

CSAR Service

Consolidated Management Report

4th Quarter 2003

Management Summary

This is the consolidated Management report for the fourth quarter 2003 of the CSAR HPC facility for UK Academia and Industry, which enables World-Class research and development.

The number of users has grown to a total of 470 to date.

The workload on both the Origin 3000 (Green) and the Cray T3E (Turing) has been fairly evenly spread across the mid- to high-end ranges of PEs during the fourth quarter of this year, with the T3E running at approximately three-quarters capacity.

CSAR has been granted an 18 month extension of service contract until June 30th 2006. With this extension CfS has implemented a further technology refresh introducing a 256 processor Itanium-2 (Madison) based SGI Altix which went into production service on 1st October.

The Cray T3E system Turing was removed at the end of December, having come to the end of its contractual lifespan.

CfS remains active in the UK Grid Forum.

Introduction

This Management Report includes a section for each of the main service functions:

1. Service Quality
2. HPC Services
3. Science Applications Support Services
4. Training & Education Services
5. User Registration & New User Services
6. Value-Added Services

Each section includes a status report for the period, including notable achievements and problems, also noteworthy items for the next period.

1 Service Quality

This section covers overall Customer Performance Assessment Ratings (CPARS), HPC System availability and usage, Service Quality Tokens and other information concerning issues, progress and plans for the CSAR Service.

1.1 CPARS

Table 1 gives the measure by which the quality of the CSAR Service is judged. It identifies the metrics and performance targets, with colour coding so that different levels of achievement against targets can be readily identified. Unsatisfactory actual performance will trigger corrective action.

CSAR Service - Service Quality Report - Performance Targets

Service Quality Measure	Performance Targets					
	White	Blue	Green	Yellow	Orange	Red
HPC Services Availability						
Availability in Core Time (% of time)	> 99.9%	> 99.5%	> 99.2%	> 98.5%	> 95%	95% or less
Availability out of Core Time (% of time)	> 99.8%	> 99.5%	> 99.2%	> 98.5%	> 95%	95% or less
Number of Failures in month	0	1	2 to 3	4	5	> 5
Mean Time between failures in 52 week rolling period (hours)	>750	>500	>300	>200	>150	otherwise
Fujitsu Service Availability						
Availability in Core Time (% of time)	> 99.9%	> 99.5%	> 99.2%	> 98.5%	> 95%	95% or less
Availability out of Core Time (% of time)	> 99.8%	> 99.5%	> 99.2%	> 98.5%	> 95%	95% or less
Help Desk						
Non In-depth Queries - Max Time to resolve 50% of all queries	< 1/4	< 1/2	< 1	< 2	< 4	4 or more
Non In-depth Queries - Max Time to resolve 95% of all queries	< 1/2	< 1	< 2	< 3	< 5	5 or more
Administrative Queries - Max Time to resolve 95% of all queries	< 1/2	< 1	< 2	< 3	< 5	5 or more
Help Desk Telephone - % of calls answered within 2 minutes	>98%	> 95%	> 90%	> 85%	> 80%	80% or less
Others						
Normal Media Exchange Requests - average response time	< 1/2	< 1	< 2	< 3	< 5	5 or more
New User Registration Time (working days)	< 1/2	< 1	< 2	< 3	< 4	otherwise
Management Report Delivery Times (working days)	< 1	< 5	< 10	< 12	< 15	otherwise
System Maintenance - no. of sessions taken per system in the month	0	1	2	3	4	otherwise

Table 1

Table 2 gives actual performance information for the period. Overall, the CPARS Performance Achievement for the 4th quarter 2003 was satisfactory (see Table 3), i.e. Green measured against the CPARS performance targets.

