLE,
DFTI_COMPLEX, 1, 32);
status = DftiSetValue( my_desc_handle, DFTI_PLACEMENT,
DFTI_NOT_INPLACE);
status = DftiCommitDescriptor( my_desc_handle);
status = DftiComputeForward( my_desc_handle, x_in, x_out);
status = DftiFreeDescriptor(&my_desc_handle);
/* result is x_out[0], x_out[1], ..., x_out[31] */
```

The approach adopted in Intel MKL for DFT computation uses one single data structure, the descriptor, to record flexible configuration whose parameters can be changed independently. This results in enhanced functionality and ease of use.

DFT Interface

To use the advanced DFT functions, the user needs to access the module `MKL_DFTI` through the "use" statement in Fortran; or access the header file `mk1_dfti.h` through "include" in C.

The Fortran interface provides a derived type `DFTI_DESCRIPTOR`; a number of named constants representing various names of configuration parameters and their possible values; and a number of overloaded functions through the generic functionality of Fortran 95.

The C interface provides a structure type `DFTI_DESCRIPTOR`, a macro definition

```
#define DFTI_DESCRIPTOR_HANDLE DFTI_DESCRIPTOR *
```

a number of named constants of two enumeration types

```
DFTI_CONFIG_PARAM and DFTI_CONFIG_VALUE;
```

and a number of functions, some of which accept different number of input arguments.



NOTE. *Some of the functions and/or functionality described in the subsequent sections of this chapter may not be supported by the currently available implementation of the library. You can find the complete list of the implementation-specific exceptions in the release notes to your version of the library.*

There are four main categories of DFT functions in Intel MKL:

1. **Descriptor Manipulation.** There are four functions in this category. The first one creates a DFT descriptor whose storage is allocated dynamically by the routine. This function configures the descriptor with default settings corresponding to a few input values supplied by the user.

The second "commits" the descriptor to all its setting. In practice, this usually means that all the necessary precomputation will be performed. This may include factorization of the input length and

computation of all the required twiddle factors. The third function makes an extra copy of a descriptor, and the fourth function frees up all the memory allocated for the descriptor information.

2. **DFT Computation.** There are two functions in this category. The first effects a forward DFT computation, and the second a backward DFT computation.
3. **Descriptor configuration.** There are two functions in this category. One function sets one specific value to one of the many configuration parameters that are changeable (a few are not); the other gets the current value of any one of these configuration parameters (all are readable). These parameters, though many, are handled one-at-a-time.
4. **Status Checking.** The functions described in the three categories return an integer value denoting the status of the operation. In particular, a non-zero return value always indicates a problem of some sort. Envisioned to be further enhanced in later releases of Intel MKL, DFT interface at present provides for one logical status class function and a simple status message generation function.

Status Checking Functions

All of the descriptor manipulation, DFT computation, and descriptor configuration functions return an integer value denoting the status of the operation. Two functions serve to check the status. The first function is a logical function that checks if the status reflects an error of a predefined class, and the second is an error message function that returns a character string.

ErrorClass

Checks if the status reflects an error of a predefined class.

Usage

```
! Fortran
Predicate = DftiErrorClass( Status, Error_Class )
/* C */
predicate = DftiErrorClass( status, error_class );
```

Discussion

DFT interface in Intel MKL provides a set of predefined error class listed in [Table 9-2](#). These are named constants and have the type `INTEGER` in Fortran and `long` in C.

Table 9-2 Predefined Error Class

Named Constants	Comments
<code>DFTI_NO_ERROR</code>	No error
<code>DFTI_MEMORY_ERROR</code>	Usually associated with memory allocation
<code>DFTI_INVALID_CONFIGURATION</code>	Invalid settings of one or more configuration parameters
<code>DFTI_INCONSISTENT_CONFIGURATION</code>	Inconsistent configuration or input parameters
<code>DFTI_MULTITHREADED_ERROR</code>	Usually associated with OMP routine's error return value
<code>DFTI_BAD_DESCRIPTOR</code>	Descriptor is unusable for computation
<code>DFTI_UNIMPLEMENTED</code>	Unimplemented legitimate settings; implementation dependent
<code>DFTI_MKL_INTERNAL_ERROR</code>	Internal library error

Note that the correct usage is to check if the status returns `.TRUE.` or `.FALSE.` through the use of `DFTI_ERROR_CLASS` with a specific error class. Direct comparison of a status with the predefined class is an incorrect usage.

Example 9-7 Using Status Checking Function

```

from C language:

DFTI_DESCRIPTOR_HANDLE desc;
long status, class_error, value;
char* error_message;
. . . descriptor creation and other code
status = DftiGetValue( desc, DFTI_PRECISION, &value); //
//or any DFTI function

class_error = DftiErrorClass(status, DFTI_ERROR_CLASS);
if (! class_error) {
printf ("status is not a member of Predefined Error
Class\n");
} else {
error_message = DftiErrorMessage(status);
printf("error_message = %s \n", error_message);
}
. . .
from Fortran:

type(DFTI_DESCRIPTOR), POINTER :: desc
integer value, status
character(DFTI_MAX_MESSAGE_LENGTH) error_message
logical class_error
. . . descriptor creation and other code
status = DftiGetValue( desc, DFTI_PRECISION, value)

class_error = DftiErrorClass(status, DFTI_ERROR_CLASS)
if (.not. class_error) then
print *, 'status is not a member of Predefined Error
Class '
else
error_message = DftiErrorMessage(status)
print *, 'error_message = ', error_message
endif

```

Interface and prototype

```

//Fortran interface
INTERFACE DftiErrorClass
//Note that the body provided here is to illustrate the different

```

```
//argument list and types of dummy arguments. The interface
//does not guarantee what the actual function names are.
//Users can only rely on the function name following the
//keyword INTERFACE
  FUNCTION some_actual_function_8( Status, Error_Class )
    LOGICAL some_actual_function_8
    INTEGER, INTENT(IN) :: Status, Error_Class
  END FUNCTION some_actual_function_8
END INTERFACE DftiErrorClass

/* C prototype */
long DftiErrorClass( long , long );
```

ErrorMessage

Generates an error message.

Usage

```
! Fortran
ERROR_MESSAGE = DftiErrorMessage( Status )
/* C */
error_message = DftiErrorMessage( status );
```

Discussion

The error message function generates an error message character string. The maximum length of the string in Fortran is given by the named constant `DFTI_MAX_MESSAGE_LENGTH`. The actual error message is implementation dependent. In Fortran, the user needs to use a character string of length `DFTI_MAX_MESSAGE_LENGTH` as the target. In C, the function returns a pointer to a character string, that is, a character array with the delimiter '0'.

[Example 9-7](#) shows how this function can be implemented.

Interface and prototype

```
//Fortran interface
INTERFACE DftiErrorMessage
//Note that the body provided here is to illustrate the different
//argument list and types of dummy arguments. The interface
//does not guarantee what the actual function names are.
//Users can only rely on the function name following the
//keyword INTERFACE
    FUNCTION some_actual_function_9( Status, Error_Class )
        CHARACTER(LEN=DFTI_MAX_MESSAGE_LENGTH) some_actual_function_9( Status )
        INTEGER, INTENT(IN) :: Status
    END FUNCTION some_actual_function_9
END INTERFACE DftiErrorMessage

/* C prototype */
char *DftiErrorMessage( long );
```

Descriptor Manipulation

There are four functions in this category: create a descriptor, commit a descriptor, copy a descriptor, and free a descriptor.

CreateDescriptor

Allocates memory for the descriptor data structure and instantiates it with default configuration settings.

Usage

```
! Fortran
Status = DftiCreateDescriptor( Desc_Handle,      &
                               Precision,       &
                               Forward_Domain,  &
                               Dimension,      &
                               Length )

/* C */
status = DftiCreateDescriptor( &desc_handle,
                               precision,
                               forward_domain,
                               dimension,
                               length );
```

Discussion

This function allocates memory for the descriptor data structure and instantiates it with all the default configuration settings with respect to the precision, domain, dimension, and length of the desired transform. The domain is understood to be the domain of the forward transform. Since memory is allocated dynamically, the result is actually a pointer to the created descriptor. This function is slightly different from the "initialization" routine in more traditional software packages or libraries used for computing DFT. In all likelihood, this function will not perform

any significant computation work such as twiddle factors computation, as the default configuration settings can still be changed upon user's request through the value setting function [DftiSetValue](#).

The precision and (forward) domain are specified through named constants provided in DFT interface for the configuration values. The choices for precision are `DFTI_SINGLE` and `DFTI_DOUBLE`; and the choices for (forward) domain are `DFTI_COMPLEX`, `DFTI_REAL`, and `DFTI_CONJUGATE_EVEN`. See [Table 9-5](#) for the complete table of named constants for configuration values.

Dimension is a simple positive integer indicating the dimension of the transform. Length is either a simple positive integer for one-dimensional transform, or an integer array (pointer in C) containing the positive integers corresponding to the lengths dimensions for multi-dimensional transform.

The function returns `DFTI_NO_ERROR` when completes successfully. See [“Status Checking Functions”](#) for more information on returned status.

Interface and prototype

```
!Fortran interface.
INTERFACE DftiCreateDescriptor
!Note that the body provided here is to illustrate the different
!argument list and types of dummy arguments. The interface
!does not guarantee what the actual function names are.
!Users can only rely on the function name following the
!keyword INTERFACE
FUNCTION some_actual_function_1D( Desc_Handle, Prec, Dom, Dim, Length )
    INTEGER :: some_actual_function_1D
    TYPE(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
    INTEGER, INTENT(IN) :: Prec, Dom
    INTEGER, INTENT(IN) :: Dim, Length
END FUNCTION some_actual_function_1D

FUNCTION some_actual_function_HIGHD( Desc_Handle, Prec, Dom, Dim, Length )
    INTEGER :: some_actual_function_HIGHD
    TYPE(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
    INTEGER, INTENT(IN) :: Prec, Dom
    INTEGER, INTENT(IN) :: Dim, Length(*)
```

```
END FUNCTION some_actual_function_HIGHD
END INTERFACE DftiCreateDescriptor
```

Note that the function is overloaded as the actual argument for `Length` can be a scalar or a a rank-one array.

```
/* C prototype */
long DftiCreateDescriptor( DFTI_DESCRIPTOR_HANDLE *,
                          DFTI_CONFIG_PARAM ,
                          DFTI_CONFIG_PARAM ,
                          long ,
                          ... );
```

The variable arguments facility is used to cope with the argument for lengths that can be a scalar (`long`), or an array (`long *`).

CommitDescriptor

Performs all initialization that facilitates the actual DFT computation.

Usage

```
! Fortran
Status = DftiCommitDescriptor( Desc_Handle )
/* C */
status = DftiCommitDescriptor( desc_handle );
```

Discussion

The interface requires a function that commits a previously created descriptor be invoked before the descriptor can be used for DFT computations. Typically, this committal performs all initialization that facilitates the actual DFT computation. For a modern implementation, it may involve exploring many different factorizations of the input length to search for highly efficient computation method.

Any changes of configuration parameters of a committed descriptor via the set value function (see [“Descriptor Configuration”](#)) requires a re-committal of the descriptor before a computation function can be invoked. Typically, this committal function call is immediately followed by a computation function call (see [“DFT Computation”](#)).

The function returns `DFTI_NO_ERROR` when completes successfully. See [“Status Checking Functions”](#) for more information on returned status.

Interface and prototype

```
! Fortran interface
INTERFACE DftiCommitDescriptor
!Note that the body provided here is to illustrate the different
!argument list and types of dummy arguments. The interface
!does not guarantee what the actual function names are.
!Users can only rely on the function name following the
!keyword INTERFACE
    FUNCTION some_actual function_1 ( Desc_Handle )
        INTEGER :: some_actual function_1
        TYPE(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
    END FUNCTION some_actual function_1
END INTERFACE DftiCommitDescriptor

/* C prototype */
long DftiCommitDescriptor( DFTI_DESCRIPTOR_HANDLE );
```

CopyDescriptor

Copies an existing descriptor.

Usage

```
! Fortran
```

```

Status = DftiCopyDescriptor( Desc_Handle_Original,
                             Desc_Handle_Copy )
/* C */
status = DftiCopyDescriptor( desc_handle_original,
                             &desc_handle_copy );

```

Discussion

This function makes a copy of an existing descriptor and provides a pointer to it. The purpose is that all information of the original descriptor will be maintained even if the original is destroyed via the free descriptor function `DftiFreeDescriptor`.

The function returns `DFTI_NO_ERROR` when completes successfully. See [“Status Checking Functions”](#) for more information on returned status.

Interface and prototype

```

! Fortran interface
INTERFACE DftiCopyDescriptor
! Note that the body provided here is to illustrate the different
! argument list and types of dummy arguments. The interface
! does not guarantee what the actual function names are.
! Users can only rely on the function name following the
! keyword INTERFACE
    FUNCTION some_actual_function_2( Desc_Handle_Original,
                                     Desc_Handle_Copy )

        INTEGER :: some_actual_function_2
        TYPE(DFTI_DESCRIPTOR), POINTER :: Desc_Handle_Original, Desc_Handle_Copy
    END FUNCTION some_actual_function_2
END INTERFACE DftiCopyDescriptor

/* C prototype */
long DftiCopyDescriptor( DFTI_DESCRIPTOR_HANDLE, DFTI_DESCRIPTOR_HANDLE * );

```

FreeDescriptor

Frees memory allocated for a descriptor.

Usage

```
! Fortran
Status = DftiFreeDescriptor( Desc_Handle )
/* C */
status = DftiFreeDescriptor( &desc_handle );
```

Discussion

This function frees up all memory space allocated for a descriptor.

The function returns `DFTI_NO_ERROR` when completes successfully. See [“Status Checking Functions”](#) for more information on returned status.

Interface and prototype

```
! Fortran interface
INTERFACE DftiFreeDescriptor
//Note that the body provided here is to illustrate the different
//argument list and types of dummy arguments. The interface
//does not guarantee what the actual function names are.
//Users can only rely on the function name following the
//keyword INTERFACE
    FUNCTION some_actual_function_3( Desc_Handle )
        INTEGER :: some_actual_function_3
        TYPE(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
    END FUNCTION some_actual_function_3
END INTERFACE DftiFreeDescriptor

/* C prototype */
long DftiFreeDescriptor( DFTI_DESCRIPTOR_HANDLE * );
```

DFT Computation

There are two functions in this category: compute the forward transform, and compute the backward transform.

ComputeForward

Computes the forward DFT.

Usage

```
! Fortran
Status = DftiComputeForward( Desc_Handle, X_inout )
Status = DftiComputeForward( Desc_Handle, X_in, X_out )
Status = DftiComputeForward( Desc_Handle, X_inout, Y_inout )
Status = DftiComputeForward( Desc_Handle, X_in, Y_in, X_out, Y_out )
/* C */
status = DftiComputeForward( desc_handle, x_inout );
status = DftiComputeForward( desc_handle, x_in, x_out );
status = DftiComputeForward( desc_handle, x_inout, y_inout );
status = DftiComputeForward( desc_handle, x_in, y_in, x_out, y_out );
```

Discussion

As soon as a descriptor is configured and committed successfully, actual computation of DFT can be performed. The `DftiComputeForward` function computes the forward DFT. By default, this is the transform using the factor $e^{-i2\pi/n}$ (instead of the one with a positive sign). Because of the flexibility in configuration, input data can be represented in various ways as well as output result can be placed differently. Consequently, the number of input parameters as well as their type vary. This variation is accommodated by the generic function facility of Fortran 95. Data and result parameters are all declared as assumed-size rank-1 array `DIMENSION(0:*)`.

The function returns `DFTI_NO_ERROR` when completes successfully. See

[“Status Checking Functions”](#) for more information on returned status.

Interface and prototype

```
//Fortran interface.
INTERFACE DftiComputeFoward
//Note that the body provided here is to illustrate the different
//argument list and types of dummy arguments. The interface
//does not guarantee what the actual function names are.
//Users can only rely on the function name following the
//keyword INTERFACE
// One argument single precision complex
FUNCTION some_actual_function_4_C( Desc_Handle, X )
    INTEGER :: some_actual_function_4_C
    TYPE(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
    COMPLEX, INTENT(INOUT) :: X(*)
END FUNCTION some_actual_function_4_C
// One argument double precision complex
FUNCTION some_actual_function_4_Z( Desc_Handle, X )
    INTEGER :: some_actual_function_4_Z
    TYPE(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
    COMPLEX (Kind((0D0,0D0))), INTENT(INOUT) :: X(*)
END FUNCTION some_actual_function_4_Z
// One argument single precision real
FUNCTION some_actual_function_4_R( Desc_Handle, X )
    INTEGER :: some_actual_function_4_R
    TYPE(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
    REAL, INTENT(INOUT) :: X(*)
END FUNCTION some_actual_function_4_R
// One argument double precision real
...
// Two argument single precision complex
...
...
// Four argument double precision real
FUNCTION some_actual_function_4_DDDD( Desc_Handle, X1_In, X2_In,
                                     Y1_Out, Y2_Out )
```

```

    INTEGER :: some_actual_function_4_DDDD
    TYPE(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
    REAL (Kind(OD0)), INTENT(IN) :: X1_In(*), X2_In(*)
    REAL (Kind(OD0)), INTENT(OUT) :: Y1_Out(*), Y2_Out(*)
  END FUNCTION some_actual_function_4_DDDD
END INTERFACE DftiComputeFoward

/* C prototype */
long DftiComputeForward( DFTI_DESCRIPTOR_HANDLE,
                        void *,
                        ... );

```

The implementations of DFT interface expect the data be treated as data stored linearly in memory with a regular "stride" pattern (discussed more fully in [“Strides”](#), see also [3]). The function expects the starting address of the first element. Hence we use the assume-size declaration in Fortran.

The descriptor by itself contains sufficient information to determine exactly how many arguments and of what type should be present. The implementation could use this information to check against possible input inconsistency.

ComputeBackward

Computes the backward DFT.

Usage

```

! Fortran
Status = DftiComputeBackward( Desc_Handle, X_inout )
Status = DftiComputeBackward( Desc_Handle, X_in, X_out )
Status = DftiComputeBackward( Desc_Handle, X_inout, Y_inout )
Status = DftiComputeBackward( Desc_Handle, X_in, Y_in, X_out, Y_out )
/* C */
status = DftiComputeBackward( desc_handle, x_inout );
status = DftiComputeBackward( desc_handle, x_in, x_out );

```

```
status = DftiComputeBackward( desc_handle, x_inout, y_inout );
status = DftiComputeBackward( desc_handle, x_in, y_in, x_out, y_out );
```

Discussion

As soon as a descriptor is configured and committed successfully, actual computation of DFT can be performed. The `DftiComputeBackward` function computes the backward DFT. By default, this is the transform using the factor $e^{i2\pi/n}$ (instead of the one with a negative sign). Because of the flexibility in configuration, input data can be represented in various ways as well as output result can be placed differently. Consequently, the number of input parameters as well as their type vary. This variation is accommodated by the generic function facility of Fortran 95. Data and result parameters are all declared as assumed-size rank-1 array `DIMENSION(0:*)`.

The function returns `DFTI_NO_ERROR` when completes successfully. See [“Status Checking Functions”](#) for more information on returned status.

Interface and prototype

```
//Fortran interface.
INTERFACE DftiComputeBackward
//Note that the body provided here is to illustrate the different
//argument list and types of dummy arguments. The interface
//does not guarantee what the actual function names are.
//Users can only rely on the function name following the
//keyword INTERFACE
// One argument single precision complex
FUNCTION some_actual_function_5_C( Desc_Handle, X )
  INTEGER :: some_actual_function_5_C
  TYPE(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
  COMPLEX, INTENT(INOUT) :: X(*)
END FUNCTION some_actual_function_5_C
// One argument double precision complex
FUNCTION some_actual_function_5_Z( Desc_Handle, X )
  INTEGER :: some_actual_function_5_Z
  TYPE(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
  COMPLEX (Kind((0D0,0D0))), INTENT(INOUT) :: X(*)
```

```
END FUNCTION some_actual_function_5_Z
// One argument single precision real
FUNCTION some_actual_function_5_R( Desc_Handle, X )
  INTEGER :: some_actual_function_5_R
  TYPE(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
  REAL, INTENT(INOUT) :: X(*)
END FUNCTION some_actual_function_5_R
// One argument double precision real
...
// Two argument single precision complex
...
...
// Four argument double precision real
FUNCTION some_actual_function_5_DDDD( Desc_Handle, X1_In, X2_In,
                                     Y1_Out, Y2_Out )
  INTEGER :: some_actual_function_5_DDDD
  TYPE(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
  REAL (Kind(OD0)), INTENT(IN) :: X1_In(*), X2_In(*)
  REAL (Kind(OD0)), INTENT(OUT) :: Y1_Out(*), Y2_Out(*)
END FUNCTION some_actual_function_5_DDDD
END INTERFACE DftiComputeBackward

/* C prototype */
long DftiComputeBackward( DFTI_DESCRIPTOR_HANDLE,
                          void *,
                          ... );
```

The implementations of DFT interface expect the data be treated as data stored linearly in memory with a regular "stride" pattern (discussed more fully in [“Strides”](#), see also [3]). The function expects the starting address of the first element. Hence we use the assume-size declaration in Fortran.

The descriptor by itself contains sufficient information to determine exactly how many arguments and of what type should be present. The implementation could use this information to check against possible input inconsistency.

Descriptor Configuration

There are two functions in this category: the value setting function `DftiSetValue` sets one particular configuration parameter to an appropriate value, and the value getting function `DftiGetValue` reads the values of one particular configuration parameter. While all configuration parameters are readable, a few of them cannot be set by user. Some of these contain fixed information of a particular implementation such as version number, or dynamic information, but nevertheless are derived by the implementation during execution of one of the functions.

Table 9-3 **Settable Configuration Parameters**

Named Constants	Value Type	Comments
<i>Most common configurations, no default, must be set explicitly</i>		
<code>DFTI_PRECISION</code>	Named constant	Precision of computation
<code>DFTI_FORWARD_DOMAIN</code>	Named constant	Domain for the forward transform
<code>DFTI_DIMENSION</code>	Integer scalar	Dimension of the transform
<code>DFTI_LENGTHS</code>	Integer scalar/array	Lengths of each dimension
<i>Common configurations including multiple transform and data representation</i>		
<code>DFTI_NUMBER_OF_TRANSFORMS</code>	Integer scalar	For multiple number of transforms
<code>DFTI_FORWARD_SIGN</code>	Named constant	The definition for forward transform
<code>DFTI_FORWARD_SCALE</code>	Floating-point scalar	Scale factor for forward transform
<code>DFTI_BACKWARD_SCALE</code>	Floating-point scalar	Scale factor for backward transform
<code>DFTI_PLACEMENT</code>	Named constant	Placement of the computation result
<code>DFTI_COMPLEX_STORAGE</code>	Named constant	Storage method, complex domain data
<code>DFTI_REAL_STORAGE</code>	Named constant	Storage method, real domain data
<code>DFTI_CONJUGATE_EVEN_STORAGE</code>	Named constant	Storage method, conjugate even domain data
<code>DFTI_DESCRIPTOR_NAME</code>	Character string	No longer than <code>DFTI_MAX_NAME_LENGTH</code>
<code>DFTI_PACKED_FORMAT</code>	Named constant	Packed format, real domain data

Table 9-3 **Settable Configuration Parameters** (continued)

Named Constants	Value Type	Comments
<i>Configurations regarding stride of data</i>		
<code>DFTI_INPUT_DISTANCE</code>	Integer scalar	Multiple transforms, distance of first elements
<code>DFTI_OUTPUT_DISTANCE</code>	Integer scalar	Multiple transforms, distance of first elements
<code>DFTI_INPUT_STRIDES</code>	Integer array	Stride information of input data
<code>DFTI_OUTPUT_STRIDES</code>	Integer array	Stride information of output data
<i>Advanced configuration</i>		
<code>DFTI_INITIALIZATION_EFFORT</code>	Named constant	Dynamic search for computation method
<code>DFTI_ORDERING</code>	Named constant	Scrambling of data order
<code>DFTI_WORKSPACE</code>	Named constant	Computation without auxiliary storage
<code>DFTI_TRANSPOSE</code>	Named constant	Scrambling of dimension

Each of these configuration parameters is identified by a named constant in the `MKL_DFTI` module. In C, these named constants have the enumeration type `DFTI_CONFIG_PARAM`. The list of configuration parameters whose values can be set by user is given in [Table 9-3](#); the list of configuration parameters that are read-only is given in [Table 9-4](#). All parameters are readable. Most of these parameters are self-explanatory, while some others are discussed more fully in the description of the relevant functions.

Table 9-4 **Read-Only Configuration Parameters**

Named Constants	Value Type	Comments
<code>DFTI_COMMIT_STATUS</code>	Name constant	Whether descriptor has been committed
<code>DFTI_VERSION</code>	String	Intel MKL library version number
<code>DFTI_FORWARD_ORDERING</code>	Integer pointer	Pointer to an integer array (see “Ordering”)
<code>DFTI_BACKWARD_ORDERING</code>	Integer pointer	Pointer to an integer array (see “Ordering”)

The configuration parameters are set by various values. Some of these values are specified by native data types such as an integer value (for example, number of simultaneous transforms requested), or a single-precision number (for example, the scale factor one would like to apply on a forward transform).

Other configuration values are discrete in nature (for example, the domain of the forward transform) and are thus provided in the DFTI module as named constants. In C, these named constants have the enumeration type `DFTI_CONFIG_VALUE`. The complete list of named constants used for this kind of configuration values is given in [Table 9-5](#).

Table 9-5 **Named Constant Configuration Values**

Named Constant	Comments
<code>DFTI_SINGLE</code>	Single precision
<code>DFTI_DOUBLE</code>	Double precision
<code>DFTI_COMPLEX</code>	Complex domain
<code>DFTI_REAL</code>	Real domain
<code>DFTI_CONJUGATE_EVEN</code>	Conjugate even domain
<code>DFTI_NEGATIVE</code>	Sign used to define the forward transform
<code>DFTI_POSITIVE</code>	Sign used to define the forward transform
<code>DFTI_INPLACE</code>	Output overwrites input
<code>DFTI_NOT_INPLACE</code>	Output does not overwrite input
<code>DFTI_COMPLEX_COMPLEX</code>	Storage method (see “Storage schemes”)
<code>DFTI_REAL_REAL</code>	Storage method (see “Storage schemes”)
<code>DFTI_COMPLEX_REAL</code>	Storage method (see “Storage schemes”)
<code>DFTI_REAL_COMPLEX</code>	Storage method (see “Storage schemes”)
<code>DFTI_HIGH</code>	A high setting, related to initialization effort
<code>DFTI_MEDIUM</code>	A medium setting, related to initialization effort
<code>DFTI_LOW</code>	A low setting, related to initialization effort
<code>DFTI_COMMITTED</code>	Committal status of a descriptor
<code>DFTI_UNCOMMITTED</code>	Committal status of a descriptor
<code>DFTI_ORDERED</code>	Data ordered in both forward and backward domains
<code>DFTI_BACKWARD_SCRAMBLED</code>	Data scrambled in backward domain (by forward transform)

Table 9-5 **Named Constant Configuration Values** (continued)

Named Constant	Comments
<code>DFTI_FORWARD_SCRAMBLED</code>	Data scrambled in forward domain (by backward transform)
<code>DFTI_ALLOW</code>	Allow certain request or usage if useful
<code>DFTI_AVOID</code>	Avoid certain request or usage if practical
<code>DFTI_NONE</code>	Used to specify no transposition
<code>DFTI_CCS_FORMAT</code>	Packed format, real data (see “Packed formats”)
<code>DFTI_PACK_FORMAT</code>	Packed format, real data (see “Packed formats”)
<code>DFTI_PERM_FORMAT</code>	Packed format, real data (see “Packed formats”)
<code>DFTI_VERSION_LENGTH</code>	Number of characters for library version length
<code>DFTI_MAX_NAME_LENGTH</code>	Maximum descriptor name length
<code>DFTI_MAX_MESSAGE_LENGTH</code>	Maximum status message length

[Table 9-6](#) lists the possible values for those configuration parameters that are discrete in nature.

Table 9-6 **Settings for Discrete Configuration Parameters**

Named Constant	Possible Values
<code>DFTI_PRECISION</code>	<code>DFTI_SINGLE</code> , or <code>DFTI_DOUBLE</code> (no default)
<code>DFTI_FORWARD_DOMAIN</code>	<code>DFTI_COMPLEX</code> , or <code>DFTI_REAL</code> , or <code>DFTI_CONJUGATE_EVEN</code> (no default)
<code>DFTI_FORWARD_SIGN</code>	<code>DFTI_NEGATIVE</code> (default), or <code>DFTI_POSITIVE</code>
<code>DFTI_PLACEMENT</code>	<code>DFTI_INPLACE</code> (default), or <code>DFTI_NOT_INPLACE</code>
<code>DFTI_COMPLEX_STORAGE</code>	<code>DFTI_COMPLEX_COMPLEX</code> (default), or <code>DFTI_COMPLEX_REAL</code> , or <code>DFTI_REAL_REAL</code>
<code>DFTI_REAL_STORAGE</code>	<code>DFTI_REAL_REAL</code> (default), or <code>DFTI_REAL_COMPLEX</code>
<code>DFTI_CONJUGATE_EVEN_STORAGE</code>	<code>DFTI_COMPLEX_COMPLEX</code> , or

Table 9-6 Settings for Discrete Configuration Parameters (continued)

Named Constant	Possible Values
	DFTI_COMPLEX_REAL (default), or
	DFTI_REAL_REAL (1-D transform only)
DFTI_PACKED_FORMAT	DFTI_CCS_FORMAT (default) or,
	DFTI_PACK_FORMAT or,
	DFTI_PERM_FORMAT

[Table 9-7](#) lists the default values of the settable configuration parameters.

Table 9-7 Default Configuration Values of Settable Parameters

Named Constants	Default Value
DFTI_NUMBER_OF_TRANSFORMS	1
DFTI_FORWARD_SIGN	DFTI_NEGATIVE
DFTI_FORWARD_SCALE	1.0
DFTI_BACKWARD_SCALE	1.0
DFTI_PLACEMENT	DFTI_INPLACE
DFTI_COMPLEX_STORAGE	DFTI_COMPLEX_COMPLEX
DFTI_REAL_STORAGE	DFTI_REAL_REAL
DFTI_CONJUGATE_EVEN_STORAGE	DFTI_COMPLEX_REAL
DFTI_PACKED_FORMAT	DFTI_CCS_FORMAT
DFTI_DESCRIPTOR_NAME	no name, string of zero length
DFTI_INPUT_DISTANCE	0
DFTI_OUTPUT_DISTANCE	0
DFTI_INPUT_STRIDES	Tightly packed according to dimension, lengths, and storage
DFTI_OUTPUT_STRIDES	Same as above. See “Strides” for details
DFTI_INITIALIZATION_EFFORT	DFTI_MEDIUM
DFTI_ORDERING	DFTI_ORDERED
DFTI_WORKSPACE	DFTI_ALLOW
DFTI_TRANSPOSE	DFTI_NONE

SetValue

Sets one particular configuration parameter with the specified configuration value.

Usage

```
! Fortran
Status = DftiSetValue( Desc_Handle,      &
                      Config_Param,     &
                      Config_Val )

/* C */
status = DftiSetValue( desc_handle,
                      config_param,
                      config_val );
```

Discussion

This function sets one particular configuration parameter with the specified configuration value. The configuration parameter is one of the named constants listed in [Table 9-3](#), and the configuration value is the corresponding appropriate type, which can be a named constant or a native type. See [“Configuration Settings”](#) for details of the meaning of the setting.

The function returns `DFTI_NO_ERROR` when completes successfully. See [“Status Checking Functions”](#) for more information on returned status.

Interface and prototype

```
! Fortran interface
INTERFACE DftiSetValue
//Note that the body provided here is to illustrate the different
//argument list and types of dummy arguments. The interface
//does not guarantee what the actual function names are.
//Users can only rely on the function name following the
```

```
//keyword INTERFACE
  FUNCTION some_actual_function_6_INTVAL( Desc_Handle, Config_Param,
INTVAL )
    INTEGER :: some_actual_function_6_INTVAL
    Type(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
    INTEGER, INTENT(IN) :: Config_Param
    INTEGER, INTENT(IN) :: INTVAL
  END FUNCTION some_actual_function_6_INTVAL

  FUNCTION some_actual_function_6_SGLVAL( Desc_Handle, Config_Param,
SGLVAL )
    INTEGER :: some_actual_function_6_SGLVAL
    Type(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
    INTEGER, INTENT(IN) :: Config_Param
    REAL, INTENT(IN) :: SGLVAL
  END FUNCTION some_actual_function_6_SGLVAL

  FUNCTION some_actual_function_6_DBLVAL( Desc_Handle, Config_Param,
DBLVAL )
    INTEGER :: some_actual_function_6_DBLVAL
    Type(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
    INTEGER, INTENT(IN) :: Config_Param
    REAL (KIND(OD0)), INTENT(IN) :: DBLVAL
  END FUNCTION some_actual_function_6_DBLVAL

  FUNCTION some_actual_function_6_INTVEC( Desc_Handle, Config_Param,
INTVEC )
    INTEGER :: some_actual_function_6_INTVEC
    Type(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
    INTEGER, INTENT(IN) :: Config_Param
    INTEGER, INTENT(IN) :: INTVEC(*)
  END FUNCTION some_actual_function_6_INTVEC

  FUNCTION some_actual_function_6_CHARS( Desc_Handle, Config_Param,
CHARS )
    INTEGER :: some_actual_function_6_CHARS
    Type(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
```

```

    INTEGER, INTENT(IN) :: Config_Param
    CHARACTER(*), INTENT(IN) :: CHARS
    END FUNCTION some_actual_function_6_CHARS
END INTERFACE DftiSetValue

```

```

/* C prototype */
long DftiSetValue( DFTI_DESCRIPTOR_HANDLE,
                  DFTI_CONFIG_PARAM ,
                  ... );

```

GetValue

Gets the configuration value of one particular configuration parameter.

Usage

```

! Fortran
Status = DftiGetValue( Desc_Handle,      &
                      Config_Param,     &
                      Config_Val )

/* C */
status = DftiGetValue( desc_handle,
                      config_param,
                      &config_val );

```

Discussion

This function gets the configuration value of one particular configuration parameter. The configuration parameter is one of the named constants listed in [Table 9-3](#) and [Table 9-4](#), and the configuration value is the corresponding appropriate type, which can be a named constant or a native type.

The function returns `DFTI_NO_ERROR` when completes successfully. See [“Status Checking Functions”](#) for more information on returned status.

Interface and prototype

```

! Fortran interface
INTERFACE DftiGetValue
//Note that the body provided here is to illustrate the different
//argument list and types of dummy arguments. The interface
//does not guarantee what the actual function names are.
//Users can only rely on the function name following the
//keyword INTERFACE
    FUNCTION some_actual_function_7_INTVAL( Desc_Handle, Config_Param,
INTVAL )
        INTEGER :: some_actual_function_7_INTVAL
        Type(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
        INTEGER, INTENT(IN) :: Config_Param
        INTEGER, INTENT(OUT) :: INTVAL
    END FUNCTION DFTI_GET_VALUE_INTVAL

    FUNCTION some_actual_function_7_SGLVAL( Desc_Handle, Config_Param,
SGLVAL )
        INTEGER :: some_actual_function_7_SGLVAL
        Type(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
        INTEGER, INTENT(IN) :: Config_Param
        REAL, INTENT(OUT) :: SGLVAL
    END FUNCTION some_actual_function_7_SGLVAL

    FUNCTION some_actual_function_7_DBLVAL( Desc_Handle, Config_Param,
DBLVAL )
        INTEGER :: some_actual_function_7_DBLVAL
        Type(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
        INTEGER, INTENT(IN) :: Config_Param
        REAL (KIND(OD0)), INTENT(OUT) :: DBLVAL
    END FUNCTION some_actual_function_7_DBLVAL

    FUNCTION some_actual_function_7_INTVEC( Desc_Handle, Config_Param,
INTVEC )
        INTEGER :: some_actual_function_7_INTVEC
        Type(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
        INTEGER, INTENT(IN) :: Config_Param

```

```

    INTEGER, INTENT(OUT) :: INTVEC(*)
END FUNCTION some_actual_function_7_INTVEC

FUNCTION some_actual_function_7_INTPNT( Desc_Handle, Config_Param,
INTPNT )
    INTEGER :: some_actual_function_7_INTPNT
    Type(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
    INTEGER, INTENT(IN) :: Config_Param
    INTEGER, DIMENSION(*), POINTER :: INTPNT
END FUNCTION some_actual_function_7_INTPNT

FUNCTION some_actual_function_7_CHARS( Desc_Handle, Config_Param,
CHARS )
    INTEGER :: some_actual_function_7_CHARS
    Type(DFTI_DESCRIPTOR), POINTER :: Desc_Handle
    INTEGER, INTENT(IN) :: Config_Param
    CHARACTER(*), INTENT(OUT):: CHARS
END FUNCTION some_actual_function_7_CHARS
END INTERFACE DftiGetValue

/* C prototype */
long DftiGetValue( DFTI_DESCRIPTOR_HANDLE,
                  DFTI_CONFIG_PARAM ,
                  ... );

```

Configuration Settings

Precision of transform

The configuration parameter `DFTI_PRECISION` denotes the floating-point precision in which the transform is to be carried out. A setting of `DFTI_SINGLE` stands for single precision, and a setting of `DFTI_DOUBLE` stands for double precision. The data is meant to be presented in this precision; the computation will be carried out in this precision; and the result will be delivered in this precision. This is one of the four settable configuration parameters that do not have default values. The user must set them explicitly, most conveniently at the call to descriptor creation function [DftiCreateDescriptor](#).

Forward domain of transform

The general form of the discrete Fourier transform is

$$Z_{k_1, k_2, \dots, k_d} = \sigma \times \sum_{j_d=0}^{n_d-1} \dots \sum_{j_2=0}^{n_2-1} \sum_{j_1=0}^{n_1-1} w_{j_1, j_2, \dots, j_d} \exp \left(\delta i 2\pi \sum_{l=1}^d j_l k_l / n_l \right) \quad (7.1)$$

for $k_l = 0, \pm 1, \pm 2, \dots$, where σ is an arbitrary real-valued scale factor and $\delta = \pm 1$. By default, the forward transform is defined by $\sigma = 1$ and $\delta = -1$. In most common situations, the domain of the forward transform, that is, the set where the input (periodic) sequence $\{w_{j_1, j_2, \dots, j_d}\}$ belongs, can be either the set of complex-valued sequences, real-valued sequences, and complex-valued conjugate even sequences. The configuration parameter `DFTI_FORWARD_DOMAIN` indicates the domain for the forward transform. Note that this implicitly specifies the domain for the backward transform because of mathematical property of the DFT. See [Table 9-8](#) for details.

Table 9-8 Correspondence of Forward and Backward Domain

	Forward Domain	Implied Backward Domain
Complex	<code>(DFTI_COMPLEX)</code>	Complex
Real	<code>(DFTI_REAL)</code>	Conjugate Even
Conjugate Even	<code>(DFTI_CONJUGATE_EVEN)</code>	Real

On transforms in the real domain, some software packages only offer one "real-to-complex" transform. This in essence omits the conjugate even domain for the forward transform. The forward domain configuration parameter `DFTI_FORWARD_DOMAIN` is the second of four configuration parameters without default value.

Transform dimension and lengths

The dimension of the transform is a positive integer value represented in an integer scalar of type `Integer`. For one-dimensional transform, the transform length is specified by a positive integer value represented in an integer scalar of type `Integer`. For multi-dimensional (≥ 2) transform, the

lengths of each of the dimension is supplied in an integer array.

`DFTI_DIMENSION` and `DFTI_LENGTHS` are the remaining two of four configuration parameters without default.

As mentioned, these four configuration parameters do not have default value. They are most conveniently set at the descriptor creation function. For dimension and length configuration, they can only be set in the descriptor creation function, and not by the function `DftiSetValue`. The other two configuration values can be changed through the function `DftiSetValue`, although this is not deemed common.



CAUTION. *Changing the dimension and length would likely render the stride value inappropriate. Unless certain of otherwise, the user is advised to reconfigure the stride (see [“Strides”](#)).*

Number of transforms

In some situations, the user may need to perform a number of DFT transforms of the same dimension and lengths. The most common situation would be to transform a number of one-dimensional data of the same length. This parameter has the default value of 1, and can be set to positive integer value by an `Integer` data type in Fortran and `long` data type in C. Data sets have no common elements. The distance parameter is obligatory if multiple number is more than one.

Sign and scale

The general form of the discrete Fourier transform is given by (7.1), for $k_1 = 0, \pm 1, \pm 2, \dots$, where σ is an arbitrary real-valued scale factor and $\delta = \pm 1$. By default, the forward transform is defined by $\sigma = 1$ and $\delta = -1$, and the backward transform is defined by $\sigma = 1$ and $\delta = 1$. The user can change the definition of forward transform via setting the sign δ to be `DFTI_NEGATIVE` (default) or `DFTI_POSITIVE`. The sign of the backward transform is implicitly defined to be the negative of the sign for the forward transform.

The forward transform and backward transform are each associated with a scale factor σ of its own with default value of 1. The user can set one or both of them via the two configuration parameters `DFTI_FORWARD_SCALE` and `DFTI_BACKWARD_SCALE`. For example, for a one-dimensional transform of length n , one can use the default scale of 1 for the forward transform while setting the scale factor for backward transform to be $1/n$, making the backward transform the inverse of the forward transform.

The scale factor configuration parameter should be set by a real floating-point data type of the same precision as the value for `DFTI_PRECISION`.



NOTE. *The sign configuration is not supported. The forward transform is defined as $\delta = -1$.*

Placement of result

By default, the computational functions overwrite the input data with the output result. That is, the default setting of the configuration parameter `DFTI_PLACEMENT` is `DFTI_INPLACE`. The user can change that by setting it to `DFTI_NOT_INPLACE`.

Packed formats

The result of the forward transform (i.e. in the frequency-domain) of real data is represented in several possible packed formats: *Pack*, *Perm*, or *CCS*. The data can be packed due to the symmetry property of the DFT transform of a real data.

The *CCS* format stores the values of the first half of the output complex signal resulted from the forward DFT. Note that the signal stored in *CCS* format is one complex element longer. In *CCS* format, the output samples of the DFT are arranged as shown in [Table 9-9](#).

The **pack** format is a compact representation of a complex conjugate-symmetric sequence. The disadvantage of this format is that it is not the natural format used by the real DFT algorithms (“natural” in the sense that array is natural for complex DFTs). In **pack** format, the output samples of the DFT are arranged as shown in [Table 9-9](#).

The **perm** format is an arbitrary permutation of the **pack** format for even lengths and one is the same as the **pack** format for odd lengths. In **perm** format, the output samples of the DFT are arranged as shown in [Table 9-9](#).

Table 9-9 Packed Format Output Samples

For (N = S*2)										
DFT Real	0	1	2	3	...	N-2	N-1	N	N+1	
	0	1						(2S-1)	2S	(2S+1)
CCS	R_0	0	R_1	I_1	...	$R_{N/2-1}$	$I_{N/2-1}$	$R_{N/2}$	$I_{N/2}$	0
Pack	R_0	R_1	I_1	R_2	...	$I_{N/2-1}$	$R_{N/2}$			
Perm	R_0	$R_{N/2}$	R_1	I_1	...	$R_{N/2-1}$	$I_{N/2-1}$			

For (N = S*2 + 1)											
DFT Real	0	1	2	3	...	N-4	N-3	N-2	N-1	N	N+1
	0	1							2S	(2S+1)	
CCS	R_0	0	R_1	I_1	...	I_{S-2}	R_{S-1}	I_{S-1}	R_S	I_S	
Pack	R_0	R_1	I_1	R_2	...	R_{S-1}	I_{S-1}	R_{S-1}	I_S		
Perm	R_0	R_1	I_1	R_2	...	R_{S-1}	I_{S-1}	R_{S-1}	I_S		

See also [Table 9-10](#) and [Table 9-11](#).

Storage schemes

For each of the three domains **DFTI_COMPLEX**, **DFTI_REAL**, and **DFTI_CONJUGATE_EVEN** (for the forward as well as the backward operator), a subset of the four storage schemes **DFTI_COMPLEX_COMPLEX**, **DFTI_COMPLEX_REAL**, **DFTI_REAL_COMPLEX**, and **DFTI_REAL_REAL**. Specific examples are presented here to illustrate the storage schemes. See the document [\[3\]](#) for the rationale behind this definition of the storage schemes.



NOTE. *The data is stored in the Fortran style only, that is, the real and imaginary parts are stored side by side.*

Storage scheme for complex domain

This setting is recorded in the configuration parameter `DFTI_COMPLEX_STORAGE`. The three values that can be set are `DFTI_COMPLEX_COMPLEX`, `DFTI_COMPLEX_REAL`, and `DFTI_REAL_REAL`. Consider a one-dimensional n -length transform of the form

$$z_k = \sum_{j=0}^{n-1} w_j e^{-i2\pi jk/n}, \quad w_j, z_k \in \mathbb{C}.$$

Assume the stride has default value (unit stride) and `DFTI_PLACEMENT` has the default in-place setting.

1. **`DFTI_COMPLEX_COMPLEX` storage scheme.** A typical usage will be as follows.