CSAR Service - Service Quality Report - Actual Performance Achievement

Service Quality Measure	2003											
	Jan	Feb	March	April	May	June	July	Aug	Sept	Oct	Nov	Dec
HPC Services Availability												
Availability in Core Time (% of time)	99.46%	99.73%	100%	99.74%	97.66%	99.25%	98.83%	98.95%	96.62%	98.84%	98.95%	98.75%
Availability out of Core Time (% of time)	99.89%	100.00%	99.81%	99.81%	99.33%	99.9%	99.57%	100%	98.48%	99.28%	97.74%	98.3%
Number of Failures in month	3	1	1	1	4	1	2	2	4	4	3	5
Mean Time between failures in 52 week rolling period (hours)	487	487	515	548	461	548	487	461	417	365	337	283
Help Desk												
Non In-depth Queries - Max Time to resolve 50% of all queries	<0.25	<0.25	<0.25	<0.25	<0.25	<0.25	<0.25	<0.25	<0.25	<0.25	<0.25	<0.25
Non In-depth Queries - Max Time to resolve 95% of all queries	<0.5	<1	<2	<3	<1	<2	<1	<0.5	<5	<2	<1	<1
Administrative Queries - Max Time to resolve 95% of all queries	<1	<0.5	<1	<0.5	<0.5	<0.5	<0.5	<1	<1	<1	<1	<1
Help Desk Telephone - % of calls answered within 2 minutes	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
Others												
Normal Media Exchange Requests - average response time	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5
New User Registration Time (working days)	0	0	0	0	0	0	0	0	0	0	0	0
Management Report Delivery Times (working days)	10	10	10	10	10	10	10	10	10	10	10	10
System Maintenance - no. of sessions taken per system in the month	2	2	2	2	2	2	2	2	2	2	2	2

Table 2

Notes:

- HPC Services Availability has been calculated using the following formulae, based on the relative NPB performance of Turing, Fermat and Green at installation:

$$\text{Turing availability} \times 143 / (143 + 40 + 233) + [\text{Fermat availability} \times 40 / (143 + 40 + 233) + \text{Green availability} \times 233 / (143 + 40 + 233)]$$
- Mean Time Between Failures for Service Credits is formally calculated from Go-Live Date.

Table 3 gives Service Credit values for each month to date. These are accounted on a quarterly basis, formally from the Go-Live Date. The values are calculated according to agreed Service Credit Ratings and Weightings.


CSAR Service - Service Quality Report - Service Credits

Service Quality Measure	2003											
	Jan	Feb	March	April	May	June	July	Aug	Sept	Oct	Nov	Dec
HPC Services Availability												
Availability in Core Time (% of time)	0	-0.039	-0.058	-0.039	0.078	0	0.039	0.039	0.078	0.039	0.039	0.039
Availability out of Core Time (% of time)	-0.047	-0.047	-0.047	-0.047	0	-0.047	-0.039	-0.047	0.078	-0.039	0.078	0.078
Number of Failures in month	0	-0.008	-0.008	-0.008	0.008	-0.008	0	0	0.008	0.008	0	0.0004
Mean Time between failures in 52 week rolling period (hours)	0	0	-0.008	-0.008	0	-0.008	0	0	0	0	0	0.0002
Help Desk												
Non In-depth Queries - Max Time to resolve 50% of all queries	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019
Non In-depth Queries - Max Time to resolve 95% of all queries	-0.019	-0.016	0	0.016	-0.016	0	-0.016	-0.019	0.0312	0	-0.016	-0.016
Administrative Queries - Max Time to resolve 95% of all queries	-0.016	-0.019	-0.016	0	-0.019	-0.019	-0.019	-0.016	-0.01551	-0.01551	-0.016	-0.016
Help Desk Telephone - % of calls answered within 2 minutes	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004
Others												
Normal Media Exchange Requests - average response time	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002
New User Registration Time (working days)	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019
Management Report Delivery Times (working days)	0	0	0	0	0	0	0	0	0	0	0	0
System Maintenance - no. of sessions taken per system in the month	0	0	0	0	0	0	0	0	0	0	0	0
Monthly Total & overall Service Quality Rating for each period:	-0.06	-0.09	-0.09	-0.07	0.00	-0.06	-0.04	-0.04	0.07	-0.03	0.02	0.02
Quarterly Service Credits:	-0.24			-0.13			-0.01			0.02		
Annual Service credit	-0.36											

Table 3


1.2 No. of Registered Users

The current position at the end of the quarter is that there are 470 registered users of the CSAR Service. The number of users saw a drop during October and November due to a number of projects which came to an end.