```
COMPLEX :: X(0:n-1)
...some other code...
Status = DftiComputeForward( Desc_Handle, X )
```

On input,

$$X(j) = w_j, \quad j = 0, 1, \dots, n-1.$$

On output,

$$X(k) = z_k, \quad k = 0, 1, \dots, n-1.$$

2. **`DFTI_COMPLEX_REAL` storage scheme.** A typical usage will be as follows.

```
REAL :: X(0:2*n-1)
...some other code...
Status = DftiComputeForward( Desc_Handle, X )
```

On input,

$$X(2*j) = \text{Re}(w_j), X(2*j+1) = \text{Im}(w_j), j = 0, 1, \dots, n-1.$$

On output,

$$X(2*k) = \text{Re}(z_k), X(2*k+1) = \text{Im}(z_k), k = 0, 1, \dots, n-1.$$

The notations $\text{Re}(w_j)$ and $\text{Im}(w_j)$ are the real and imaginary parts of the complex number w_j .

3. DFTI_REAL_REAL storage scheme. A typical usage will be as follows.

```
REAL :: X(0:n-1), Y(0:n-1)
...some other code...
Status = DftiComputeForward( Desc_Handle, X, Y )
```

On input,

$$X(j) = \text{Re}(w_j), Y(j) = \text{Im}(w_j), j = 0, 1, \dots, n-1.$$

On output,

$$X(k) = \text{Re}(z_k), Y(k) = \text{Im}(z_k), k = 0, 1, \dots, n-1.$$

Storage scheme for the real and conjugate even domains

This setting for the storage schemes for these domains are recorded in the configuration parameters `DFTI_REAL_STORAGE` and `DFTI_CONJUGATE_EVEN`. Since a forward real domain corresponds to a conjugate even backward domain, they are considered together. The example uses a one-dimensional real to conjugate even transform. In-place computation is assumed whenever possible (that is, when the input data type matches with the output data type).

Consider a one-dimensional n -length transform of the form

$$z_k = \sum_{j=0}^{n-1} w_j e^{-i2\pi jk/n}, \quad w_j \in \mathbf{R}, z_k \in \mathbf{C}.$$

There is a symmetry:

For even n : $z(n/2+i) = \text{conjg}(z(n/2-i))$, $1 \leq i \leq n/2 - 1$, and moreover $z(0)$ and $z(n/2)$ are real values.

For odd n : $z(m+i) = \text{conjg}(z(m-i+1))$, $1 \leq i \leq m$, and moreover $z(0)$ is real value.

$$m = \text{floor}(n/2).$$

Table 9-10 Comparison of the Storage Effects of Complex-to-Complex and Real-to-Complex DFTs for Forward Transform

Input Vectors			Output Vectors					
Complex DFT		Real DFT	complex DFT			real DFT		
Complex Data		Real Data	Complex Data		Real Data			
Real	Imaginary		Real	Imaginary	CCS	Pack	Perm	
w0	0.000000	w0	z0	0.000000	z0	z0	z0	
w1	0.000000	w1	Re(z1)	Im(z1)	0.000000	Re(z1)	z4	
w2	0.000000	w2	Re(z2)	Im(z2)	Re(z1)	Im(z1)	Re(z1)	
w3	0.000000	w3	Re(z3)	Im(z3)	Im(z1)	Re(z2)	Im(z1)	
w4	0.000000	w4	z4	0.000000	Re(z2)	Im(z2)	Re(z2)	
w5	0.000000	w5	Re(z3)	-Im(z3)	Im(z2)	Re(z3)	Im(z2)	
w6	0.000000	w6	Re(z2)	-Im(z2)	Re(z3)	Im(z3)	Re(z3)	
w7	0.000000	w7	Re(z1)	-Im(z1)	Im(z3)	z4	Im(z3)	
					z4			
					0.000000			

Input Vectors			Output Vectors					
Complex DFT		Real DFT	complex DFT			real DFT		
Complex Data		Real Data	Complex Data		Real Data			
Real	Imaginary		Real	Imaginary	CCS	Pack	Perm	
w0	0.000000	w0	z0	0.000000	z0	z0	z0	
w1	0.000000	w1	Re(z1)	Im(z1)	0.000000	Re(z1)	Re(z1)	
w2	0.000000	w2	Re(z2)	Im(z2)	Re(z1)	Im(z1)	Im(z1)	

N=7

Input Vectors			Output Vectors				
Complex DFT		Real DFT	complex DFT		real DFT		
Complex Data		Real Data	Complex Data		Real Data		
Real	Imaginary		Real	Imaginary	CCS	Pack	Perm
w3	0.000000	w3	Re(z3)	Im(z3)	Im(z1)	Re(z2)	Re(z2)
w4	0.000000	w4	Re(z3)	-Im(z3)	Re(z2)	Im(z2)	Im(z2)
w5	0.000000	w5	Re(z2)	-Im(z2)	Im(z2)	Re(z3)	Re(z3)
w6	0.000000	w6	Re(z1)	-Im(z1)	Re(z3)	Im(z3)	Im(z3)
					Im(z3)		

Table 9-11 Comparison of the Storage Effects of Complex-to-Complex and Complex-to-Real DFTs for Backward Transform

N=8								
Output Vectors			Input Vectors					
Complex DFT		Real DFT	complex DFT					
Complex Data		Real Data	Complex Data					
Real	Imaginary		Real	Imaginary	CCS	Pack	Perm	
w0	0.000000	w0	z0	0.000000	z0	z0	z0	
w1	0.000000	w1	Re(z1)	Im(z1)	0.000000	Re(z1)	z4	
w2	0.000000	w2	Re(z2)	Im(z2)	Re(z1)	Im(z1)	Re(z1)	
w3	0.000000	w3	Re(z3)	Im(z3)	Im(z1)	Re(z2)	Im(z1)	
w4	0.000000	w4	z4		Re(z2)	Im(z2)	Re(z2)	
w5	0.000000	w5	Re(z3)	-Im(z3)	Im(z2)	Re(z3)	Im(z2)	
w6	0.000000	w6	Re(z2)	-Im(z2)	Re(z3)	Im(z3)	Re(z3)	
w7	0.000000	w7	Re(z1)	-Im(z1)	Im(z3)	z4	Im(z3)	
					z4			
					0.000000			

N=7								
Output Vectors			Input Vectors					
Complex DFT		Real DFT	complex DFT			real DFT		
Complex Data		Real Data	Complex Data			Real Data		
Real	Imaginary		Real	Imaginary	CCS	Pack	Perm	
w0	0.000000	w0	z0	0.000000	z0	z0	z0	
w1	0.000000	w1	Re(z1)	Im(z1)	0.000000	Re(z1)	Re(z1)	
w2	0.000000	w2	Re(z2)	Im(z2)	Re(z1)	Im(z1)	Im(z1)	
w3	0.000000	w3	Re(z3)	Im(z3)	Im(z1)	Re(z2)	Re(z2)	
w4	0.000000	w4	Re(z3)	-Im(z3)	Re(z2)	Im(z2)	Im(z2)	
w5	0.000000	w5	Re(z2)	-Im(z2)	Im(z2)	Re(z3)	Re(z3)	

N=7

Output Vectors			Input Vectors				
Complex DFT	Real DFT	Real DFT	complex DFT	real DFT			
Complex Data		Real Data	Complex Data		Real Data		
Real	Imaginary		Real	Imaginary	CCS	Pack	Perm
w6	0.000000	w6	Re(z1)	-Im(z1)	Re(z3)	Im(z3)	Im(z3)
					Im(z3)		

Assume that the stride has the default value (unit stride).

This complex conjugate-symmetric vector can be stored in the complex array of size $m+1$ or in the real array of size $2m+2$ or $2m$ depending on packed format.

1. **DFTI_REAL_REAL** for real domain, **DFTI_COMPLEX_COMPLEX** for conjugate even domain. A typical usage will be as follows.

```
// m = floor( n/2 )
REAL :: X(0:n-1)
COMPLEX :: Y(0:m)
...some other code...
...out of place transform...
Status = DftiComputeForward( Desc_Handle, X, Y )
```

On input,

$$X(j) = w_j, j = 0, 1, \dots, n-1.$$

On output,

$$Y(k) = z_k, k = 0, 1, \dots, m.$$

2. **DFTI_REAL_REAL** for real domain, **DFTI_COMPLEX_REAL** for conjugate even domain. A typical usage will be as follows.

```
// m = floor( n/2 )
REAL :: X(0:2*m+1)
...some other code...
...assuming inplace...
Status = DftiComputeForward( Desc_Handle, X )
```

On input,

$$x(j) = w_j, j = 0, 1, \dots, n-1.$$

On output,

Output data stored in one of formats: Pack, Perm or CCS (see [“Packed formats”](#)).

CCS format: $x(2*k) = \text{Re}(z_k)$, $x(2*k+1) = \text{Im}(z_k)$, $k = 0, 1, \dots, m$.

Pack format: even n : $x(0) = \text{Re}(z_0)$, $x(2*k-1) = \text{Re}(z_k)$, $x(2*k) = \text{Im}(z_k)$, $k = 1, \dots, m-1$, and $x(n-1) = \text{Re}(z_m)$

odd n : $x(0) = \text{Re}(z_0)$, $x(2*k-1) = \text{Re}(z_k)$, $x(2*k) = \text{Im}(z_k)$, $k = 1, \dots, m$

Perm format: even n : $x(0) = \text{Re}(z_0)$, $x(1) = \text{Re}(z_m)$, $x(2*k) = \text{Re}(z_k)$, $x(2*k+1) = \text{Im}(z_k)$, $k = 1, \dots, m-1$,

odd n : $x(0) = \text{Re}(z_0)$, $x(2*k-1) = \text{Re}(z_k)$, $x(2*k) = \text{Im}(z_k)$, $k = 1, \dots, m$.

3. DFTI_REAL_REAL for real domain, DFTI_REAL_REAL for conjugate even domain. This storage scheme for conjugate even domain is applicable for one-dimensional transform only. A typical usage will be as follows.

```
// m = floor( n/2 )
REAL :: X(0:n-1)
...some other code...
...assuming inplace...
Status = DftiComputeForward( Desc_Handle, X )
```

On input,

$$x(j) = w_j, j = 0, 1, \dots, n-1.$$

On output,

$$x(k) = \text{Re}(z_k), k = 0, 1, \dots, m.$$

and

$$x(n-k) = \text{Im}(z_k), k = 1, 2, \dots, m-1.$$

4. DFTI_REAL_COMPLEX for real domain, DFTI_COMPLEX_COMPLEX for conjugate even domain. A typical usage will be as follows.

```
// m = floor( n/2 )
COMPLEX :: X(0:n-1)
...some other code...
...inplace transform...
```

```
Status = DftiComputeForward( Desc_Handle, X )
```

On input,

$$X(j) = w_j, j = 0, 1, \dots, n-1.$$

That is, the imaginary parts of $X(j)$ are zero. On output,

$$Y(k) = z_k, k = 0, 1, \dots, m.$$

where m is $\lfloor n/2 \rfloor$.

5. DFTI_REAL_COMPLEX for real domain, DFTI_COMPLEX_REAL for conjugate even domain. A typical usage will be as follows.

```
// m = floor( n/2 )
COMPLEX :: X(0:n-1)
REAL    :: Y(0:2*m+1)
...some other code...
...not inplace...
Status = DftiComputeForward( Desc_Handle, X, Y )
```

On input,

$$X(j) = w_j, j = 0, 1, \dots, n-1.$$

On output,

Output data stored in one of formats: Pack, Perm or CCS (see [“Packed formats”](#)).

CCS format: $Y(2*k) = \text{Re}(z_k)$, $Y(2*k+1) = \text{Im}(z_k)$, $k = 0, 1, \dots, m$.

Pack format: even n : $Y(0) = \text{Re}(z_0)$, $Y(2*k-1) = \text{Re}(z_k)$, $Y(2*k) = \text{Im}(z_k)$, $k = 1, \dots, m-1$, and $Y(n-1) = \text{Re}(z_m)$

odd n : $Y(0) = \text{Re}(z_0)$, $Y(2*k-1) = \text{Re}(z_k)$, $Y(2*k) = \text{Im}(z_k)$, $k = 1, \dots, m$

Perm format: even n : $Y(0) = \text{Re}(z_0)$, $Y(1) = \text{Re}(z_m)$, $Y(2*k) = \text{Re}(z_k)$, $Y(2*k+1) = \text{Im}(z_k)$, $k = 1, \dots, m-1$,

odd n : $Y(0) = \text{Re}(z_0)$, $Y(2*k-1) = \text{Re}(z_k)$, $Y(2*k) = \text{Im}(z_k)$, $k = 1, \dots, m$.

6. DFTI_REAL_COMPLEX for real domain, DFTI_REAL_REAL for conjugate even domain. This storage scheme for conjugate even domain is applicable for one-dimensional transform only. A typical usage will be as follows.

```
// m = floor( n/2 )
COMPLEX :: X(0:n-1)
```

```

REAL :: Y(0:n-1)
...some other code...
...not inplace...
Status = DftiComputeForward( Desc_Handle, X, Y )

```

On input,

$$X(j) = w_j, j = 0, 1, \dots, n-1.$$

On output,

$$Y(k) = \text{Re}(z_k), k = 0, 1, \dots, m.$$

and

$$Y(n-k) = \text{Im}(z_k), k = 1, 2, \dots, m-1.$$

Input and output distances

DFT interface in Intel MKL allows the computation of multiple number of transforms. Consequently, the user needs to be able to specify the data distribution of these multiple sets of data. This is accomplished by the distance between the first data element of the consecutive data sets. This parameter is obligatory if multiple number is more than one. Data sets don't have any common elements. The following example illustrates the specification. Consider computing the forward DFT on three 32-length complex sequences stored in `X(0:31, 1)`, `X(0:31, 2)`, and `X(0:31, 3)`. Suppose the results are to be stored in the locations `Y(0:31, k)`, $k = 1, 2, 3$, of the array `Y(0:63, 3)`. Thus the input distance is 32, while the output distance is 64. Notice that the data and result parameters in computation functions are all declared as assumed-size rank-1 array `DIMENSION(0:*)`. Therefore two-dimensional array must be transformed to one-dimensional array by `EQUIVALENCE` statement or other facilities of Fortran. Here is the code fragment:

```

Complex :: X_2D(0:31,3), Y_2D(0:63, 3)
Complex :: X(96), Y(192)
Equivalence (X_2D, X)
Equivalence (Y_2D, Y)
.....
Status = DftiCreateDescriptor(Desc_Handle, DFTI_SINGLE,
                             DFTI_COMPLEX, 1, 32)
Status = DftiSetValue(Desc_Handle, DFTI_NUMBER_OF_TRANSFORM, 3)

```

```
Status = DftiSetValue(Desc_Handle, DFTI_INPUT_DISTANCE, 32)
Status = DftiSetValue(Desc_Handle, DFTI_OUTPUT_DISTANCE, 64)
Status = DftiSetValue(Desc_Handle, DFTI_PLACEMENT, DFTI_NOT_INPLACE)
Status = DftiCommitDescriptor(Desc_Handle)
Status = DftiComputeForward(Desc_Handle, X, Y)
Status = DftiFreeDescriptor(Desc_Handle)
```

Strides

In addition to supporting transforms of multiple number of datasets, DFT interface supports non-unit stride distribution of data within each data set. Consider the following situation where a 32-length DFT is to be computed on the sequence x_j , $0 \leq j < 32$. The actual location of these values are in $\mathbf{x}(5), \mathbf{x}(7), \dots, \mathbf{x}(67)$ of an array $\mathbf{x}(1:68)$. The stride accommodated by DFT interface consists of a displacement from the first element of the data array L_0 , (4 in this case), and a constant distance of consecutive elements L_1 (2 in this case). Thus for the Fortran array \mathbf{x}

$$x_j = \mathbf{x}(1 + L_0 + L_1 * j) = \mathbf{x}(5 + L_1 * j) .$$

This stride vector (4,2) is provided by a length-2 rank-1 integer array:

```
COMPLEX :: X(68)
INTEGER :: Stride(2)
.....
Status = DftiCreateDescriptor(Desc_Handle, DFTI_SINGLE,
                             DFTI_COMPLEX, 1, 32)

Stride = (/ 4, 2 /)
Status = DftiSetValue(Desc_Handle, DFTI_INPUT_STRIDE, Stride)
Status = DftiSetValue(Desc_Handle, DFTI_OUTPUT_STRIDE, Stride)
Status = DftiCommitDescriptor(Desc_Handle)
Status = DftiComputeForward(Desc_Handle, X)
Status = DftiFreeDescriptor(Desc_Handle)
```

In general, for a d -dimensional transform, the stride is provided by a $d+1$ -length integer vector $(L_0, L_1, L_2, \dots, L_d)$ with the meaning:

L_0 = displacement from the first array element

L_1 = distance between consecutive data elements in the first dimension

L_2 = distance between consecutive data elements in the second dimension

... = ...

L_d = distance between consecutive data elements in the d -th dimension.

A d -dimensional data sequence

$$x_{j_1, j_2, \dots, j_d}, \quad 0 \leq j_i < J_i, \quad 1 \leq i \leq d$$

will be stored in the rank-1 array x by the mapping

$$x_{j_1, j_2, \dots, j_d} = x(\text{first index} + L_0 + j_1 L_1 + j_2 L_2 + \dots + j_d L_d).$$

For multiple transforms, the value L_0 applies to the first data sequence, and $L_j, j = 1, 2, \dots, d$ apply to all the data sequences.

In the case of a single one-dimensional sequence, L_j is simply the usual stride. The default setting of strides in the general multi-dimensional situation corresponds to the case where the sequences are distributed tightly into the array:

$$L_1 = 1, L_2 = J_1, L_3 = J_1 J_2, \dots, L_d = \prod_{i=1}^{d-1} J_i$$

Both the input data and output data have a stride associated with it. The default is set in accordance with the data to be stored contiguously in memory in a way that is natural to the language.

Finally, consider a contrived example where a 20-by-40 two-dimensional DFT is computed explicitly using one-dimensional transforms. Notice that the data and result parameters in computation functions are all declared as assumed-size rank-1 array `DIMENSION(0:*)`. Therefore two-dimensional array must be transformed to one-dimensional array by `EQUIVALENCE` statement or other facilities of Fortran.