1.3 Service Quality Tokens

The graph below illustrates the monthly usage trend of Service Quality Tokens:



Over the course of the quarter the position is that as a management tool the Service Quality Tokens have been available to enable the users to provide qualitative feedback about all aspects of the service. This feedback is used as a mechanism to initiate change in the service where appropriate.

A the end of the quarter there was one black mark allocated to the service, as shown in the table below. A development queue has subsequently been made available for use on Newton.

SUMMARY OF SERVICE QUALITY TOKEN USAGE

No of Stars or Marks	Consortia	Date Allocated	Reason Given
1 black mark	csn003	19/12/03	No development queue on Newton

2 HPC Services Usage

Usage information is given in tabular form, and in graphical format. The system usage information covers:


- CPU usage
- User Disk allocation
- HSM/tape usage

This is illustrated in a number of graphs including:


- a) MPP (T3E) Usage by month, showing usage each month of CPU (T3E PE Elapsed Hours), split by Research Council and giving the equivalent GFLOP-Years as per NPB. The Baseline Capacity is shown by an overlaid horizontal line.
- b) SMP (Origin) Usage by month, showing usage each month in CPU Hours, split by Research Council and giving the equivalent GFLOP-Years as per NPB. The Baseline Capacity is shown by an overlaid horizontal line.
- c) High Performance Disk (T3E) allocated for User Data by month, showing the allocated space each month in GBytes, split by Research Council. The Baseline Capacity (1 Terabyte) is shown by an overlaid horizontal line.
- d) Medium Performance Disk (Origin) allocated for User Data by month, showing the allocated space each month in GBytes, split by Research Council. The Baseline Capacity (1.5 Terabytes) is shown by an overlaid horizontal line.
- e) HSM/Tape Usage (T3E) by month, showing the volumes held each in GBytes, split by Research Council. The Baseline Capacity (16 Terabytes) available will be shown by an overlaid horizontal line.


2.1 Service Usage Charts

The graphs below show recent monthly PE, CPU, disk and HSM allocations and usage.




Turing PE usage is shown by Research Council during the last 12 months of service in the above chart.







Usage of the two batch SGI Origin systems, Fermat and Green, is shown by Research Council during the last 12 months of service in the preceding two charts.




The above chart, and the one below, show the wait time trend in hours on the Turing system.




The following two charts show average wait times in hours for the quarter on the Origin 3000 (Green).







The next series of four charts show the percentage PE time utilisation by the major job categories on the Turing, Green, Fermat and Newton systems for the 4th quarter 2003.




The workload on Turing for the fourth quarter was predominantly in the 33 to 64 PE range.




On Green, the 65 to 128 PE range has seen the greatest percentage of workload during this quarter.




On Fermat the highest concentration of work was in the 17 to 32 PE range.




Newton's workload during its first quarter in production has been spread fairly evenly across the system.



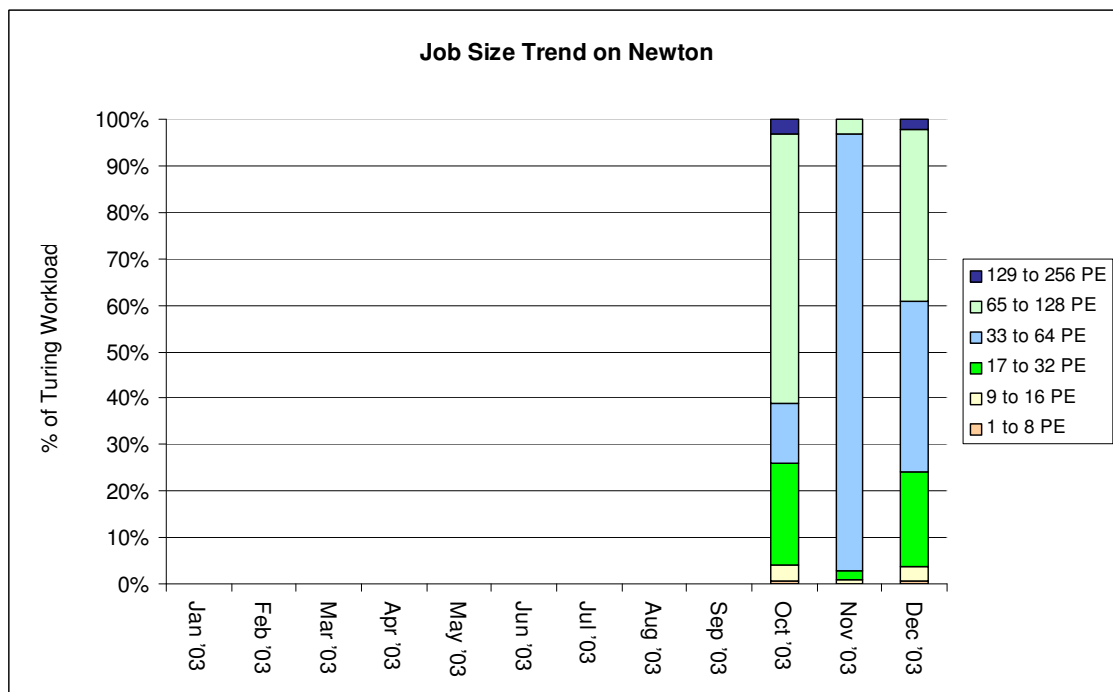
During this quarter there has been increasing usage of the T3E for jobs in the mid-range of PEs.



Usage on Green tended more to the higher-end range of PEs during the 4th quarter.



The trend on Fermat during this quarter is that the greatest proportion of the workload was primarily in the mid-range of PEs.




During its first quarter in production, Newton's usage was reasonably spread across the machine.

2.2 System Usage Graphs


In all the Usage Charts, the baseline varies dependant on the number of days in each month, within a 365-day year.

2.2.1 Baseline System


The graph below shows the Gflop Years utilisation on the CSAR systems by Research Council for the last 12 months. BBSRC has become a significant user of Green.




The graph below shows the same service utilisation by major system. This quarter has seen the introduction of the new SGI Altix 3700 system Newton into the CSAR service.




The next chart shows the historic interactive usage of the 'baseline' Fermat system (equivalent to 16@250Mhz CPUs) up to the end of February 2003, at which point the interactive usage was transferred to Wren and Fermat became a batch-only system.. Eight of the higher speed 500Mhz CPUs in the Origin 300 system (Wren) deliver the baseline capacity equivalent to that which was previously available on Fermat for interactive usage.



The next series of graphs illustrates the usage of the disk and HSM resources of the system.




The graph above illustrates the historic allocation of the High Performance Disk on Turing.



This graph illustrates the historic allocation of the Medium Performance Disk on Fermat.

The next graph shows the historic HSM usage by Research Council funded projects, which has exceeded the overall Baseline of 16 Terabytes, and now totals 48 Terabytes.



2.2.2 Guest System Usage

There is currently no Guest System usage.

2.3 Service Status, Issues and Plans

Status

The service has been reasonably utilised throughout the fourth quarter of 2003, with usage exceeding baseline.

During the quarter there was a relatively balanced spread of work across all major systems.

The end of the fourth quarter has seen the removal of the Cray T3E system Turing due to its reaching the end of its contract lifespan. The SGI Altix 3700 Itanium-2 system Newton, which was introduced into the service at the beginning of October, has seen a steady increase of work across the quarter, with several codes having been ported from Turing with encouraging results.

Issues

There have been some outages during the latter part of the fourth quarter due to issues with the SAN. These issues were subsequently identified by the vendor, and patches to address the problems encountered have now been applied to the Operating Systems of all affected machines. All systems continue to be closely monitored to ensure maximum stability of the CSAR service.

Plans

Work continues with porting codes to the Altix system Newton, with promising results being seen in scaling and speed when compared to other systems.