```
! Fortran
Complex :: X_2D(20,40),
Complex :: X(800)
Equivalence (X_2D, X)
INTEGER :: STRIDE(2)
type(DFTI_DESCRIPTOR), POINTER :: Desc_Handle_Dim1
type(DFTI_DESCRIPTOR), POINTER :: Desc_Handle_Dim2
...
Status = DftiCreateDescriptor( Desc_Handle_Dim1, DFTI_SINGLE,
                             DFTI_COMPLEX, 1, 20 )
```

```

Status = DftiCreateDescriptor( Desc_Handle_Dim2, DFTI_SINGLE,
                               DFTI_COMPLEX, 1, 40 )

! perform 40 one-dimensional transforms along 1st dimension
Status = DftiSetValue( Desc_Handle_Dim1, DFTI_NUMBER_OF_TRANSFORMS, 40 )
Status = DftiSetValue( Desc_Handle_Dim1, DFTI_INPUT_DISTANCE, 20 )
Status = DftiSetValue( Desc_Handle_Dim1, DFTI_OUTPUT_DISTANCE, 20 )
Status = DftiCommitDescriptor( Desc_Handle_Dim1 )
Status = DftiComputeForward( Desc_Handle_Dim1, X )

! perform 20 one-dimensional transforms along 2nd dimension
Stride(1) = 0; Stride(2) = 20
Status = DftiSetValue( Desc_Handle_Dim2, DFTI_NUMBER_OF_TRANSFORMS, 20 )
Status = DftiSetValue( Desc_Handle_Dim2, DFTI_INPUT_DISTANCE, 1 )
Status = DftiSetValue( Desc_Handle_Dim2, DFTI_OUTPUT_DISTANCE, 1 )
Status = DftiSetValue( Desc_Handle_Dim2, DFTI_INPUT_STRIDES, Stride )
Status = DftiSetValue( Desc_Handle_Dim2, DFTI_OUTPUT_STRIDES, Stride )
Status = DftiCommitDescriptor( Desc_Handle_Dim2 )
Status = DftiComputeForward( Desc_Handle_Dim2, X )
Status = DftiFreeDescriptor( Desc_Handle_Dim1 )
Status = DftiFreeDescriptor( Desc_Handle_Dim2 )

/* C */
float _Complex x[20][40];
long stride[2];
DFTI_DESCRIPTOR_HANDLE Desc_Handle_Dim1;
DFTI_DESCRIPTOR_HANDLE Desc_Handle_Dim2;
...
status = DftiCreateDescriptor( &desc_handle_dim1, DFTI_SINGLE,
                               DFTI_COMPLEX, 1, 20 );
status = DftiCreateDescriptor( &desc_handle_dim2, DFTI_SINGLE,
                               DFTI_COMPLEX, 1, 40 );

/* perform 40 one-dimensional transforms along 1st dimension */
/* note that the 1st dimension data are not unit-stride */
stride[0] = 0; stride[1] = 40;
status = DftiSetValue( desc_handle_dim1, DFTI_NUMBER_OF_TRANSFORMS, 40 );

```



```

status = DftiSetValue( desc_handle_dim1, DFTI_INPUT_DISTANCE, 1 );
status = DftiSetValue( desc_handle_dim1, DFTI_OUTPUT_DISTANCE, 1 );
status = DftiSetValue( desc_handle_dim1, DFTI_INPUT_STRIDES, stride );
status = DftiSetValue( desc_handle_dim1, DFTI_OUTPUT_STRIDES, stride );
status = DftiCommitDescriptor( desc_handle_dim1 );
status = DftiComputeForward( desc_handle_dim1, x );

/* perform 20 one-dimensional transforms along 2nd dimension */
/* note that the 2nd dimension is unit stride */
status = DftiSetValue( desc_handle_dim2, DFTI_NUMBER_OF_TRANSFORMS, 20 );
status = DftiSetValue( desc_handle_dim2, DFTI_INPUT_DISTANCE, 40 );
status = DftiSetValue( desc_handle_dim2, DFTI_OUTPUT_DISTANCE, 40 );
status = DftiCommitDescriptor( desc_handle_dim2 );
status = DftiComputeForward( desc_handle_dim2, x );
status = DftiFreeDescriptor( &Desc_Handle_Dim1 );
status = DftiFreeDescriptor( &Desc_Handle_Dim2 );

```

Initialization Effort

In modern approaches to constructing fast algorithms (FFT) for DFT computations, one often has a flexibility of spending more effort in initializing (preparing for) an FFT algorithm to buy higher efficiency in the computation on actual data to follow. Advanced DFT functions in Intel MKL accommodate this situation through the configuration parameter `DFTI_INITIALIZATION_EFFORT`. The three configuration values are `DFTI_LOW`, `DFTI_MEDIUM` (default), and `DFTI_HIGH`. Note that specific implementations of DFT interface may or may not make use of this setting (see *MKL Release Notes* for implementation details).

Ordering

It is well known that a number of FFT algorithms apply an explicit permutation stage that is time consuming [4]. The exclusion of this step is similar to applying DFT to input data whose order is scrambled, or allowing a scrambled order of the DFT results. In applications such as convolution and power spectrum calculation, the order of result or data is unimportant and thus permission of scrambled order is attractive if it leads to higher performance. Three following options are available in Intel MKL:

1. **DFTI_ORDERED**: Forward transform data ordered, backward transform data ordered (default option).
2. **DFTI_BACKWARD_SCRAMBLED**: Forward transform data ordered, backward transform data scrambled.
3. **DFTI_FORWARD_SCRAMBLED**: Forward transform data scrambled, backward transform data ordered.

[Table 9-12](#) tabulates the effect on this configuration setting.

Table 9-12 Scrambled Order Transform

	DftiComputeForward	DftiComputeBackward
DFTI_ORDERING	Input →Output	Input →Output
DFTI_ORDERED	ordered →ordered	ordered →ordered
DFTI_BACKWARD_SCRAMBLED	ordered →scrambled	scrambled →ordered
DFTI_FORWARD_SCRAMBLED	scrambled →ordered	ordered →scrambled

Note that meaning of the latter two options are "allow scrambled order if practical." There are situations where in fact allowing out of order data gives no performance advantage, and thus an implementation may choose to ignore the suggestion. Strictly speaking, the normal order is also a scrambled order, the trivial one.

When the ordering setting is other than the default **DFTI_ORDERED**, the user may need to know the actual ordering of the input and output data. The ordering of the data in the forward domain is obtained through reading (getting) the configuration parameter **DFTI_FORWARD_ORDERING**; and the ordering of the data in the reverse domain is obtained through reading (getting) the configuration parameter **DFTI_BACKWARD_ORDERING**. The configuration values are integer vectors, thus provided by pointer to any integer array. We now describe how these integer values specify the actual scrambling of data.

All scramblings involved are digit reversal along one single dimension. Precisely, a length J is factored into K ordered factors D_1, D_2, \dots, D_K . Any index i , $0 \leq i < n$, can be expressed uniquely as K digits i_1, i_2, \dots, i_K where $0 \leq i_l < D_l$ and

$$i = i_1 + i_2 D_1 + i_3 D_1 D_2 + \dots + i_K D_1 D_2 \dots D_{K-1} .$$

A digit reversal permutation $\text{scram}(i)$ is given by

$$\text{scram}(i) = i_K + i_{K-1}D_K + i_{K-2}D_KD_{K-1} + \dots + i_1D_KD_{K-1} \dots D_2$$

Factoring J into one factor J leads to no scrambling at all, that is, $\text{scram}(i) = i$. Note that the factoring does not need to correspond exactly to the number of "butterfly" stages to be carried out. In fact, the computation routine in its initialization stage determines if a scrambled order in some or all of the dimensions can result in performance gain. The digits of the digit reversal are recorded and stored in the descriptor. These digits can be obtained by calling a corresponding inquiry routine that returns a pointer to an integer array. The first element is $K^{(1)}$, which is the number of digits for the first dimension, followed by $K^{(1)}$ values of the corresponding digits. If the dimension is higher than one, the next integer value is $K^{(2)}$, etc.

Simple permutation such as mod- p sort [4] is a special case of digit reversal. Hence this option could be useful for high-performance implementation of one-dimensional DFT via a "six-step" or "four-step" framework [4].

The user can check the scrambling setting on the forward data and reverse data. This information is returned as an integer vector containing a number of sequence $(K, D_1, D_2, \dots, D_K)$, one for each dimension. Thus the first element indicates how many D 's will follow. The inquiry routine allocates memory, fills it with this information, and returns a pointer to the memory location.

Workspace

Some FFT algorithms do not require a scratch space for permutation purposes. The user can choose between the setting of `DFTI_ALLOW` (default) and `DFTI_AVOID` for the option `DFTI_WORKSPACE`. Note that the setting `DFTI_AVOID` is meant to be "avoid if practical," hence allowing the implementation the flexibility to use workspace regardless of the setting.

Transposition

This is an option that allows for the result of a high-dimensional transform to be presented in a transposed manner. The default setting is `DFTI_NONE` and can be set to `DFTI_ALLOW`. Similar to that of scrambled order, sometimes in higher dimension transform, performance can be gained if the

result is delivered in a transposed manner. DFT interface offers an option for the output be returned in a transposed form if performance gain is possible. Since the generic stride specification is naturally suited for representation of transposition, this option allows the strides for the output to be possibly different from those originally specified by the user. Consider an example where a two-dimensional result \mathbf{Y}_{j_1, j_2} , $0 \leq j_i < n_i$, is expected. Originally the user specified that the result be distributed in the (flat) array \mathbf{y} in with generic strides $L_1 = 1$ and $L_2 = n_1$. With the transposition option, the computation may actually return the result into \mathbf{y} with stride $L_1 = n_2$ and $L_2 = 1$. These strides can be obtained from an appropriate inquiry function. Note also that in dimension 3 and above, transposition means an arbitrary permutation of the dimension.

Routine and Function Arguments



The major arguments in the BLAS routines are vector and matrix, whereas VML functions work on vector arguments only.

The sections that follow discuss each of these arguments and provide examples.

Vector Arguments in BLAS

Vector arguments are passed in one-dimensional arrays. The array dimension (length) and vector increment are passed as integer variables. The length determines the number of elements in the vector. The increment (also called stride) determines the spacing between vector elements and the order of the elements in the array in which the vector is passed.

A vector of length n and increment $incx$ is passed in a one-dimensional array x whose values are defined as

$x(1), x(1+|incx|), \dots, x(1+(n-1)*|incx|)$

If $incx$ is positive, then the elements in array x are stored in increasing order. If $incx$ is negative, the elements in array x are stored in decreasing order with the first element defined as $x(1+(n-1)*|incx|)$. If $incx$ is zero, then all elements of the vector have the same value, $x(1)$. The dimension of the one-dimensional array that stores the vector must always be at least

$idimx = 1 + (n-1)*|incx|$

Example A-1 One-dimensional Real Array

Let $x(1:7)$ be the one-dimensional real array

$x = (1.0, 3.0, 5.0, 7.0, 9.0, 11.0, 13.0)$.

If $incx = 2$ and $n = 3$, then the vector argument with elements in order from first to last is $(1.0, 5.0, 9.0)$.

If $incx = -2$ and $n = 4$, then the vector elements in order from first to last is $(13.0, 9.0, 5.0, 1.0)$.

If $incx = 0$ and $n = 4$, then the vector elements in order from first to last is $(1.0, 1.0, 1.0, 1.0)$.

One-dimensional substructures of a matrix, such as the rows, columns, and diagonals, can be passed as vector arguments with the starting address and increment specified. In Fortran, storing the m by n matrix is based on column-major ordering where the increment between elements in the same column is 1 , the increment between elements in the same row is m , and the increment between elements on the same diagonal is $m + 1$.

Example A-2 Two-dimensional Real Matrix

Let a be the real 5×4 matrix declared as `REAL A (5,4)`.

To scale the third column of a by 2.0, use the BLAS routine `sscal` with the following calling sequence:

```
call sscal (5, 2.0, a(1,3), 1).
```

To scale the second row, use the statement:

```
call sscal (4, 2.0, a(2,1), 5).
```

To scale the main diagonal of A by 2.0, use the statement:

```
call sscal (5, 2.0, a(1,1), 6).
```



NOTE. *The default vector argument is assumed to be 1.*

Vector Arguments in VML

Vector arguments of VML mathematical functions are passed in one-dimensional arrays with unit vector increment. It means that a vector of length n is passed contiguously in an array \mathbf{a} whose values are defined as $a[0], a[1], \dots, a[n-1]$ (for C- interface).

To accommodate for arrays with other increments, or more complicated indexing, VML contains auxiliary pack/unpack functions that gather the array elements into a contiguous vector and then scatter them after the computation is complete.

Generally, if the vector elements are stored in a one-dimensional array \mathbf{a} as $a[m_0], a[m_1], \dots, a[m_{n-1}]$

and need to be regrouped into an array \mathbf{y} as

$y[k_0], y[k_1], \dots, y[k_{n-1}]$,

VML pack/unpack functions can use one of the following indexing methods:

Positive Increment Indexing

$k_j = \text{incy} * j, m_j = \text{inca} * j, \quad j = 0, \dots, n-1$

Constraint: $\text{incy} > 0$ and $\text{inca} > 0$.

For example, setting $\text{incy} = 1$ specifies gathering array elements into a contiguous vector.

This method is similar to that used in BLAS, with the exception that negative and zero increments are not permitted.

Index Vector Indexing

$k_j = \text{iy}[j], m_j = \text{ia}[j], \quad j = 0, \dots, n-1,$

where ia and iy are arrays of length n that contain index vectors for the input and output arrays \mathbf{a} and \mathbf{y} , respectively.

Mask Vector Indexing

Indices k_j, m_j are such that:

$\text{my}[k_j] \neq 0, \text{ma}[m_j] \neq 0, \quad j = 0, \dots, n-1,$

where ma and my are arrays that contain mask vectors for the input and output arrays \mathbf{a} and \mathbf{y} , respectively.

Matrix Arguments

Matrix arguments of the Intel® Math Kernel Library routines can be stored in either one- or two-dimensional arrays, using the following storage schemes:

- conventional full storage (in a two-dimensional array)
- packed storage for Hermitian, symmetric, or triangular matrices (in a one-dimensional array)
- band storage for band matrices (in a two-dimensional array).

Full storage is the following obvious scheme: a matrix A is stored in a two-dimensional array a , with the matrix element a_{ij} stored in the array element $a(i, j)$.

If a matrix is *triangular* (upper or lower, as specified by the argument $uplo$), only the elements of the relevant triangle are stored; the remaining elements of the array need not be set.

Routines that handle symmetric or Hermitian matrices allow for either the upper or lower triangle of the matrix to be stored in the corresponding elements of the array:

- if $uplo = 'U'$, a_{ij} is stored in $a(i, j)$ for $i \leq j$,
other elements of a need not be set.
- if $uplo = 'L'$, a_{ij} is stored in $a(i, j)$ for $j \leq i$,
other elements of a need not be set.

Packed storage allows you to store symmetric, Hermitian, or triangular matrices more compactly: the relevant triangle (again, as specified by the argument $uplo$) is packed by columns in a one-dimensional array ap :

- if $uplo = 'U'$, a_{ij} is stored in $ap(i+j(j-1)/2)$ for $i \leq j$
if $uplo = 'L'$, a_{ij} is stored in $ap(i+(2*n-j)*(j-1)/2)$ for $j \leq i$.

In descriptions of LAPACK routines, arrays with packed matrices have names ending in p .

Band storage is as follows: an m by n band matrix with kl non-zero sub-diagonals and ku non-zero super-diagonals is stored compactly in a two-dimensional array ab with $kl+ku+1$ rows and n columns. Columns of the matrix are stored in the corresponding columns of the array, and diagonals of the matrix are stored in rows of the array. Thus,

a_{ij} is stored in $ab(kl+ku+1+i-j, j)$ for $\max(n, j-ku) \leq i \leq \min(n, j+kl)$.

Use the band storage scheme only when kl and ku are much less than the matrix size n . (Although the routines work correctly for all values of kl and ku , it's inefficient to use the band storage if your matrices are not really banded).

When a general band matrix is supplied for *LU factorization*, space must be allowed to store kl additional super-diagonals generated by fill-in as a result of row interchanges. This means that the matrix is stored according to the above scheme, but with $kl + ku$ super-diagonals.

The band storage scheme is illustrated by the following example, when $m = n = 6$, $kl = 2$, $ku = 1$:

Banded matrix A	Band storage of A
$\begin{bmatrix} a_{11} & a_{12} & 0 & 0 & 0 & 0 \\ a_{21} & a_{22} & a_{23} & 0 & 0 & 0 \\ a_{31} & a_{32} & a_{33} & a_{34} & 0 & 0 \\ 0 & a_{42} & a_{43} & a_{44} & a_{45} & 0 \\ 0 & 0 & a_{53} & a_{54} & a_{55} & a_{56} \\ 0 & 0 & 0 & a_{64} & a_{65} & a_{66} \end{bmatrix}$	$\begin{matrix} * & * & * & + & + & + \\ * & * & + & + & + & + \\ * & a_{12} & a_{23} & a_{34} & a_{45} & a_{56} \\ a_{11} & a_{22} & a_{33} & a_{44} & a_{55} & a_{66} \\ a_{21} & a_{32} & a_{43} & a_{54} & a_{65} & * \\ a_{31} & a_{42} & a_{53} & a_{64} & * & * \end{matrix}$

Array elements marked * are not used by the routines; elements marked + need not be set on entry, but are required by the LU factorization routines to store the results. The input array will be overwritten on exit by the details of the LU factorization as follows:

$$\begin{matrix} * & * & * & u_{14} & u_{25} & u_{36} \\ * & * & u_{13} & u_{24} & u_{35} & u_{46} \\ * & u_{12} & u_{23} & u_{34} & u_{45} & u_{56} \\ u_{11} & u_{22} & u_{33} & u_{44} & u_{55} & u_{66} \\ m_{21} & m_{32} & m_{43} & m_{54} & m_{65} & * \\ m_{31} & m_{42} & m_{53} & m_{64} & * & * \end{matrix}$$

where u_{ij} are the elements of the upper triangular matrix U, and m_{ij} are the multipliers used during factorization.

Triangular band matrices are stored in the same format, with either $kl=0$ if upper triangular, or $ku=0$ if lower triangular. For symmetric or Hermitian band matrices with k sub-diagonals or super-diagonals, you need to store only the upper or lower triangle, as specified by the argument *uplo*:

if *uplo* = 'U', a_{ij} is stored in $ab(k+1+i-j, j)$ for $\max(1, j-k) \leq i \leq j$
 if *uplo* = 'L', a_{ij} is stored in $ab(1+i-j, j)$ for $j \leq i \leq \min(n, j+k)$.

In descriptions of LAPACK routines, arrays that hold matrices in band storage have names ending in *b*.

In Fortran, column-major ordering of storage is assumed. This means that elements of the same column occupy successive storage locations.

Three quantities are usually associated with a two-dimensional array argument: its leading dimension, which specifies the number of storage locations between elements in the same row, its number of rows, and its number of columns. For a matrix in full storage, the leading dimension of the array must be at least as large as the number of rows in the matrix.

A character transposition parameter is often passed to indicate whether the matrix argument is to be used in normal or transposed form or, for a complex matrix, if the conjugate transpose of the matrix is to be used. The values of the transposition parameter for these three cases are the following:

- 'N' or 'n' normal (no conjugation, no transposition)
- 'T' or 't' transpose
- 'C' or 'c' conjugate transpose.

Example A-3 Two-Dimensional Complex Array

Suppose $A(1:5, 1:4)$ is the complex two-dimensional array presented by matrix

$$\begin{bmatrix} (1.1, 0.11) & (1.2, 0.12) & (1.3, 0.13) & (1.4, 0.14) \\ (2.1, 0.21) & (2.2, 0.22) & (2.3, 0.23) & (2.4, 0.24) \\ (3.1, 0.31) & (3.2, 0.32) & (3.3, 0.33) & (3.4, 0.34) \\ (4.1, 0.41) & (4.2, 0.42) & (4.3, 0.43) & (4.4, 0.44) \\ (5.1, 0.51) & (5.2, 0.52) & (5.3, 0.53) & (5.4, 0.54) \end{bmatrix}$$

Let $transa$ be the transposition parameter, m be the number of rows, n be the number of columns, and lda be the leading dimension. Then if $transa = 'N'$, $m = 4$, $n = 2$, and $lda = 5$, the matrix argument would be

$$\begin{bmatrix} (1.1, 0.11) & (1.2, 0.12) \\ (2.1, 0.21) & (2.2, 0.22) \\ (3.1, 0.31) & (3.2, 0.32) \\ (4.1, 0.41) & (4.2, 0.42) \end{bmatrix}$$

If $transa = 'T'$, $m = 4$, $n = 2$, and $lda = 5$, the matrix argument would be

$$\begin{bmatrix} (1.1, 0.11) & (2.1, 0.21) & (3.1, 0.31) & (4.1, 0.41) \\ (1.2, 0.12) & (2.2, 0.22) & (3.2, 0.32) & (4.2, 0.42) \end{bmatrix}$$

If $transa = 'C'$, $m = 4$, $n = 2$, and $lda = 5$, the matrix argument would be

$$\begin{bmatrix} (1.1, -0.11) & (2.1, -0.21) & (3.1, -0.31) & (4.1, -0.41) \\ (1.2, -0.12) & (2.2, -0.22) & (3.2, -0.32) & (4.2, -0.42) \end{bmatrix}$$

Note that care should be taken when using a leading dimension value which is different from the number of rows specified in the declaration of the two-dimensional array. For example, suppose the array A above is declared as $\text{COMPLEX } A(5, 4)$.

continued *

Then if `transa = 'N'`, `m = 3`, `n = 4`, and `lda = 4`, the matrix argument will be

$$\begin{bmatrix} (1.1, 0.11) & (5.1, 0.51) & (4.2, 0.42) & (3.3, 0.33) \\ (2.1, 0.21) & (1.2, 0.12) & (5.2, 0.52) & (4.3, 0.43) \\ (3.1, 0.31) & (2.2, 0.22) & (1.3, 0.13) & (5.3, 0.53) \end{bmatrix}$$

Code Examples

B

This appendix presents code examples of using BLAS routines and functions.

Example B-1 Using BLAS Level 1 Function



[?dot](#)
[description](#)

The following example illustrates a call to the BLAS Level 1 function `sdot`. This function performs a vector-vector operation of computing a scalar product of two single-precision real vectors x and y .

Parameters

n Specifies the order of vectors x and y .
incx Specifies the increment for the elements of x .
incy Specifies the increment for the elements of y .