It is planned to make the /hold facility accessible via the Storage Area Network (SAN) towards the end of January. This will mean that /hold is then available to users of the Altix system Newton, alongside users of the Origin systems who currently have direct access to /hold.

An upgrade to the next release 2.4 of the Altix Linux (Propack) operating system is planned for the first quarter 2004, offering benefits including greater stability for the Altix system Newton.

3 Project Management, Documentation and User Feedback

This section covers aspects relating to the registration of projects and users, the management of projects and resources, topics associated with documentation and user feedback.

3.1 Project Applications

5 applications for new CSAR projects were received in the fourth quarter of 2003, requesting a total of 71928 service tokens.

3.2 New Projects

6 new CSAR projects were started with 25699 service tokens being awarded in total.

3.3 Queries

Overall, 313 Class 1, 2 and 3 queries relating to the CSAR service were received by the Helpdesk between 1st October and 30th December 2003.

3.4 Service Quality Tokens

Five black marks were awarded during the period. This was a complaint about the lack of a development queue on Newton despite requests for one. This was reduced to one black mark after agreement was reached on how to best implement the request.

3.5 Annual Report

The CfS Annual Report is currently being prepared for publication.

3.6 CSAR Focus

The Winter edition is being prepared.

3.7 CSAR Website

Fiona Cook continues to manage the creation of the new CSAR website.

3.8 New Altix System

The user service started on 1st October as planned. Initially jobs were restricted to a single partition, but cross-partition jobs were possible from around mid-October. The maximum job size is currently 244 processors. Approximately 150 users were registered to use the Altix by 31st December.

3.9 Turing Retirement

The Cray T3E system Turing was retired on 31st December 2003. The machine had been in service since the beginning of the CSAR service.

4 Scientific Application Support Services

4.1 Training and Education

The new course timetable for CSAR/HPC courses for Semester 1 of the 2003/2004 academic year is now available at <http://www.csar.cfs.ac.uk/using/courses>.

Several new courses have been developed including courses on the new Altix system and the Unified Model.

4.2 Consortia Support/Software

Work has concentrated on porting and optimising codes both for the new Altix service and for the removal of Turing. There continues to be great activity in this area. Installation of software on the Altix continues.

Some highlights of the porting and timing process (ongoing):

- o Newton is ~3 times faster than Green for most user codes, with some (eg Gaussian) up to 5 times faster
- o The Unified Model is 25% faster than on a p690 system
- o The molecular dynamics codes NAMD and AMBER are both ~50% faster than on a p690 system

There are a number of consortia that we have provided optimization work for:

CgLES

CgLES is a C/C++ code which required porting to Newton but there was also a request for optimization work to be done for Green. After analysis of the code it was found that the compiler did a poor job of dereferencing arrays. This needed to be done manually on 4 of the key functions in the code. This resulted in an increase in performance of between 200% and 300%. This left a code that was heavily dominated by communication; in fact on 14 processors, 50% of runtime was communication. Further work was done to convert communications to single-sided MPI communication which reduced the MPI to 20% of the runtime of the code. This work was performed by Neil Stringfellow.

Kai Luo

This began as code porting to Green (bugs in the Intel compiler are preventing this being ported to Newton). Minor serial code optimizations provided no significant gains. Communication also proved to be a problem. The communication involved numerous sends from one processor to its neighbour in order to perform halo exchanges followed by receives. This was replaced by sendrecv using a derived type; this also opens up the possibility of single copy communication (a significant parallel optimization), and will also allow further optimization to convert the code to single-sided communication. On 128 processors, the code runs about 12 times faster on Newton than on Turing. This work was performed by Neil Stringfellow.