```
program dot_main
real x(10), y(10), sdot, res
integer n, incx, incy, i
external sdot
n = 5
incx = 2
incy = 1
do i = 1, 10
  x(i) = 2.0e0
  y(i) = 1.0e0
end do
```

continued *

Example B-1 Using BLAS Level 1 Function (continued)

```
res = sdot (n, x, incx, y, incy)
print*, 'SDOT = ', res
end
```

As a result of this program execution, the following line is printed:

```
SDOT = 10.000
```

Example B-2 Using BLAS Level 1 Routine



[?copy](#)
[description](#)

The following example illustrates a call to the BLAS Level 1 routine `scopy`. This routine performs a vector-vector operation of copying a single-precision real vector `x` to a vector `y`.

Parameters

<code>n</code>	Specifies the order of vectors <code>x</code> and <code>y</code> .
<code>incx</code>	Specifies the increment for the elements of <code>x</code> .
<code>incy</code>	Specifies the increment for the elements of <code>y</code> .

```
program copy_main
real x(10), y(10)
integer n, incx, incy, i
n = 3
incx = 3
incy = 1
do i = 1, 10
  x(i) = i
end do
call scopy (n, x, incx, y, incy)
print*, 'Y = ', (y(i), i = 1, n)
end
```

As a result of this program execution, the following line is printed:

```
Y = 1.00000 4.00000 7.00000
```

Example B-3 Using BLAS Level 2 Routine



[?ger](#)
[description](#)

The following example illustrates a call to the BLAS Level 2 routine `sger`. This routine performs a matrix-vector operation

$$a := \alpha x y' + a.$$
Parameters

alpha Specifies a scalar *alpha*.
x *m*-element vector.
y *n*-element vector.
a *m* by *n* matrix.

```
program ger_main
real a(5,3), x(10), y(10), alpha
integer m, n, incx, incy, i, j, lda
m = 2
n = 3
lda = 5
incx = 2
incy = 1
alpha = 0.5
do i = 1, 10
  x(i) = 1.0
  y(i) = 1.0
end do
do i = 1, m
  do j = 1, n
    a(i,j) = j
  end do
end do
call sger (m, n, alpha, x, incx, y, incy, a, lda)
print*, 'Matrix A: '
do i = 1, m
  print*, (a(i,j), j = 1, n)
end do
end
```

continued *

Example B-3 Using BLAS Level 2 Routine (continued)

As a result of this program execution, matrix *a* is printed as follows:

Matrix A:

```
1.50000 2.50000 3.50000
1.50000 2.50000 3.50000
```

Example B-4 Using BLAS Level 3 Routine



[?symm](#)
[description](#)

The following example illustrates a call to the BLAS Level 3 routine *ssymm*. This routine performs a matrix-matrix operation

```
c := alpha*a*b' + beta*c.
```

Parameters

<i>alpha</i>	Specifies a scalar <i>alpha</i> .
<i>beta</i>	Specifies a scalar <i>beta</i> .
<i>a</i>	Symmetric matrix.
<i>b</i>	<i>m</i> by <i>n</i> matrix.
<i>c</i>	<i>m</i> by <i>n</i> matrix.

```
program symm_main
real a(3,3), b(3,2), c(3,3), alpha, beta
integer m, n, lda, ldb, ldc, i, j
character uplo, side
uplo = 'u'
side = 'l'
m = 3
n = 2
lda = 3
ldb = 3
ldc = 3
alpha = 0.5
beta = 2.0
```

continued *

Example B-4 Using BLAS Level 3 Routine (continued)

```
do i = 1, m
  do j = 1, m
    a(i,j) = 1.0
  end do
end do
do i = 1, m
  do j = 1, n
    c(i,j) = 1.0
    b(i,j) = 2.0
  end do
end do
call ssymm (side, uplo, m, n, alpha, a, lda, b, ldb,
beta, c, ldc)
print*, 'Matrix C: '
do i = 1, m
  print*, (c(i,j), j = 1, n)
end do
end
```

As a result of this program execution, matrix *c* is printed as follows:

Matrix C:

```
5.00000 5.00000
5.00000 5.00000
5.00000 5.00000
```

Example B-5 Calling a Complex BLAS Level 1 Function from C

The following example illustrates a call from a C program to the complex BLAS Level 1 function `zdotc()`. This function computes the dot product of two double-precision complex vectors.

In this example, the complex dot product is returned in the structure `c`.

```
#define N 5
void main()
{
    int n, inca = 1, incb = 1, i;
    typedef struct{ double re; double im; } complex16;
    complex16 a[N], b[N], c;
    void zdotc();
    n = N;
    for( i = 0; i < n; i++ ){
        a[i].re = (double)i; a[i].im = (double)i * 2.0;
        b[i].re = (double)(n - i); b[i].im = (double)i * 2.0;
    }
    zdotc( &c, &n, a, &inca, b, &incb );
    printf( "The complex dot product is: ( %6.2f, %6.2f
)\n", c.re, c.im );
}
```



NOTE. Instead of calling BLAS directly from C programs, you might wish to use the CBLAS interface; this is the supported way of calling BLAS from C. For more information about CBLAS, see Appendix C, [“CBLAS Interface to the BLAS”](#).

CBLAS Interface to the BLAS



This appendix presents CBLAS, the C interface to the Basic Linear Algebra Subprograms (BLAS) implemented in Intel[®] MKL.

Similar to BLAS, the CBLAS interface includes the following levels of functions:

- [Level 1 CBLAS](#) (vector-vector operations)
- [Level 2 CBLAS](#) (matrix-vector operations)
- [Level 3 CBLAS](#) (matrix-matrix operations).
- [Sparse CBLAS](#) (operations on sparse vectors).

To obtain the C interface, the Fortran routine names are prefixed with `cblas_` (for example, `dasum` becomes `cblas_dasum`). Names of all CBLAS functions are in lowercase letters.

Complex functions `?dotc` and `?dotu` become CBLAS subroutines (void functions); they return the complex result via a void pointer, added as the last parameter. CBLAS names of these functions are suffixed with `_sub`. For example, the BLAS function `cdotc` corresponds to `cblas_cdotc_sub`.

CBLAS Arguments

The arguments of CBLAS functions obey the following rules:

- Input arguments are declared with the `const` modifier.
- Non-complex scalar input arguments are passed by value.
- Complex scalar input arguments are passed as void pointers.
- Array arguments are passed by address.
- Output scalar arguments are passed by address.

- BLAS character arguments are replaced by the appropriate enumerated type.
- Level 2 and Level 3 routines acquire an additional parameter of type `CBLAS_ORDER` as their first argument. This parameter specifies whether two-dimensional arrays are row-major (`CblasRowMajor`) or column-major (`CblasColMajor`).

Enumerated Types

The CBLAS interface uses the following enumerated types:

```
enum CBLAS_ORDER {
    CblasRowMajor=101, /* row-major arrays */
    CblasColMajor=102}; /* column-major arrays */
enum CBLAS_TRANSPOSE {
    CblasNoTrans=111, /* trans='N' */
    CblasTrans=112, /* trans='T' */
    CblasConjTrans=113}; /* trans='C' */
enum CBLAS_UPLO {
    CblasUpper=121, /* uplo = 'U' */
    CblasLower=122}; /* uplo = 'L' */
enum CBLAS_DIAG {
    CblasNonUnit=131, /* diag = 'N' */
    CblasUnit=132}; /* diag = 'U' */
enum CBLAS_SIDE {
    CblasLeft=141, /* side = 'L' */
    CblasRight=142}; /* side = 'R' */
```

Level 1 CBLAS

This is an interface to [BLAS Level 1 Routines and Functions](#), which perform basic vector-vector operations.

?asum

```
float cblas_sasum(const int N, const float *X, const int incX);
double cblas_dasum(const int N, const double *X, const int
incX);
float cblas_scasum(const int N, const void *X, const int incX);
double cblas_dzasum(const int N, const void *X, const int
incX);
```

?axpy

```
void cblas_saxpy(const int N, const float alpha, const float
*X, const int incX, float *Y, const int incY);
void cblas_daxpy(const int N, const double alpha, const double
*X, const int incX, double *Y, const int incY);
void cblas_caxpy(const int N, const void *alpha, const void *X,
const int incX, void *Y, const int incY);
void cblas_zaxpy(const int N, const void *alpha, const void *X,
const int incX, void *Y, const int incY);
```

?copy

```
void cblas_scopy(const int N, const float *X, const int incX,
float *Y, const int incY);
void cblas_dcopy(const int N, const double *X, const int incX,
double *Y, const int incY);
void cblas_ccopy(const int N, const void *X, const int incX,
void *Y, const int incY);
void cblas_zcopy(const int N, const void *X, const int incX,
void *Y, const int incY);
```

?dot

```
float cblas_sdot(const int N, const float *X, const int incX,
const float *Y, const int incY);
double cblas_ddot(const int N, const double *X, const int incX,
const double *Y, const int incY);
```

?sdot

```
float cblas_sdsdot(const int N, const float *SB, const float
*SX, const int incX, const float *SY, const int incY);
```

```
double cblas_dsdot(const int N, const float *SX, const int
incX, const float *SY, const int incY);
```

?dotc

```
void cblas_cdote_sub(const int N, const void *X, const int
incX, const void *Y, const int incY, void *dotc);
```

```
void cblas_zdotc_sub(const int N, const void *X, const int
incX, const void *Y, const int incY, void *dotc);
```

?dotu

```
void cblas_cdotu_sub(const int N, const void *X, const int
incX, const void *Y, const int incY, void *dotu);
```

```
void cblas_zdotu_sub(const int N, const void *X, const int
incX, const void *Y, const int incY, void *dotu);
```

?norm2

```
float cblas_snorm2(const int N, const float *X, const int incX);
```

```
double cblas_dnorm2(const int N, const double *X, const int
incX);
```

```
float cblas_scnrm2(const int N, const void *X, const int incX);
```

```
double cblas_dznrm2(const int N, const void *X, const int
incX);
```

?rot

```
void cblas_srot(const int N, float *X, const int incX, float
*Y, const int incY, const float c, const float s);
```

```
void cblas_drot(const int N, double *X, const int incX, double
*Y, const int incY, const double c, const double s);
```

?rotg

```
void cblas_srotg(float *a, float *b, float *c, float *s);
```

```
void cblas_drotg(double *a, double *b, double *c, double *s);
```

?rotm

```
void cblas_srotm(const int N, float *X, const int incX, float
*Y, const int incY, const float *P);
```

```
void cblas_drotm(const int N, double *X, const int incX, double
*Y, const int incY, const double *P);
```

?rotmg

```
void cblas_srotmg(float *d1, float *d2, float *b1, const float
b2, float *P);
```

```
void cblas_drotmg(double *d1, double *d2, double *b1, const
double b2, double *P);
```

?scal

```
void cblas_sscal(const int N, const float alpha, float *X,
const int incX);
void cblas_dscal(const int N, const double alpha, double *X,
const int incX);
void cblas_cscal(const int N, const void *alpha, void *X, const
int incX);
void cblas_zscal(const int N, const void *alpha, void *X, const
int incX);
void cblas_csscal(const int N, const float alpha, void *X,
const int incX);
void cblas_zdscal(const int N, const double alpha, void *X,
const int incX);
```

?swap

```
void cblas_sswap(const int N, float *X, const int incX, float
*Y, const int incY);
void cblas_dswap(const int N, double *X, const int incX, double
*Y, const int incY);
void cblas_cswap(const int N, void *X, const int incX, void *Y,
const int incY);
void cblas_zswap(const int N, void *X, const int incX, void *Y,
const int incY);
```

i?amax

```
CBLAS_INDEX cblas_isamax(const int N, const float *X, const int
incX);
CBLAS_INDEX cblas_idamax(const int N, const double *X, const
int incX);
CBLAS_INDEX cblas_icamax(const int N, const void *X, const int
incX);
CBLAS_INDEX cblas_izamax(const int N, const void *X, const int
incX);
```

i?amin

```
CBLAS_INDEX cblas_isamin(const int N, const float *X, const int
incX);
CBLAS_INDEX cblas_idamin(const int N, const double *X, const
int incX);
CBLAS_INDEX cblas_icamin(const int N, const void *X, const int
incX);
CBLAS_INDEX cblas_izamin(const int N, const void *X, const int
incX);
```

Level 2 CBLAS

This is an interface to [BLAS Level 2 Routines](#), which perform basic matrix-vector operations. Each C routine in this group has an additional parameter of type `CBLAS_ORDER` (the first argument) that determines whether the two-dimensional arrays use column-major or row-major storage.

?gbmv

```
void cblas_sgbmv(const enum CBLAS_ORDER order, const enum
CBLAS_TRANSPOSE TransA, const int M, const int N, const int KL,
const int KU, const float alpha, const float *A, const int lda,
const float *X, const int incX, const float beta, float *Y,
const int incY);
```

```
void cblas_dgbmv(const enum CBLAS_ORDER order, const enum
CBLAS_TRANSPOSE TransA, const int M, const int N, const int KL,
const int KU, const double alpha, const double *A, const int
lda, const double *X, const int incX, const double beta, double
*Y, const int incY);
```

```
void cblas_cgbmv(const enum CBLAS_ORDER order, const enum
CBLAS_TRANSPOSE TransA, const int M, const int N, const int KL,
const int KU, const void *alpha, const void *A, const int lda,
const void *X, const int incX, const void *beta, void *Y, const
int incY);
```

```
void cblas_zgbmv(const enum CBLAS_ORDER order, const enum
CBLAS_TRANSPOSE TransA, const int M, const int N, const int KL,
const int KU, const void *alpha, const void *A, const int lda,
const void *X, const int incX, const void *beta, void *Y, const
int incY);
```

?gemv

```
void cblas_sgemv(const enum CBLAS_ORDER order, const enum
CBLAS_TRANSPOSE TransA, const int M, const int N, const float
alpha, const float *A, const int lda, const float *X, const int
incX, const float beta, float *Y, const int incY);
```

```
void cblas_dgemv(const enum CBLAS_ORDER order, const enum
CBLAS_TRANSPOSE TransA, const int M, const int N, const double
alpha, const double *A, const int lda, const double *X, const
int incX, const double beta, double *Y, const int incY);
```

```
void cblas_cgemv(const enum CBLAS_ORDER order, const enum
CBLAS_TRANSPOSE TransA, const int M, const int N, const void
*alpha, const void *A, const int lda, const void *X, const int
incX, const void *beta, void *Y, const int incY);
```



```
void cblas_zgemv(const enum CBLAS_ORDER order, const enum
CBLAS_TRANSPOSE TransA, const int M, const int N, const void
*alpha, const void *A, const int lda, const void *X, const int
incX, const void *beta, void *Y, const int incY);
```

?ger

```
void cblas_sger(const enum CBLAS_ORDER order, const int M,
const int N, const float alpha, const float *X, const int incX,
const float *Y, const int incY, float *A, const int lda);
```

```
void cblas_dger(const enum CBLAS_ORDER order, const int M,
const int N, const double alpha, const double *X, const int
incX, const double *Y, const int incY, double *A, const int
lda);
```

?gerc

```
void cblas_cgerc(const enum CBLAS_ORDER order, const int M,
const int N, const void *alpha, const void *X, const int incX,
const void *Y, const int incY, void *A, const int lda);
```

```
void cblas_zgerc(const enum CBLAS_ORDER order, const int M,
const int N, const void *alpha, const void *X, const int incX,
const void *Y, const int incY, void *A, const int lda);
```

?geru

```
void cblas_cgeru(const enum CBLAS_ORDER order, const int M,
const int N, const void *alpha, const void *X, const int incX,
const void *Y, const int incY, void *A, const int lda);
```

```
void cblas_zgeru(const enum CBLAS_ORDER order, const int M,
const int N, const void *alpha, const void *X, const int incX,
const void *Y, const int incY, void *A, const int lda);
```

?hbm

```
void cblas_chbmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const int K, const void *alpha,
const void *A, const int lda, const void *X, const int incX,
const void *beta, void *Y, const int incY);
```

```
void cblas_zhbmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const int K, const void *alpha,
const void *A, const int lda, const void *X, const int incX,
const void *beta, void *Y, const int incY);
```

?hem

```
void cblas_chemv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const void *alpha, const void *A,
const int lda, const void *X, const int incX, const void *beta,
void *Y, const int incY);
```

```
void cblas_zhemv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const void *alpha, const void *A,
const int lda, const void *X, const int incX, const void *beta,
void *Y, const int incY);
```

?her

```
void cblas_cher(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const float alpha, const void *X,
const int incX, void *A, const int lda);
```

```
void cblas_zher(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const double alpha, const void
*X, const int incX, void *A, const int lda);
```

?her2

```
void cblas_cher2(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const void *alpha, const void *X,
const int incX, const void *Y, const int incY, void *A, const
int lda);
```

```
void cblas_zher2(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const void *alpha, const void *X,
const int incX, const void *Y, const int incY, void *A, const
int lda);
```

?hpmv

```
void cblas_chpmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const void *alpha, const void
*Ap, const void *X, const int incX, const void *beta, void *Y,
const int incY);
```

```
void cblas_zhpmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const void *alpha, const void
*Ap, const void *X, const int incX, const void *beta, void *Y,
const int incY);
```

?hpr

```
void cblas_chpr(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const float alpha, const void *X,
const int incX, void *A);
```

```
void cblas_zhpr(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const double alpha, const void
*X, const int incX, void *A);
```

?hpr2

```
void cblas_chpr2(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const void *alpha, const void *X,
const int incX, const void *Y, const int incY, void *Ap);
```

```
void cblas_zhpr2(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const void *alpha, const void *X,
const int incX, const void *Y, const int incY, void *Ap);
```

?sbmv

```
void cblas_ssbmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const int K, const float alpha,
const float *A, const int lda, const float *X, const int incX,
const float beta, float *Y, const int incY);
```

```
void cblas_dsbmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const int K, const double alpha,
const double *A, const int lda, const double *X, const int
incX, const double beta, double *Y, const int incY);
```

?spmv

```
void cblas_sspmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const float alpha, const float
*Ap, const float *X, const int incX, const float beta, float
*Y, const int incY);
```

```
void cblas_dspmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const double alpha, const double
*Ap, const double *X, const int incX, const double beta, double
*Y, const int incY);
```

?spr

```
void cblas_sspr(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const float alpha, const float
*X, const int incX, float *Ap);
```

```
void cblas_dspr(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const double alpha, const double
*X, const int incX, double *Ap);
```

?spr2

```
void cblas_sspr2(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const float alpha, const float
*X, const int incX, const float *Y, const int incY, float *A);
```

```
void cblas_dspr2(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const double alpha, const double
*X, const int incX, const double *Y, const int incY, double
*A);
```

?symv

```
void cblas_ssymv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const float alpha, const float
*A, const int lda, const float *X, const int incX, const float
beta, float *Y, const int incY);
```

```
void cblas_dsymv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const double alpha, const double
*A, const int lda, const double *X, const int incX, const
double beta, double *Y, const int incY);
```

?syr

```
void cblas_ssyr(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const float alpha, const float
*X, const int incX, float *A, const int lda);
```

```
void cblas_dsyr(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const double alpha, const double
*X, const int incX, double *A, const int lda);
```

?syr2

```
void cblas_ssyr2(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const float alpha, const float
*X, const int incX, const float *Y, const int incY, float *A,
const int lda);
```

```
void cblas_dsyr2(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const int N, const double alpha, const double
*X, const int incX, const double *Y, const int incY, double *A,
const int lda);
```

?tbmv

```
void cblas_stbmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const int K, const float *A,
const int lda, float *X, const int incX);
```

```
void cblas_dtbmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const int K, const double *A,
const int lda, double *X, const int incX);
```

```
void cblas_ctbmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const int K, const void *A, const
int lda, void *X, const int incX);
```

```
void cblas_ztbmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const int K, const void *A, const
int lda, void *X, const int incX);
```

?tbsv

```
void cblas_stbsv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const int K, const float *A,
const int lda, float *X, const int incX);
```

```
void cblas_dtbsv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const int K, const double *A,
const int lda, double *X, const int incX);
```

```
void cblas_ctbsv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const int K, const void *A, const
int lda, void *X, const int incX);
```

```
void cblas_ztbsv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const int K, const void *A, const
int lda, void *X, const int incX);
```

?tpmv

```
void cblas_stpmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const float *Ap, float *X, const
int incX);
```

```
void cblas_dtpmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const double *Ap, double *X,
const int incX);
```

```
void cblas_ctpmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const void *Ap, void *X, const int
incX);
```

```
void cblas_ztpmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const void *Ap, void *X, const int
incX);
```

?tpsv

```
void cblas_stpsv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const float *Ap, float *X, const
int incX);
```

```
void cblas_dtpsv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const double *Ap, double *X, const
int incX);
```

```
void cblas_ctpsv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const void *Ap, void *X, const int
incX);
```

```
void cblas_ztpsv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const void *Ap, void *X, const int
incX);
```

?trmv