H2Mol

H2Mol is an in-house code from Queen's University Belfast. This code is part of the porting effort to Newton. Initial runs of the code showed extremely poor timings, taking longer than Green. The code is dominated by multiplications of the following form:

$$C = A \times B + B \times A$$

The major reason for the poor performance was that this was implemented as a pair of matmul operations and A is not stride one data. Matmul appears to be inlined into the code on the altix which led to really poor cache use and also exhibited a feature unique to the Intel chip, cache bank conflicts. Work was done to replace this function in a convenient manner to further development or optimization exercises, and on this resulted in the code taking half as long to run. Further optimization was performed after discovering that certain matrices were symmetric, and replacing loop structures with library calls to SCSL routines. Currently, scaling does deteriorate with higher numbers of processors, but is consistent between the two systems, with Newton more than 3 times faster than Green. Work performed by Kevin Roy.

Helium

Helium is another in-house code from Queen's University Belfast. It has also been part of the porting program to the Altix. Work has been performed to increase the performance by nearly 20% on Newton and some minor work on MPI buffering has been completed. The code is currently undergoing a more rigorous examination of its parallel performance. Currently, Newton is more than 6 times faster than Turing. Work performed by Kevin Roy.

Cse082 – EROS-UK

This is work that was begun in the previous quarter, and also reported upon in the previous report on optimization work. Further analysis of this code revealed a dependence between the problem and the number of processors that it ran on optimally. On the problems that we were given this didn't give any concerns of scalability. Serial optimization proved extremely difficult due to the complex datastructures employed. The compiler found it too difficult to provide any optimization beyond -O2. The limited support resources of this project meant that serial optimization could not be performed. Work performed by Kevin Roy.

ftLMPs

FtLMPs is a laser materials processing code which uses a 3D adaptive mesh refinement scheme. Initial work on this code has been done identifying problem areas and will be continued into the next quarter. Work performed by Kevin Roy.

CASTEP

Initial timings show that Castep is around 50% faster than HPCx on 32 processors, and scales better than on HPCx. Work is ongoing to correct some communications problems around lower processor counts – especially at 2 processors on Newton.

CPMD

Newton is about 3.5 times faster than Green.

The Channel code (figures provided by Dr Lionel Temmerman)

Initial timings show roughly a factor of 10 improvement on Newton over Turing, without any optimisation.

PChan (Dr Zhiwei Hu)

CSAR staff have ported this code from Turing to Green and Newton. After optimisation, it typically performs about 7 times faster on Newton than on Turing.

PDSN3D (Dr Zhiwei Hu)

CSAR staff have ported this code to Newton. It runs about 5 times faster than on Green up to 64 processors.

LB3D (Prof Peter Coveney)

LB3D has been ported and optimised on Newton. The code was used extensively and very successfully in the Teragyroid experiments during SC2003. There were compiler problems on Newton which prevented use of the latest compiler. Two types of code changes were made to significantly improve performance. One of these code changes, made by Kevin Roy, also solved problems in running the code on HPCx. The compiler issues on Newton are still being pursued, although they are not delaying any work on the system.

NAMD ApoA1 benchmark

Time is per timestep in seconds. The best results for single processor jobs on other machines including P690 (from the NAMD website) are around 8-9 seconds per step as opposed to 6.2 seconds per step for Newton. See <http://www.ks.uiuc.edu/Research/namd/performance.html> or the HPCx website <http://www.hpcx.ac.uk/research/hpc/NAMD.html> - note that these web pages use a slightly older version of NAMD.

NCAS Consortium

See below for the Unified Model.

Porting and optimisation work was performed on behalf of Dr Omduth Coceal, who provided the following feedback:

Re. feedback about Newton, it is almost an order of magnitude faster than Turing – so what used to take me days can now be done in hours. This is quite a major boost in terms of my research. Memory also used to be a problem on Turing – I experienced problems with postprocessing large amounts of data. This is no longer an issue on Newton.

Unified Model

Jeff Cole from the NCAS consortium has provided some information on running the UM 4.5 on various machines. For the atmosphere model, on 4 and 8 processors, both Newton and HPCs run about 3 times faster than Green. On 16 processors there is a marginal improvement on Newton, while on HPCx performance is not as good as on 8 processors.

In contrast, for the Ocean model, performance on Newton is similar to that on Green, while HPCx is significantly better. However, similar performance is obtained on Newton and HPCx by using the latest version of the compiler on Newton. This is still under investigation.

NWChem (Dr Tanja Van Mourik)

The CPU time on Newton is roughly 10% faster than on HPCx. However, the wall time on Newton is excessive. The latest version of NWChem has been installed, and the wall clock time problem on Newton is being investigated.

VASP (Results provided by Dr Dario Alfe, UCL)

Results obtained using VASP 4.4.3, using the Intel 7.1 compiler. The latest version, VASP, 4.6.7, is now being installed, and the Intel 8.0 compiler will be used for both versions. Better performance is expected with the 8.0 compiler. VASP scales well on both Newton and Green with excellent performance in comparison with HPCx.

4.3 Netsolve

Work continues on the project to get Netsolve installed and working on CSAR systems. This project is being done in collaboration with Jack Dongarra's group at the University of Tennessee.

5 Collaboration and Conferences

5.1 MRCCS Projects

5.2.1 Virtual Prototyping/Finite Elements

Lee Margetts has finished writing the prototype of the new parallel FEA code. Work is ongoing.

5.2 Events

The following events were attended by SVE staff:

- o November – Supercomputing 2003. See item 6.1.
- o November – Worldwide Universities Network conference in America.
- o December – Daresbury workshop evaluation. See item 6.4.

6 Added Value Services

6.1 International Conferences – SC2003 Phoenix

The University of Manchester were involved in many activities:

- o HPC Challenge competition winners in the 'Most Innovative Data-Intensive Application' category with TRICEPS which involved resources from CSAR, HPCx, Pittsburgh, San Diego and Argonne National Labs in Chicago.
- o Partners in the second HPC Challenge category 'Most Geographically Distributed Application' led by HLRS Stuttgart and Craig Stewart of Indiana.
- o An exhibition stand on Supercomputing and visualisation run by Pen Richardson.
- o Adrian Tate assisted with a tutorial on Co-array Fortran.
- o SGI BOF on visualisation.
- o Mike Daw was the Technical Director for the SC Global event on collaborative working.
- o Lee Margetts led a showcase demonstration on virtual prototyping at SC Global.
- o Terry Hewitt was the Chair of two SC Global events.
- o Presentations on SGI and UK e-Science stands about CSAR, Reality Grid and Triceps.
- o Further presentations on ANL, NCSA, PSC and CALTECH stands.
- o Major involvement in Teragyroid project.

6.2 Seminars

6.2.1 MRCCS/ESNW Seminar Series

The following seminars have been held during the last three months:

- o Delivering e-Science and Visualization to Medical Applications – Nigel John, Joanna Leng, Mark Riding, Manchester Visualization Centre, University of Manchester.
- o ESNW Middleware Discussion.
- o Pedro: A Tool for Bridging e-Science with Grid Services. A Case Study for Medical Informatics – Chris Garwood, University of Manchester.
- o Applications of Ontologies to Grid Middleware Interoperability – Kevin Garwood, University of Manchester.
- o A Grid Application Framework based on Web Services Specifications and Practices – Dr. Savas Parastatidis, School of Computing Science, University of Newcastle upon Tyne.

6.3 Summer School 2003

The Summer School 2003 on HPC in finite Element Analysis, jointly organised by MRCCS and the National Science Foundation (NSF) of the USA took place from 1st to 5th September. Speakers attended from France and Japan in addition to the speakers from the USA and the UK. Questionnaires showed that delegates were very pleased with this meeting. We will be in contact with participants to see how these ideas can be taken forward in the future. We are also hoping to obtain funding from EPSRC for further work.

6.4 Visualization

.SVE again supported the Machine Evaluation Workshop held annually at Daresbury. Andrew Jones gave a talk on benchmarking – why, how, pitfalls to avoid, what benchmarking means, the dangers of being too simple, etc. The talk included a limited benchmark study of Manchester's highly scalable finite element engineering code, showing that Itanium2 (and in particular the SGI Altix) was probably the best performer, faring well against systems such as the IBM p690+ (HPCx Phase 2), Cray XI, NEC SX5, and others.