```
void cblas_strmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const float *A, const int lda,
float *X, const int incX);
```

```
void cblas_dtrmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const double *A, const int lda,
double *X, const int incX);
```

```
void cblas_ctrmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const void *A, const int lda, void
*X, const int incX);
```

```
void cblas_ztrmv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const void *A, const int lda, void
*X, const int incX);
```

?trsv

```
void cblas_strsv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const float *A, const int lda,
float *X, const int incX);
```

```
void cblas_dtrsv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const double *A, const int lda,
double *X, const int incX);

void cblas_ctrsv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const void *A, const int lda, void
*X, const int incX);

void cblas_ztrsv(const enum CBLAS_ORDER order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE TransA, const enum
CBLAS_DIAG Diag, const int N, const void *A, const int lda, void
*X, const int incX);
```

Level 3 CBLAS

This is an interface to [BLAS Level 3 Routines](#), which perform basic matrix-matrix operations. Each C routine in this group has an additional parameter of type `CBLAS_ORDER` (the first argument) that determines whether the two-dimensional arrays use column-major or row-major storage.

?gemm

```
void cblas_sgemm(const enum CBLAS_ORDER Order, const enum
CBLAS_TRANSPOSE TransA, const enum CBLAS_TRANSPOSE TransB,
const int M, const int N, const int K, const float alpha, const
float *A, const int lda, const float *B, const int ldb, const
float beta, float *C, const int ldc);
```

```
void cblas_dgemm(const enum CBLAS_ORDER Order, const enum
CBLAS_TRANSPOSE TransA, const enum CBLAS_TRANSPOSE TransB,
const int M, const int N, const int K, const double alpha,
const double *A, const int lda, const double *B, const int ldb,
const double beta, double *C, const int ldc);
```

```
void cblas_cgemm(const enum CBLAS_ORDER Order, const enum
CBLAS_TRANSPOSE TransA, const enum CBLAS_TRANSPOSE TransB,
const int M, const int N, const int K, const void *alpha, const
void *A, const int lda, const void *B, const int ldb, const
void *beta, void *C, const int ldc);
```

```
void cblas_zgemm(const enum CBLAS_ORDER Order, const enum
CBLAS_TRANSPOSE TransA, const enum CBLAS_TRANSPOSE TransB,
const int M, const int N, const int K, const void *alpha, const
void *A, const int lda, const void *B, const int ldb, const
void *beta, void *C, const int ldc);
```

?hemm

```
void cblas_chemm(const enum CBLAS_ORDER Order, const enum
CBLAS_SIDE Side, const enum CBLAS_UPLO Uplo, const int M, const
int N, const void *alpha, const void *A, const int lda, const
void *B, const int ldb, const void *beta, void *C, const int
ldc);
```

```
void cblas_zhemm(const enum CBLAS_ORDER Order, const enum
CBLAS_SIDE Side, const enum CBLAS_UPLO Uplo, const int M, const
int N, const void *alpha, const void *A, const int lda, const
void *B, const int ldb, const void *beta, void *C, const int
ldc);
```


?herk

```
void cblas_cherk(const enum CBLAS_ORDER Order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE Trans, const int N,
const int K, const float alpha, const void *A, const int lda,
const float beta, void *C, const int ldc);
```

```
void cblas_zherk(const enum CBLAS_ORDER Order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE Trans, const int N,
const int K, const double alpha, const void *A, const int lda,
const double beta, void *C, const int ldc);
```

?her2k

```
void cblas_cher2k(const enum CBLAS_ORDER Order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE Trans, const int N,
const int K, const void *alpha, const void *A, const int lda,
const void *B, const int ldb, const float beta, void *C, const
int ldc);
```

```
void cblas_zher2k(const enum CBLAS_ORDER Order, const enum
CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE Trans, const int N,
const int K, const void *alpha, const void *A, const int lda,
const void *B, const int ldb, const double beta, void *C, const
int ldc);
```

?symm

```
void cblas_ssymm(const enum CBLAS_ORDER Order, const enum
CBLAS_SIDE Side, const enum CBLAS_UPLO Uplo, const int M, const
int N, const float alpha, const float *A, const int lda, const
float *B, const int ldb, const float beta, float *C, const int
ldc);
```

```
void cblas_dsymm(const enum CBLAS_ORDER Order, const enum
CBLAS_SIDE Side, const enum CBLAS_UPLO Uplo, const int M, const
int N, const double alpha, const double *A, const int lda,
const double *B, const int ldb, const double beta, double *C,
const int ldc);
```

```
void cblas_csymm(const enum CBLAS_ORDER Order, const enum
CBLAS_SIDE Side, const enum CBLAS_UPLO Uplo, const int M, const
int N, const void *alpha, const void *A, const int lda, const
void *B, const int ldb, const void *beta, void *C, const int
ldc);
```

```
void cblas_zsymm(const enum CBLAS_ORDER Order, const enum
CBLAS_SIDE Side, const enum CBLAS_UPLO Uplo, const int M, const
int N, const void *alpha, const void *A, const int lda, const
void *B, const int ldb, const void *beta, void *C, const int
ldc);
```

?syrk

```
void cblas_ssyrc(const enum CBLAS_ORDER Order, const enum CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE Trans, const int N, const int K, const float alpha, const float *A, const int lda, const float beta, float *C, const int ldc);
```

```
void cblas_dsyrc(const enum CBLAS_ORDER Order, const enum CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE Trans, const int N, const int K, const double alpha, const double *A, const int lda, const double beta, double *C, const int ldc);
```

```
void cblas_csyrc(const enum CBLAS_ORDER Order, const enum CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE Trans, const int N, const int K, const void *alpha, const void *A, const int lda, const void *beta, void *C, const int ldc);
```

```
void cblas_zsyrc(const enum CBLAS_ORDER Order, const enum CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE Trans, const int N, const int K, const void *alpha, const void *A, const int lda, const void *beta, void *C, const int ldc);
```

?syr2k

```
void cblas_ssyrc2k(const enum CBLAS_ORDER Order, const enum CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE Trans, const int N, const int K, const float alpha, const float *A, const int lda, const float *B, const int ldb, const float beta, float *C, const int ldc);
```

```
void cblas_dsyrc2k(const enum CBLAS_ORDER Order, const enum CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE Trans, const int N, const int K, const double alpha, const double *A, const int lda, const double *B, const int ldb, const double beta, double *C, const int ldc);
```

```
void cblas_csyrc2k(const enum CBLAS_ORDER Order, const enum CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE Trans, const int N, const int K, const void *alpha, const void *A, const int lda, const void *B, const int ldb, const void *beta, void *C, const int ldc);
```

```
void cblas_zsyrc2k(const enum CBLAS_ORDER Order, const enum CBLAS_UPLO Uplo, const enum CBLAS_TRANSPOSE Trans, const int N, const int K, const void *alpha, const void *A, const int lda, const void *B, const int ldb, const void *beta, void *C, const int ldc);
```

?trmm

```
void cblas_strmm(const enum CBLAS_ORDER Order, const enum
CBLAS_SIDE Side, const enum CBLAS_UPLO Uplo, const enum
CBLAS_TRANSPOSE TransA, const enum CBLAS_DIAG Diag, const int
M, const int N, const float alpha, const float *A, const int
lda, float *B, const int ldb);
```

```
void cblas_dtrmm(const enum CBLAS_ORDER Order, const enum
CBLAS_SIDE Side, const enum CBLAS_UPLO Uplo, const enum
CBLAS_TRANSPOSE TransA, const enum CBLAS_DIAG Diag, const int
M, const int N, const double alpha, const double *A, const int
lda, double *B, const int ldb);
```

```
void cblas_ctrmm(const enum CBLAS_ORDER Order, const enum
CBLAS_SIDE Side, const enum CBLAS_UPLO Uplo, const enum
CBLAS_TRANSPOSE TransA, const enum CBLAS_DIAG Diag, const int
M, const int N, const void *alpha, const void *A, const int
lda, void *B, const int ldb);
```

```
void cblas_ztrmm(const enum CBLAS_ORDER Order, const enum
CBLAS_SIDE Side, const enum CBLAS_UPLO Uplo, const enum
CBLAS_TRANSPOSE TransA, const enum CBLAS_DIAG Diag, const int
M, const int N, const void *alpha, const void *A, const int
lda, void *B, const int ldb);
```

?trsm

```
void cblas_strsm(const enum CBLAS_ORDER Order, const enum
CBLAS_SIDE Side, const enum CBLAS_UPLO Uplo, const enum
CBLAS_TRANSPOSE TransA, const enum CBLAS_DIAG Diag, const int
M, const int N, const float alpha, const float *A, const int
lda, float *B, const int ldb);
```

```
void cblas_dtrsm(const enum CBLAS_ORDER Order, const enum
CBLAS_SIDE Side, const enum CBLAS_UPLO Uplo, const enum
CBLAS_TRANSPOSE TransA, const enum CBLAS_DIAG Diag, const int
M, const int N, const double alpha, const double *A, const int
lda, double *B, const int ldb);
```

```
void cblas_ctrsm(const enum CBLAS_ORDER Order, const enum
CBLAS_SIDE Side, const enum CBLAS_UPLO Uplo, const enum
CBLAS_TRANSPOSE TransA, const enum CBLAS_DIAG Diag, const int
M, const int N, const void *alpha, const void *A, const int
lda, void *B, const int ldb);
```

```
void cblas_ztrsm(const enum CBLAS_ORDER Order, const enum
CBLAS_SIDE Side, const enum CBLAS_UPLO Uplo, const enum
CBLAS_TRANSPOSE TransA, const enum CBLAS_DIAG Diag, const int
M, const int N, const void *alpha, const void *A, const int
lda, void *B, const int ldb);
```

Sparse CBLAS

This is an interface to [Sparse BLAS Routines and Functions](#), which perform a number of common vector operations on sparse vectors stored in compressed form.

Note that all index parameters, *indx*, are in C-type notation and vary in the range $[0 \dots N-1]$.

?axpyi

```
void cblas_saxpyi(const int N, const float alpha,
const float *X, const int *indx, float *Y);
void cblas_daxpyi(const int N, const double alpha,
const double *X, const int *indx, double *Y);
void cblas_caxpyi(const int N, const void *alpha,
const void *X, const int *indx, void *Y);
void cblas_zaxpyi(const int N, const void *alpha,
const void *X, const int *indx, void *Y);
```

?doti

```
float cblas_sdoti(const int N, const float *X,
const int *indx, const float *Y);
double cblas_ddoti(const int N, const double *X,
const int *indx, const double *Y);
```

?dotci

```
void cblas_cdotci_sub(const int N, const void *X, const int
*indx, const void *Y, void *dotui);
void cblas_zdotci_sub(const int N, const void *X, const int
*indx, const void *Y, void *dotui);
```

?dotui

```
void cblas_cdotui_sub(const int N, const void *X, const int
*indx, const void *Y, void *dotui);
void cblas_zdotui_sub(const int N, const void *X, const int
*indx, const void *Y, void *dotui);
```

?gthr

```
void cblas_sgthr(const int N, const float *Y, float *X,
const int *indx);
void cblas_dgthr(const int N, const double *Y, double *X,
const int *indx);
```

```
void cblas_cgthr(const int N, const void *Y, void *X,  
const int *indx);
```

```
void cblas_zgthr(const int N, const void *Y, void *X,  
const int *indx);
```

?gthrz

```
void cblas_sgthrz(const int N, float *Y, float *X,  
const int *indx);
```

```
void cblas_dgthrz(const int N, double *Y, double *X,  
const int *indx);
```

```
void cblas_cgthrz(const int N, void *Y, void *X,  
const int *indx);
```

```
void cblas_zgthrz(const int N, void *Y, void *X,  
const int *indx);
```

?roti

```
void cblas_sroti(const int N, float *X, const int *indx,  
float *Y, const float c, const float s);
```

```
void cblas_droti(const int N, double *X, const int *indx,  
double *Y, const double c, const double s);
```

?sctr

```
void cblas_ssctr(const int N, const float *X, const int *indx,  
float *Y);
```

```
void cblas_dsctr(const int N, const double *X, const int *indx,  
double *Y);
```

```
void cblas_csctr(const int N, const void *X, const int *indx,  
void *Y);
```

```
void cblas_zsctr(const int N, const void *X, const int *indx,  
void *Y);
```

Glossary

A^H	Denotes the conjugate of a general matrix A . <i>See also</i> conjugate matrix.
A^T	Denotes the transpose of a general matrix A . <i>See also</i> transpose.
band matrix	A general m by n matrix A such that $a_{ij} = 0$ for $ i - j > l$, where $1 < l < \min(m, n)$. For example, any tridiagonal matrix is a band matrix.
band storage	A special storage scheme for band matrices. A matrix is stored in a two-dimensional array: columns of the matrix are stored in the corresponding columns of the array, and <i>diagonals</i> of the matrix are stored in rows of the array.
BLAS	Abbreviation for Basic Linear Algebra Subprograms. These subprograms implement vector, matrix-vector, and matrix-matrix operations.
Bunch-Kaufman factorization	Representation of a real symmetric or complex Hermitian matrix A in the form $A = PUDU^H P^T$ (or $A = PLDL^H P^T$) where P is a permutation matrix, U and L are upper and lower triangular matrices with unit diagonal, and D is a Hermitian block-diagonal matrix with 1-by-1 and 2-by-2 diagonal blocks. U and L have 2-by-2 unit diagonal blocks corresponding to the 2-by-2 blocks of D .

c	When found as the first letter of routine names, c indicates the usage of single-precision complex data type.
CBLAS	C interface to the BLAS. <i>See</i> BLAS.
Cholesky factorization	Representation of a symmetric positive-definite or, for complex data, Hermitian positive-definite matrix A in the form $A = U^H U$ or $A = L L^H$, where L is a lower triangular matrix and U is an upper triangular matrix.
condition number	The number $\kappa(A)$ defined for a given square matrix A as follows: $\kappa(A) = \ A \ \ A^{-1} \ $.
conjugate matrix	The matrix A^H defined for a given general matrix A as follows: $(A^H)_{ij} = (a_{ji})^*$.
conjugate number	The conjugate of a complex number $z = a + bi$ is $z^* = a - bi$.
d	When found as the first letter of routine names, d indicates the usage of double-precision real data type.
dot product	The number denoted $x \cdot y$ and defined for given vectors x and y as follows: $x \cdot y = \sum_i x_i y_i$. Here x_i and y_i stand for the i th elements of x and y , respectively.
double precision	A floating-point data type. On Intel® processors, this data type allows you to store real numbers x such that $2.23 \cdot 10^{-308} < x < 1.79 \cdot 10^{308}$. For this data type, the machine precision ϵ is approximately 10^{-15} , which means that double-precision numbers usually contain no more than 15 significant decimal digits. For more information, refer to <i>Pentium® Processor Family Developer's Manual, Volume 3: Architecture and Programming Manual</i> .
eigenvalue	<i>See</i> eigenvalue problem.

eigenvalue problem	A problem of finding non-zero vectors x and numbers λ (for a given square matrix A) such that $Ax = \lambda x$. Here the numbers λ are called the <i>eigenvalues</i> of the matrix A and the vectors x are called the <i>eigenvectors</i> of the matrix A .
eigenvector	<i>See</i> eigenvalue problem.
elementary reflector (Householder matrix)	Matrix of a general form $H = I - \tau vv^T$, where v is a column vector and τ is a scalar. In LAPACK elementary reflectors are used, for example, to represent the matrix Q in the QR factorization (the matrix Q is represented as a product of elementary reflectors).
factorization	Representation of a matrix as a product of matrices. <i>See also</i> Bunch-Kaufman factorization, Cholesky factorization, LU factorization, LQ factorization, QR factorization, Schur factorization.
FFTs	Abbreviation for Fast Fourier Transforms. <i>See</i> Chapter 3 of this book.
full storage	A storage scheme allowing you to store matrices of any kind. A matrix A is stored in a two-dimensional array \mathbf{a} , with the matrix element a_{ij} stored in the array element $\mathbf{a}(i, j)$.
Hermitian matrix	A square matrix A that is equal to its conjugate matrix A^H . The conjugate A^H is defined as follows: $(A^H)_{ij} = (a_{ji})^*$.
I	<i>See</i> identity matrix.
identity matrix	A square matrix I whose diagonal elements are 1, and off-diagonal elements are 0. For any matrix A , $AI = A$ and $IA = A$.
in-place	Qualifier of an operation. A function that performs its operation in-place takes its input from an array and returns its output to the same array.

inverse matrix	The matrix denoted as A^{-1} and defined for a given square matrix A as follows: $AA^{-1} = A^{-1}A = I$. A^{-1} does not exist for singular matrices A .
LQ factorization	Representation of an m by n matrix A as $A = LQ$ or $A = (L\ 0)Q$. Here Q is an n by n orthogonal (unitary) matrix. For $m \leq n$, L is an m by m lower triangular matrix with real diagonal elements; for $m > n$,
	$L = \begin{bmatrix} L_1 \\ L_2 \end{bmatrix}$
	where L_1 is an n by n lower triangular matrix, and L_2 is a rectangular matrix.
LU factorization	Representation of a general m by n matrix A as $A = PLU$, where P is a permutation matrix, L is lower triangular with unit diagonal elements (lower trapezoidal if $m > n$) and U is upper triangular (upper trapezoidal if $m < n$).
machine precision	The number ϵ determining the precision of the machine representation of real numbers. For Intel® architecture, the machine precision is approximately 10^{-7} for single-precision data, and approximately 10^{-15} for double-precision data. The precision also determines the number of significant decimal digits in the machine representation of real numbers. <i>See also</i> double precision and single precision.
MKL	Abbreviation for Math Kernel Library.
orthogonal matrix	A real square matrix A whose transpose and inverse are equal, that is, $A^T = A^{-1}$, and therefore $AA^T = A^T A = I$. All eigenvalues of an orthogonal matrix have the absolute value 1.
packed storage	A storage scheme allowing you to store symmetric, Hermitian, or triangular matrices more compactly. The upper or lower triangle of a matrix is packed by columns in a one-dimensional array.

positive-definite matrix	A square matrix A such that $Ax \cdot x > 0$ for any non-zero vector x . Here \cdot denotes the dot product.
QR factorization	Representation of an m by n matrix A as $A = QR$, where Q is an m by m orthogonal (unitary) matrix, and R is n by n upper triangular with real diagonal elements (if $m \geq n$) or trapezoidal (if $m < n$) matrix.
s	When found as the first letter of routine names, s indicates the usage of single-precision real data type.
Schur factorization	Representation of a square matrix A in the form $A = ZTZ^H$. Here T is an upper quasi-triangular matrix (for complex A , triangular matrix) called the Schur form of A ; the matrix Z is orthogonal (for complex A , unitary). Columns of Z are called Schur vectors.
single precision	A floating-point data type. On Intel [®] processors, this data type allows you to store real numbers x such that $1.18 \cdot 10^{-38} < x < 3.40 \cdot 10^{38}$. For this data type, the machine precision (ϵ) is approximately 10^{-7} , which means that single-precision numbers usually contain no more than 7 significant decimal digits. For more information, refer to <i>Pentium[®] Processor Family Developer's Manual, Volume 3: Architecture and Programming Manual</i> .
singular matrix	A matrix whose determinant is zero. If A is a singular matrix, the inverse A^{-1} does not exist, and the system of equations $Ax = b$ does not have a unique solution (that is, there exist no solutions or an infinite number of solutions).
singular value	The numbers defined for a given general matrix A as the eigenvalues of the matrix AA^H . <i>See also</i> SVD.
SMP	Abbreviation for Symmetric MultiProcessing. The MKL offers performance gains through parallelism provided by the SMP feature.

sparse BLAS	Routines performing basic vector operations on sparse vectors. Sparse BLAS routines take advantage of vectors' sparsity: they allow you to store only non-zero elements of vectors. <i>See</i> BLAS.
sparse vectors	Vectors in which most of the components are zeros.
storage scheme	The way of storing matrices. <i>See</i> full storage, packed storage, and band storage.
SVD	Abbreviation for Singular Value Decomposition. <i>See also</i> Singular value decomposition section in Chapter 5.
symmetric matrix	A square matrix A such that $a_{ij} = a_{ji}$.
transpose	The transpose of a given matrix A is a matrix A^T such that $(A^T)_{ij} = a_{ji}$ (rows of A become columns of A^T , and columns of A become rows of A^T).
trapezoidal matrix	A matrix A such that $A = (A_1A_2)$, where A_1 is an upper triangular matrix, A_2 is a rectangular matrix.
triangular matrix	A matrix A is called an upper (lower) triangular matrix if all its subdiagonal elements (superdiagonal elements) are zeros. Thus, for an upper triangular matrix $a_{ij} = 0$ when $i > j$; for a lower triangular matrix $a_{ij} = 0$ when $i < j$.
tridiagonal matrix	A matrix whose non-zero elements are in three diagonals only: the leading diagonal, the first subdiagonal, and the first super-diagonal.
unitary matrix	A complex square matrix A whose conjugate and inverse are equal, that is, that is, $A^H = A^{-1}$, and therefore $AA^H = A^HA = I$. All eigenvalues of a unitary matrix have the absolute value 1.
VML	Abbreviation for Vector Mathematical Library. <i>See</i> Chapter 6 of this book.
z	When found as the first letter of routine names, z indicates the usage of double-precision complex data type.

Index

Routines

?asum, 2-5
?axpy, 2-6
?axpyi, 2-116
?bdsqr, 5-94, 5-98
?copy, 2-7
?dot, 2-8
?dotc, 2-9
?dotci, 2-119
?doti, 2-118
?dotu, 2-10
?dotui, 2-120
?fft1d, 3-4, 3-8
?fft1dc, 3-5, 3-10, 3-15
?fft2d, 3-19, 3-22, 3-28
?fft2dc, 3-20, 3-24, 3-29
?gbbrd, 5-79
?gbcon, 4-65
?gbmv, 2-24
?gbrfs, 4-95
?gbtf2, 6-21
?gbtrf, 4-10
?gbtrs, 4-36
?gebak, 5-193
?gebal, 5-190
?gebd2, 6-23
?gebrd, 5-76
?gecon, 4-63
?gees, 5-379
?geesx, 5-384
?geev, 5-390
?geevx, 5-394
?gehd2, 6-26
?gehrd, 5-178
?gelq2, 6-29
?gelqf, 5-25, 5-36
?gels, 5-279
?gelsd, 5-289
?gelss, 5-286
?gelsy, 5-282
?gemm, 2-83
?gemv, 2-27
?geql2, 6-31
?geqpf, 5-11, 5-14
?geqr2, 6-33
?geqrf, 5-8, 5-48, 5-60, 5-68, 5-71
?ger, 2-30
?gerc, 2-31
?gerfs, 4-92, 4-98
?gerq2, 6-35
?geru, 2-33
?gesc2, 6-37

?gesdd, 5-405
?gesvd, 5-400
?getc2, 6-39
?getf2, 6-40
?getrf, 4-7
?getri, 4-133
?getrs, 4-34
?ggbak, 5-233
?ggbal, 5-230
?gges, 5-482
?ggesx, 5-489
?ggev, 5-497
?ggevx, 5-502
?ggglm, 5-296
?gghrd, 5-226
?gglse, 5-293
?ggsvd, 5-409
?ggsvp, 5-267
?gtcon, 4-67
?gthr, 2-121
?gthrz, 2-122
?gttrf, 4-12
?gttrs, 4-38, 6-42
?hbev, 5-348
?hbevd, 5-353
?hbevx, 5-361
?hbgst, 5-169
?hbgv, 5-463
?hbgvd, 5-469
?hbgvx, 5-477
?hbrd, 5-130
?hecon, 4-80
?heev, 5-301
?heevd, 5-306
?heevr, 5-322
?heevx, 5-313
?hegst, 5-160
?hegv, 5-419
?hegvd, 5-425
?hegvx, 5-434
?hemm, 2-86
?hemv, 2-38
?her, 2-40
?her2, 2-42
?her2k, 2-92
?herfs, 4-116
?herk, 2-89
?hetrd, 5-111
?hetrf, 4-25
?hetri, 4-141
?hetrs, 4-51
?hgeqz, 5-235
?hpcon, 4-84
?hpev, 5-329
?hpevd, 5-334
?hpevx, 5-342
?hpgst, 5-164
?hpgv, 5-442
?hpgvd, 5-448
?hpgvx, 5-456
?hpmv, 2-44
?hpr, 2-47
?hpr2, 2-49
?hprfs, 4-122
?hptrd, 5-122
?hptrf, 4-31
?hptri, 4-145
?hptrs, 4-55
?hsein, 5-199
?hseqr, 5-195
?labrd, 6-45
?lacgv, 6-1
?lacon, 6-48
?lacpy, 6-50

?lacrm, 6-2
?lact, 6-3
?ladiv, 6-51
?lae2, 6-52
?laebz, 6-54
?laed0, 6-60
?laed1, 6-64
?laed2, 6-67
?laed3, 6-70
?laed4, 6-73
?laed5, 6-74
?laed6, 6-75
?laed7, 6-78
?laed8, 6-83
?laed9, 6-87
?laeda, 6-89
?laein, 6-92
?laesy, 6-4
?laev2, 6-95
?laexc, 6-98
?lag2, 6-100
?lags2, 6-103
?lagtf, 6-105
?lagtm, 6-108
?lagts, 6-110
?lagv2, 6-113
?lahqr, 6-115
?lahrd, 6-118
?laic1, 6-121
?laln2, 6-124
?lals0, 6-127
?lalsa, 6-132
?lalsd, 6-136
?lamc1, 6-140
?lamc2, 6-141
?lamc3, 6-142
?lamc4, 6-143
?lamc5, 6-144
?lamch, 6-139
?lamrg, 6-145
?langb, 6-146
?lange, 6-148
?langt, 6-149
?lanhb, 6-154
?lanhe, 6-163
?lanhp, 6-158
?lanhs, 6-151
?lansb, 6-152
?lansp, 6-156
?lanst/?lanht, 6-159
?lansy, 6-161
?lantb, 6-164
?lantp, 6-167
?lantr, 6-169
?lanv2, 6-171
?lapll, 6-172
?lapmt, 6-174
?lapy2, 6-175
?lapy3, 6-176
?laqgb, 6-176
?laqge, 6-179
?laqp2, 6-181
?laqps, 6-183
?laqsb, 6-185
?laqsp, 6-187
?laqsy, 6-189
?laqtr, 6-191
?lar1v, 6-194
?lar2v, 6-196
?larf, 6-198
?larfb, 6-200
?larfg, 6-202
?larft, 6-204
?larfx, 6-207

?largv, 6-208
?larnv, 6-210
?larrb, 6-212
?larre, 6-214
?larf, 6-216
?larv, 6-218
?lartg, 6-221
?lartv, 6-223
?laruv, 6-224
?larz, 6-225
?larzb, 6-227
?larzt, 6-230
?las2, 6-233
?lascl, 6-234
?lasd0, 6-236
?lasd1, 6-238
?lasd2, 6-241
?lasd3, 6-246
?lasd4, 6-249
?lasd5, 6-251
?lasd6, 6-252
?lasd7, 6-257
?nrm2, 2-11
?opgtr, 5-119
?opmtr, 5-120
?orgbr, 5-82
?orghr, 5-180
?orglq, 5-28, 5-38, 5-40
?orgqr, 5-17, 5-50, 5-52
?orgtr, 5-107
?ormbr, 5-85
?ormhr, 5-182
?ormlq, 5-30, 5-42, 5-45, 5-54, 5-57, 5-62, 5-65
?ormqr, 5-19
?ormtr, 5-109
?pbcon, 4-74
?pbrfs, 4-107
?pbstf, 5-172
?pbtrf, 4-18
?pbtrs, 4-45
?pocon, 4-70
?porfs, 4-101, 4-110
?potrf, 4-14
?potri, 4-135
?potrs, 4-41
?ppcon, 4-72
?pprfs, 4-104
?pptrf, 4-16
?pptri, 4-137
?pptrs, 4-43
?ptcon, 4-76
?pteqr, 5-146
?pttrf, 4-20
?pttrs, 4-47
?rot, 2-12
?rot (complex), 6-6
?rotg, 2-14
?roti, 2-123
?rotm, 2-15
?rotmg, 2-17
?sbev, 5-346
?sbevd, 5-350
?sbevx, 5-357
?sbgst, 5-166
?sbgv, 5-460
?sbgvd, 5-466
?sbgvx, 5-473
?sbmv, 2-51
?sbtrd, 5-128
?scal, 2-18
?sctr, 2-124
?spcon, 4-82
?spev, 5-327
?spevd, 5-331

?spevx, 5-338
?spgst, 5-162
?spgv, 5-439
?spgvd, 5-445
?spgvx, 5-452
?spm, 2-54, 6-7
?spr, 2-56, 6-9
?spr2, 2-58
?sprfs, 4-119
?sptrd, 5-117
?sptrf, 4-28
?sptri, 4-143
?sptrs, 4-53
?stebz, 5-141, 5-149
?stein, 5-152
?steqr, 5-134, 5-137
?sterf, 5-132
?stev, 5-365
?stevd, 5-367
?stevr, 5-374
?stevx, 5-370
?sum1, 6-20
?swap, 2-20
?sycon, 4-78, 5-154
?syev, 5-299
?syevd, 5-303
?syevr, 5-317
?syevx, 5-309
?sygst, 5-158
?sygv, 5-416
?sygvd, 5-422
?sygvx, 5-429
?symm, 2-96
?symv, 2-60
?symv (complex), 6-11
?syr, 2-62
?syr (complex), 6-13
?syr2, 2-64
?syr2k, 2-103
?syrfs, 4-113
?syrk, 2-100
?sytrd, 5-105
?sytrf, 4-22
?sytri, 4-139
?sytrs, 4-49
?tbcon, 4-90
?tbmv, 2-66
?tbsv, 2-69
?tbtrs, 4-61
?tgevc, 5-242
?tgexc, 5-247
?tgsen, 5-250
?tgsja, 5-271
?tgsna, 5-261
?tgsyl, 5-256
?tpcon, 4-88
?tpmv, 2-72
?tprfs, 4-127
?tps, 2-75
?tptri, 4-148
?tprts, 4-59
?trcon, 4-86
?trevc, 5-205
?trexc, 5-215
?trmm, 2-107
?trmv, 2-77
?trrfs, 4-124
?trsen, 5-217
?trsm, 2-110
?trsna, 5-210
?trsv, 2-79
?trsyl, 5-222
?trtri, 4-147
?trtrs, 4-57

?ungbr, 5-88
 ?unghr, 5-185
 ?unglq, 5-32
 ?ungqr, 5-21
 ?ungtr, 5-113
 ?unmbr, 5-91
 ?unmhr, 5-187
 ?unmlq, 5-34
 ?unmqr, 5-23
 ?unmtr, 5-115
 ?upgtr, 5-124
 ?upmtr, 5-125

A

absolute value of a vector element
 largest, 2-21
 smallest, 2-22
 accuracy modes, in VML, 7-2
 adding magnitudes of the vector elements, 2-5
 arguments
 matrix, A-4
 sparse vector, 2-114
 vector, A-1

B

balancing a matrix, 5-190
 band storage scheme, A-4
 Bernoulli, 8-48
 bidiagonal matrix, 5-74
 Binomial, 8-52
 BLAS Level 1 functions
 ?asum, 2-4, 2-5
 ?dot, 2-4, 2-8
 ?dotc, 2-4, 2-9
 ?dotu, 2-4, 2-10
 ?nrm2, 2-4, 2-11
 code example, B-1, B-2
 i?amax, 2-4, 2-21

 i?amin, 2-4, 2-22
 BLAS Level 1 routines
 ?axpy, 2-4, 2-6
 ?copy, 2-4, 2-7
 ?rot, 2-4, 2-12
 ?rotg, 2-4, 2-14
 ?rotm, 2-15
 ?rotmg, 2-17
 ?scal, 2-4, 2-18
 ?swap, 2-4, 2-20
 code example, B-2
 BLAS Level 2 routines
 ?gbmv, 2-23, 2-24
 ?gemv, 2-23, 2-27
 ?ger, 2-23, 2-30
 ?gerc, 2-23, 2-31
 ?geru, 2-23, 2-33
 ?hbm, 2-23, 2-35
 ?hemv, 2-23, 2-38
 ?her, 2-23, 2-40
 ?her2, 2-23, 2-42
 ?hpmv, 2-23, 2-44
 ?hpr, 2-23, 2-47
 ?hpr2, 2-23, 2-49
 ?sbmv, 2-23, 2-51
 ?spmv, 2-23, 2-54
 ?spr, 2-23, 2-56
 ?spr2, 2-23, 2-58
 ?symv, 2-23, 2-60
 ?syr, 2-23, 2-62
 ?syr2, 2-23, 2-64
 ?tbmv, 2-24, 2-66
 ?tbsv, 2-24, 2-69
 ?tpmv, 2-24, 2-72
 ?tpsv, 2-24, 2-75
 ?trmv, 2-24, 2-77
 ?trsv, 2-24, 2-79
 code example, B-3, B-4
 BLAS Level 3 routines
 ?gemm, 2-82, 2-83
 ?hemm, 2-82, 2-86
 ?her2k, 2-82, 2-92
 ?herk, 2-82, 2-89

- ?symm, 2-82, 2-96
 - ?syr2k, 2-82, 2-103
 - ?syrk, 2-82, 2-100
 - ?trmm, 2-82, 2-107
 - ?trsm, 2-82, 2-110
 - code example, B-4, B-5
 - BLAS routines
 - matrix arguments, A-4
 - routine groups, 1-6, 2-1
 - vector arguments, A-1
 - block-splitting method, 8-6
 - Bunch-Kaufman factorization, 4-7
 - Hermitian matrix, 4-25
 - packed storage, 4-31
 - symmetric matrix, 4-22
 - packed storage, 4-28
- C**
-
- C interface, 3-3
 - Cauchy, 8-33
 - CBLAS, 1
 - arguments, 1
 - level 1 (vector operations), 3
 - level 2 (matrix-vector operations), 6
 - level 3 (matrix-matrix operations), 14
 - sparse BLAS, 18
 - Cholesky factorization
 - Hermitian positive-definite matrix, 4-14
 - band storage, 4-18
 - packed storage, 4-16
 - symmetric positive-definite matrix, 4-14
 - band storage, 4-18
 - packed storage, 4-16
 - code examples
 - BLAS Level 1 function, B-1
 - BLAS Level 1 routine, B-2
 - BLAS Level 2 routine, B-3
 - BLAS Level 3 routine, B-4
 - CommitDescriptor, 9-17
 - complex-to-complex one-dimensional FFTs, 3-3
 - complex-to-complex two-dimensional FFTs,
 - 3-18
 - complex-to-real one-dimensional FFTs, 3-12
 - complex-to-real two-dimensional FFTs, 3-27
 - Computational Routines, 5-6
 - ComputeBackward, 9-23
 - ComputeForward, 9-21
 - condition number
 - band matrix, 4-65
 - general matrix, 4-63
 - Hermitian matrix, 4-80
 - packed storage, 4-84
 - Hermitian positive-definite matrix, 4-70
 - band storage, 4-74
 - packed storage, 4-72
 - tridiagonal, 4-76
 - symmetric matrix, 4-78, 5-154
 - packed storage, 4-82
 - symmetric positive-definite matrix, 4-70
 - band storage, 4-74
 - packed storage, 4-72
 - tridiagonal, 4-76
 - triangular matrix, 4-86
 - band storage, 4-90
 - packed storage, 4-88
 - tridiagonal matrix, 4-67
 - configuration parameters, in DFTI, 9-6
 - Continuous Distribution Generators, 8-20
 - converting a sparse vector into compressed storage form, 2-121
 - and writing zeros to the original vector, 2-122
 - converting compressed sparse vectors into full storage form, 2-124
 - CopyDescriptor, 9-18
 - copying vectors, 2-7
 - CopyStream, 8-12
 - CreateDescriptor, 9-15
- D**
-
- data structure requirements for FFTs, 3-3

data type
in VML, 7-2
shorthand, 1-8

DeleteStream, 8-11

Descriptor configuration, in DFTI, 9-10

Descriptor Manipulation, in DFTI, 9-9

DFT computation, 9-10

DFT routines
descriptor configuration
 GetValue, 9-33
 SetValue, 9-31
descriptor manipulation
 CommitDescriptor, 9-17
 CopyDescriptor, 9-18
 CreateDescriptor, 9-15
 FreeDescriptor, 9-20
DFT computation
 ComputeBackward, 9-23
 ComputeForward, 9-21
status checking
 ErrorClass, 9-11
 ErrorMessage, 9-13

dimension, A-1

Discrete Distribution Generators, 8-20

Discrete Fourier Transform
CommitDescriptor, 9-17
ComputeBackward, 9-23
ComputeForward, 9-21
CopyDescriptor, 9-18
CreateDescriptor, 9-15
ErrorClass, 9-11
ErrorMessage, 9-13
FreeDescriptor, 9-20
GetValue, 9-33
SetValue, 9-31

dot product
complex vectors, conjugated, 2-9
complex vectors, unconjugated, 2-10
real vectors, 2-8
sparse complex vectors, 2-120
sparse complex vectors, conjugated, 2-119
sparse real vectors, 2-118

Driver Routines, 5-278

E

eigenvalue problems
general matrix, 5-174, 5-225
generalized form, 5-157
Hermitian matrix, 5-101
symmetric matrix, 5-101

eigenvalues. *See* eigenvalue problems

eigenvectors. *See* eigenvalue problems

error diagnostics, in VML, 7-6

Error reporting routine, XERBLA, 2-1

ErrorClass, 9-11

ErrorMessage, 9-13

errors in solutions of linear equations
general matrix, 4-92, 4-98
 band storage, 4-95
Hermitian matrix, 4-116
 packed storage, 4-122
Hermitian positive-definite matrix, 4-101, 4-110
 band storage, 4-107
 packed storage, 4-104
symmetric matrix, 4-113
 packed storage, 4-119
symmetric positive-definite matrix, 4-101, 4-110
 band storage, 4-107
 packed storage, 4-104
triangular matrix, 4-124
 band storage, 4-130
 packed storage, 4-127

Euclidean norm
of a vector, 2-11

Exponential, 8-26

F

factorization
See also triangular factorization
Bunch-Kaufman, 4-7

-
- Cholesky, 4-7
 - LU, 4-7
 - orthogonal (LQ, QR), 5-7
 - fast Fourier transforms, 1-3, 9-1
 - C interface, 3-3
 - data storage types, 3-2
 - data structure requirements, 3-3
 - routines
 - ?fft1d, 3-4, 3-8, 3-13
 - ?fft1dc, 3-5, 3-10, 3-15
 - ?fft2d, 3-19, 3-22, 3-28
 - ?fft2dc, 3-20, 3-24, 3-29
 - FFT. *See* fast Fourier transforms
 - finding
 - element of a vector with the largest absolute value, 2-21
 - element of a vector with the smallest absolute value, 2-22
 - font conventions, 1-8
 - forward or inverse FFTs, 3-4, 3-5, 3-19, 3-20
 - Fourier transforms
 - mixed radix, 9-1
 - multi-dimensional, 9-1
 - FreeDescriptor, 9-20
 - full storage scheme, A-4
 - function name conventions, in VML, 7-2
- G**
-
- gathering sparse vector's elements into
 - compressed form, 2-121
 - and writing zeros to these elements, 2-122
 - Gaussian, 8-23
 - general matrix
 - eigenvalue problems, 5-174, 5-225
 - estimating the condition number, 4-63
 - band storage, 4-65
 - inverting the matrix, 4-133
 - LQ factorization, 5-25, 5-36
 - LU factorization, 4-7
 - band storage, 4-10
 - matrix-vector product, 2-27
 - band storage, 2-24
 - QR factorization, 5-8, 5-48, 5-60, 5-68, 5-71
 - with pivoting, 5-11, 5-14
 - rank-1 update, 2-30
 - rank-1 update, conjugated, 2-31
 - rank-1 update, unconjugated, 2-33
 - scalar-matrix-matrix product, 2-83
 - solving systems of linear equations, 4-34
 - band storage, 4-36
 - generalized eigenvalue problems, 5-157
 - See also* LAPACK routines, generalized eigenvalue problems
 - complex Hermitian-definite problem, 5-160
 - band storage, 5-169
 - packed storage, 5-164
 - real symmetric-definite problem, 5-158
 - band storage, 5-166
 - packed storage, 5-162
 - Geometric, 8-50
 - GetBrngProperties, 8-63
 - GetStreamStateBrng, 8-19
 - GetValue, 9-33
 - GFSR, 8-4
 - Givens rotation
 - modified Givens transformation parameters, 2-17
 - of sparse vectors, 2-123
 - parameters, 2-14
 - Gumbel, 8-41
- H**
-
- Hermitian matrix, 5-101, 5-157
 - Bunch-Kaufman factorization, 4-25
 - packed storage, 4-31
 - estimating the condition number, 4-80
 - packed storage, 4-84
 - generalized eigenvalue problems, 5-157
 - inverting the matrix, 4-141
 - packed storage, 4-145
 - matrix-vector product, 2-38
 - band storage, 2-35

- packed storage, 2-44
- rank-1 update, 2-40
 - packed storage, 2-47
- rank-2 update, 2-42
 - packed storage, 2-49
- rank-2k update, 2-92
- rank-n update, 2-89
- scalar-matrix-matrix product, 2-86
- solving systems of linear equations, 4-51
 - packed storage, 4-55

Hermitian positive-definite matrix

- Cholesky factorization, 4-14
 - band storage, 4-18
 - packed storage, 4-16
- estimating the condition number, 4-70
 - band storage, 4-74
 - packed storage, 4-72
- inverting the matrix, 4-135
 - packed storage, 4-137
- solving systems of linear equations, 4-41
 - band storage, 4-45
 - packed storage, 4-43

Hypergeometric, 8-54

I

- i?amax, 2-21
- i?amin, 2-22
- i?max1, 6-15
- ilaenv, 6-16
- increment, A-1
- inverse matrix. *See* inverting a matrix
- inverting a matrix
 - general matrix, 4-133
 - Hermitian matrix, 4-141
 - packed storage, 4-145
 - Hermitian positive-definite matrix, 4-135
 - packed storage, 4-137
 - symmetric matrix, 4-139
 - packed storage, 4-143
 - symmetric positive-definite matrix, 4-135
 - packed storage, 4-137

- triangular matrix, 4-147
 - packed storage, 4-148

L

LAPACK routines

condition number estimation

- ?gbcon, 4-65
- ?gecon, 4-63
- ?gtcon, 4-67
- ?hecon, 4-80
- ?hpcon, 4-84
- ?pbcon, 4-74
- ?pocon, 4-70
- ?ppcon, 4-72
- ?ptcon, 4-76
- ?spcon, 4-82
- ?sycon, 4-78, 5-154
- ?tbcon, 4-90
- ?tpcon, 4-88
- ?trcon, 4-86

generalized eigenvalue problems

- ?hbgst, 5-169
- ?hegst, 5-160
- ?hpgst, 5-164
- ?pbstf, 5-172
- ?sbgst, 5-166
- ?spgst, 5-162
- ?sygst, 5-158

LQ factorization

- ?gelqf, 5-25, 5-36
- ?orglq, 5-28, 5-38, 5-40
- ?ormlq, 5-30, 5-42, 5-45, 5-54, 5-57, 5-62, 5-65
- ?unglq, 5-32
- ?unmlq, 5-34

matrix inversion

- ?getri, 4-133
- ?hetri, 4-141
- ?hptri, 4-145
- ?potri, 4-135
- ?pptri, 4-137
- ?sptri, 4-143
- ?sytri, 4-139

-
- ?tptri, 4-148
 - ?trtri, 4-147
 - nonsymmetric eigenvalue problems
 - ?gebak, 5-193
 - ?gebal, 5-190
 - ?gehrd, 5-178
 - ?hsein, 5-199
 - ?hseqr, 5-195
 - ?orghr, 5-180
 - ?ormhr, 5-182
 - ?trevc, 5-205
 - ?trexc, 5-215
 - ?trsen, 5-217
 - ?trsna, 5-210
 - ?unghr, 5-185
 - ?unmhr, 5-187
 - QR factorization
 - ?geqpf, 5-11, 5-14
 - ?geqrf, 5-8, 5-48, 5-60, 5-68, 5-71
 - ?orgqr, 5-17, 5-50, 5-52
 - ?ormqr, 5-19
 - ?ungqr, 5-21
 - ?unmqr, 5-23
 - singular value decomposition
 - ?bdsqr, 5-94, 5-98
 - ?gbbrd, 5-79
 - ?gebrd, 5-76
 - ?orgbr, 5-82
 - ?ormbr, 5-85
 - ?ungbr, 5-88
 - ?unmbr, 5-91
 - solution refinement and error estimation
 - ?gbrfs, 4-95
 - ?gerfs, 4-92, 4-98
 - ?herfs, 4-116
 - ?hprfs, 4-122
 - ?pbrfs, 4-107
 - ?porfs, 4-101, 4-110
 - ?pprfs, 4-104
 - ?sprfs, 4-119
 - ?syrf, 4-113
 - ?tbrfs, 4-130
 - ?tprfs, 4-127
 - ?trrfs, 4-124
 - solving linear equations
 - ?gbtrs, 4-36
 - ?getrs, 4-34
 - ?gttrs, 4-38, 6-42
 - ?hetrs, 4-51
 - ?hptrs, 4-55
 - ?pbtrs, 4-45
 - ?potrs, 4-41
 - ?pptrs, 4-43
 - ?pttrs, 4-47
 - ?sptrs, 4-53
 - ?sytrs, 4-49
 - ?tbtrs, 4-61
 - ?tptrs, 4-59
 - ?trtrs, 4-57
 - Sylvester's equation
 - ?trsyl, 5-222
 - symmetric eigenvalue problems
 - ?hbevd, 5-353
 - ?hbtrd, 5-130
 - ?heevd, 5-306
 - ?hetrd, 5-111
 - ?hpevd, 5-334
 - ?hptrd, 5-122
 - ?opgtr, 5-119
 - ?opmtr, 5-120
 - ?orgtr, 5-107
 - ?ormtr, 5-109
 - ?pteqr, 5-146
 - ?sbevd, 5-350
 - ?sbtrd, 5-128
 - ?spevd, 5-331
 - ?sptrd, 5-117
 - ?stebz, 5-141, 5-149
 - ?stein, 5-152
 - ?steqr, 5-134, 5-137
 - ?sterf, 5-132
 - ?stevd, 5-367
 - ?syevd, 5-303
 - ?sytrd, 5-105
 - ?ungtr, 5-113
 - ?unmtr, 5-115
 - ?upgtr, 5-124
 - ?upmtr, 5-125

- triangular factorization
 - ?gbtrf, 4-10
 - ?getrf, 4-7
 - ?gttrf, 4-12
 - ?hetrf, 4-25
 - ?hprrf, 4-31
 - ?pbtrf, 4-18
 - ?potrf, 4-14
 - ?pptrf, 4-16
 - ?pttrf, 4-20
 - ?sprf, 4-28
 - ?sytrf, 4-22
 - Laplace, 8-28
 - leading dimension, A-6
 - leapfrog method, 8-6
 - LeapfrogStream, 8-13
 - length. *See* dimension
 - linear combination of vectors, 2-6
 - Linear Congruential Generator, 8-3
 - linear equations, solving
 - general matrix, 4-34
 - band storage, 4-36
 - Hermitian matrix, 4-51
 - packed storage, 4-55
 - Hermitian positive-definite matrix, 4-41
 - band storage, 4-45
 - packed storage, 4-43
 - symmetric matrix, 4-49
 - packed storage, 4-53
 - symmetric positive-definite matrix, 4-41
 - band storage, 4-45
 - packed storage, 4-43
 - triangular matrix, 4-57
 - band storage, 4-61
 - packed storage, 4-59
 - tridiagonal matrix, 4-38, 4-47, 6-42
 - Lognormal, 8-38
 - LQ factorization, 5-6
 - computing the elements of
 - orthogonal matrix Q, 5-28, 5-38, 5-40
 - unitary matrix Q, 5-32
 - lsame, 6-18
 - lsamen, 6-19
 - LU factorization, 4-7
 - band matrix, 4-10
 - tridiagonal matrix, 4-12
- ## M
-
- matrix arguments, A-4
 - column-major ordering, A-2, A-6
 - example, A-7
 - leading dimension, A-6
 - number of columns, A-6
 - number of rows, A-6
 - transposition parameter, A-6
 - matrix equation
 - $AX = B$, 2-110, 4-5, 4-33
 - matrix one-dimensional substructures, A-2
 - matrix-matrix operation
 - product
 - general matrix, 2-83
 - rank-2k update
 - Hermitian matrix, 2-92
 - symmetric matrix, 2-103
 - rank-n update
 - Hermitian matrix, 2-89
 - symmetric matrix, 2-100
 - scalar-matrix-matrix product
 - Hermitian matrix, 2-86
 - symmetric matrix, 2-96
 - triangular matrix, 2-107
 - matrix-vector operation
 - product, 2-24, 2-27
 - Hermitian matrix, 2-38
 - band storage, 2-35
 - packed storage, 2-44
 - symmetric matrix, 2-60
 - band storage, 2-51
 - packed storage, 2-54
 - triangular matrix, 2-77
 - band storage, 2-66
 - packed storage, 2-72

rank-1 update, 2-30, 2-31, 2-33
 Hermitian matrix, 2-40
 packed storage, 2-47
 symmetric matrix, 2-62
 packed storage, 2-56

rank-2 update
 Hermitian matrix, 2-42
 packed storage, 2-49
 symmetric matrix, 2-64
 packed storage, 2-58

mixed radix Fourier transforms, 9-1

multi-dimensional Fourier transforms, 9-1

Multiplicative Congruential Generator, 8-4

N

naming conventions, 1-8
 BLAS, 2-2
 LAPACK, 4-2, 5-4
 Sparse BLAS, 2-115
 VML, 7-2

NegBinomial, 8-57

NewStream, 8-9

NewStreamEx, 8-10

O

one-dimensional FFTs, 3-1
 complex sequence, 3-10, 3-11, 3-14, 3-16
 complex-to-complex, 3-3
 complex-to-real, 3-12
 computing a forward FFT, real input data,
 3-8, 3-10
 computing a forward or inverse FFT of a
 complex vector, 3-4, 3-5
 groups, 3-2
 performing an inverse FFT, complex input
 data, 3-13, 3-15
 real-to-complex, 3-7
 storage effects, 3-8, 3-13, 9-42, 9-44

orthogonal matrix, 5-74, 5-101, 5-174, 5-225

P

Packed formats, 9-38

packed storage scheme, A-4

parameters
 for a Givens rotation, 2-14
 modified Givens transformation, 2-17

platforms supported, 1-5

points
 rotation in the modified plane, 2-15
 rotation in the plane, 2-12

Poisson, 8-56

positive-definite matrix
 generalized eigenvalue problems, 5-158

product
 See also dot product
 matrix-vector
 general matrix, 2-27
 band storage, 2-24
 Hermitian matrix, 2-38
 band storage, 2-35
 packed storage, 2-44
 symmetric matrix, 2-60
 band storage, 2-51
 packed storage, 2-54
 triangular matrix, 2-77
 band storage, 2-66
 packed storage, 2-72
 scalar-matrix
 general matrix, 2-83
 Hermitian matrix, 2-86
 scalar-matrix-matrix
 general matrix, 2-83
 Hermitian matrix, 2-86
 symmetric matrix, 2-96
 triangular matrix, 2-107
 vector-scalar, 2-18

pseudorandom numbers, 8-1

Q

QR factorization, 5-6
 computing the elements of

- orthogonal matrix Q, 5-17, 5-50, 5-52
- unitary matrix Q, 5-21
- with pivoting, 5-11, 5-14
- quasi-triangular matrix, 5-174, 5-225

R

- Random Number Generators, 8-1
- random stream, 8-2
- rank-1 update
 - conjugated, general matrix, 2-31
 - general matrix, 2-30
 - Hermitian matrix, 2-40
 - packed storage, 2-47
 - symmetric matrix, 2-62
 - packed storage, 2-56
 - unconjugated, general matrix, 2-33
- rank-2 update
 - Hermitian matrix, 2-42
 - packed storage, 2-49
 - symmetric matrix, 2-64
 - packed storage, 2-58
- rank-2k update
 - Hermitian matrix, 2-92
 - symmetric matrix, 2-103
- rank-n update
 - Hermitian matrix, 2-89
 - symmetric matrix, 2-100
- Rayleigh, 8-36
- real-to-complex one-dimensional FFTs, 3-7, 9-21
- real-to-complex two-dimensional FFTs, 3-21, 9-21
- reducing generalized eigenvalue problems, 5-158
- refining solutions of linear equations
 - band matrix, 4-95
 - general matrix, 4-92, 4-98
 - Hermitian matrix, 4-116
 - packed storage, 4-122
 - Hermitian positive-definite matrix, 4-101, 4-110

- band storage, 4-107
 - packed storage, 4-104
- symmetric matrix, 4-113
 - packed storage, 4-119
- symmetric positive-definite matrix, 4-101, 4-110
 - band storage, 4-107
 - packed storage, 4-104

- RegisterBrng, 8-62
- registering a basic generator, 8-59
- rotation
 - of points in the modified plane, 2-15
 - of points in the plane, 2-12
 - of sparse vectors, 2-123
 - parameters for a Givens rotation, 2-14
 - parameters of modified Givens transformation, 2-17
- routine name conventions
 - BLAS, 2-2
 - Sparse BLAS, 2-115

S

- scalar-matrix product, 2-83, 2-86, 2-96
- scalar-matrix-matrix product, 2-86
 - general matrix, 2-83
 - symmetric matrix, 2-96
 - triangular matrix, 2-107
- scattering compressed sparse vector's elements into full storage form, 2-124
- SetValue, 9-31
- singular value decomposition, 5-74
 - See also* LAPACK routines, singular value decomposition
- SkipAheadStream, 8-16
- smallest absolute value of a vector element, 2-22
- solving linear equations. *See* linear equations
- Sparse BLAS, 2-114
 - data types, 2-115
 - naming conventions, 2-115
- Sparse BLAS routines and functions, 2-115

- ?axpyi, 2-116
 - ?dotci, 2-119
 - ?doti, 2-118
 - ?dotui, 2-120
 - ?gthr, 2-121
 - ?gthrz, 2-122
 - ?roti, 2-123
 - ?sctr, 2-124
 - sparse vectors, 2-114
 - adding and scaling, 2-116
 - complex dot product, conjugated, 2-119
 - complex dot product, unconjugated, 2-120
 - compressed form, 2-114
 - converting to compressed form, 2-121, 2-122
 - converting to full-storage form, 2-124
 - full-storage form, 2-114
 - Givens rotation, 2-123
 - norm, 2-116
 - passed to BLAS level 1 routines, 2-116
 - real dot product, 2-118
 - scaling, 2-116
 - split Cholesky factorization (band matrices), 5-172
 - Status Checking, in DFTI, 9-10
 - stream descriptor, 8-2
 - stream state, 8-7
 - stride. *See* increment
 - sum
 - of magnitudes of the vector elements, 2-5
 - of sparse vector and full-storage vector, 2-116
 - of vectors, 2-6
 - SVD (singular value decomposition), 5-74
 - swapping vectors, 2-20
 - Sylvester's equation, 5-222
 - symmetric matrix, 5-101, 5-157
 - Bunch-Kaufman factorization, 4-22
 - packed storage, 4-28
 - estimating the condition number, 4-78, 5-154
 - packed storage, 4-82
 - generalized eigenvalue problems, 5-157
 - inverting the matrix, 4-139
 - packed storage, 4-143
 - matrix-vector product, 2-60
 - band storage, 2-51
 - packed storage, 2-54
 - rank-1 update, 2-62
 - packed storage, 2-56
 - rank-2 update, 2-64
 - packed storage, 2-58
 - rank-2k update, 2-103
 - rank-n update, 2-100
 - scalar-matrix-matrix product, 2-96
 - solving systems of linear equations, 4-49
 - packed storage, 4-53
 - symmetric positive-definite matrix
 - Cholesky factorization, 4-14
 - band storage, 4-18
 - packed storage, 4-16
 - estimating the condition number, 4-70
 - band storage, 4-74
 - packed storage, 4-72
 - tridiagonal matrix, 4-76
 - inverting the matrix, 4-135
 - packed storage, 4-137
 - solving systems of linear equations, 4-41
 - band storage, 4-45
 - packed storage, 4-43
 - system of linear equations
 - with a triangular matrix, 2-79
 - band storage, 2-69
 - packed storage, 2-75
 - systems of linear equations. *See* linear equations
- T**
-
- transforms, Fourier (advanced), 9-1
 - transposition parameter, A-6
 - triangular factorization
 - band matrix, 4-10
 - general matrix, 4-7
 - Hermitian matrix, 4-25
 - packed storage, 4-31
 - Hermitian positive-definite matrix, 4-14

- band storage, 4-18
- packed storage, 4-16
- tridiagonal matrix, 4-20
- symmetric matrix, 4-22
 - packed storage, 4-28
- symmetric positive-definite matrix, 4-14
 - band storage, 4-18
 - packed storage, 4-16
 - tridiagonal matrix, 4-20
- tridiagonal matrix, 4-12
- triangular matrix, 5-174, 5-225
 - estimating the condition number, 4-86
 - band storage, 4-90
 - packed storage, 4-88
 - inverting the matrix, 4-147
 - packed storage, 4-148
 - matrix-vector product, 2-77
 - band storage, 2-66
 - packed storage, 2-72
 - scalar-matrix-matrix product, 2-107
 - solving systems of linear equations, 2-79, 4-57
 - band storage, 2-69, 4-61
 - packed storage, 2-75, 4-59
- tridiagonal matrix, 5-101
 - estimating the condition number, 4-67
 - solving systems of linear equations, 4-38, 4-47, 6-42
- two-dimensional FFTs, 3-17, 9-21
 - complex-to-complex, 3-18
 - complex-to-real, 3-27
 - computing a forward FFT, real input data, 3-22, 3-24
 - computing a forward or inverse FFT, 3-19, 3-20
 - computing an inverse FFT, complex input data, 3-28, 3-29
 - data storage types, 3-18
 - data structure requirements, 3-18
 - equations, 3-18
 - groups, 3-17
 - real-to-complex, 3-21

U

- Uniform (continuous), 8-21
- Uniform (discrete), 8-44
- UniformBits, 8-46
- unitary matrix, 5-74, 5-101, 5-174, 5-225
- updating
 - rank-1
 - general matrix, 2-30
 - Hermitian matrix, 2-40
 - packed storage, 2-47
 - symmetric matrix, 2-62
 - packed storage, 2-56
 - rank-1, conjugated
 - general matrix, 2-31
 - rank-1, unconjugated
 - general matrix, 2-33
 - rank-2
 - Hermitian matrix, 2-42
 - packed storage, 2-49
 - symmetric matrix, 2-64
 - packed storage, 2-58
 - rank-2k
 - Hermitian matrix, 2-92
 - symmetric matrix, 2-103
 - rank-n
 - Hermitian matrix, 2-89
 - symmetric matrix, 2-100
- upper Hessenberg matrix, 5-174, 5-225

V

- vector arguments, A-1
 - array dimension, A-1
 - default, A-2
 - examples, A-2
 - increment, A-1
 - length, A-1
 - matrix one-dimensional substructures, A-2
 - sparse vector, 2-114
- vector indexing, 7-6

-
- vector mathematical functions, 7-8
 - cosine, 7-23
 - cube root, 7-15
 - denary logarithm, 7-22
 - division, 7-11
 - error function value, 7-40
 - exponential, 7-20
 - four-quadrant arctangent, 7-31
 - hyperbolic cosine, 7-32
 - hyperbolic sine, 7-34
 - hyperbolic tangent, 7-35
 - inverse cosine, 7-28
 - inverse cube root, 7-16
 - inverse hyperbolic cosine, 7-36
 - inverse hyperbolic sine, 7-37
 - inverse hyperbolic tangent, 7-39
 - inverse sine, 7-29
 - inverse square root, 7-13
 - inverse tangent, 7-30
 - inversion, 7-10
 - natural logarithm, 7-21
 - power, 7-17
 - power (constant), 7-18
 - sine, 7-24
 - sine and cosine, 7-25
 - square root, 7-12
 - tangent, 7-27
 - vector pack function, 7-42
 - vector statistics functions
 - Bernoulli, 8-48
 - Binomial, 8-52
 - Cauchy, 8-33
 - CopyStream, 8-12
 - DeleteStream, 8-11
 - Exponential, 8-26
 - Gaussian, 8-23
 - Geometric, 8-50
 - GetBrngProperties, 8-63
 - GetStreamStateBrng, 8-19
 - Gumbel, 8-41
 - Hypergeometric, 8-54
 - Laplace, 8-28
 - LeapfrogStream, 8-13
 - Lognormal, 8-38
 - NegBinomial, 8-57
 - NewStream, 8-9
 - NewStreamEx, 8-10
 - Poisson, 8-56
 - Rayleigh, 8-36
 - RegisterBrng, 8-62
 - SkipAheadStream, 8-16
 - Uniform (continuous), 8-21
 - Uniform (discrete), 8-44
 - UniformBits, 8-46
 - Weibull, 8-31
 - vector unpack function, 7-44
 - vectors
 - adding magnitudes of vector elements, 2-5
 - copying, 2-7
 - dot product
 - complex vectors, 2-10
 - complex vectors, conjugated, 2-9
 - real vectors, 2-8
 - element with the largest absolute value, 2-21
 - element with the smallest absolute value, 2-22
 - Euclidean norm, 2-11
 - Givens rotation, 2-14
 - linear combination of vectors, 2-6
 - modified Givens transformation parameters, 2-17
 - rotation of points, 2-12
 - rotation of points in the modified plane, 2-15
 - sparse vectors, 2-115
 - sum of vectors, 2-6
 - swapping, 2-20
 - vector-scalar product, 2-18
 - vector-scalar product, 2-18
 - sparse vectors, 2-116
 - VML, 7-1
 - VML functions
 - mathematical functions
 - Acosh, 7-36
 - Asin, 7-29
 - Asinh, 7-37
 - Atan, 7-30

- Atan2, 7-31
- Atanh, 7-39
- Cbrt, 7-15
- Cos, 7-23
- Cosh, 7-32
- Div, 7-11
- Erf, 7-40
- Exp, 7-20
- Inv, 7-10
- InvCbrt, 7-16
- InvSqrt, 7-13
- Ln, 7-21
- Log10, 7-22
- Pow, 7-17
- Powx, 7-18
- Sin, 7-24
- SinCos, 7-25
- Sinh, 7-34
- Sqrt, 7-12
- Tan, 7-27
- Tanh, 7-35
- pack/unpack functions
 - Pack, 7-42
 - Unpack, 7-44
- service functions
 - ClearErrorCallBack, 7-58
 - ClearErrStatus, 7-54
 - GetErrorCallBack, 7-57
 - GetErrStatus, 7-53
 - GetMode, 7-50
 - SetErrorCallBack, 7-55
 - SetErrStatus, 7-51
 - SetMode, 7-47
- VSL functions
 - advanced service subroutines
 - GetBrngProperties, 8-63
 - RegisterBrng, 8-62
 - generator subroutines
 - , 8-57
 - Bernoulli, 8-48
 - Binomial, 8-52
 - Cauchy, 8-33
 - Exponential, 8-26
 - Gaussian, 8-23
 - Geometric, 8-50
 - Gumbel, 8-41
 - Hypergeometric, 8-54
 - Laplace, 8-28
 - Lognormal, 8-38
 - Poisson, 8-56
 - Rayleigh, 8-36
 - Uniform (continuous), 8-21
 - Uniform (discrete), 8-44
 - UniformBits, 8-46
 - Weibull, 8-31
 - service subroutines
 - CopyStream, 8-12
 - DeleteStream, 8-11
 - GetStreamStateBrng, 8-19
 - LeapfrogStream, 8-13
 - NewStream, 8-9
 - NewStreamEx, 8-10
 - SkipAheadStream, 8-16

W

Weibull, 8-31

X

XERBLA, error reporting routine, 2-1