


CSAR FOCUS

Edition 4, January 2000, Editor: Kaukab Jaffri, University of Manchester

CSAR

Being the first issue of the new millennium, and CSAR's first birthday, this edition of CSAR Focus will certainly be one to remember! We celebrate with more than just a cake - we cut into the Earth's surface.....



Turing Probes the Earth's Core

Ground-breaking work to determine the temperature of the Earth's core carried out by a team of scientists at University College London, using the CSAR Cray T3E

CSC



sgi

Content

EDITORIAL

ARTICLE: Turing Probes the Earth's Core	4
FEATURE: CSAR Fujitsu VPP300	8
CSAR Wins Award at SuperComputing '99	10
CSAR Applications Support News	16
MRCCS News: Report on the Summer School	17
MRCCS Seminars	20
MRCCS Summer School Lecture: Professor Tom Kilburn & the Making of "The Baby"	21
PEOPLE FOCUS	24
Cray T3E Workshop	26
SOFTWARE UPDATE: FTW - The Fastest Fourier Transform in the West!	27
The Effects of Cluster Management on Parallel Applications	30
A First Meeting-Ground with Laboratory Experiment Enabled by Turing	31
HPC/CSAR Course Calendar	35
HPC/CSAR Course Description	37
Forthcoming Events	41

「 Editorial 」

As this is the first issue of the new millennium, on behalf of CSAR, I hope you all had a good holiday.

This edition of CSAR Focus is packed with a spectrum of interesting articles to inspire you all to make a new year's resolution to contribute to the next issue. You never know, your research may even make it to the front page!

Dr Dario Alfe and colleagues at UCL have managed to do just that. They have produced ground-breaking results using CSAR's Cray T3E to find out how hot the earth's core really is. This work reached as far as the BBC and Channel 4 News! The article features on page 4.

Also in this issue, we have an article by Professor Ken Taylor from Queen's University Belfast on his groups' recent breakthrough research on multiphoton processes carried out on the Cray T3E.

Our **Feature** is on CSAR's vector processor, the Fuji machine, (Fujitsu VPP300), and Manjit Boparai from CSC gives us an update of the capabilities of the machine and encourages CSAR users to make use of its resources.

Dr John Brooke, Head of CSAR Applications Support, writes about the award won by Manchester Computing at the SC99 Conference in the USA.

If you would like to contribute to the next issue of CSAR Focus which is June 2000, please contact me at the CSAR Helpdesk. The deadline for submission of articles will be announced on the CSAR News pages: <http://www.csar.cfs.ac.uk/general/news.shtml> in the coming months.

「 We welcome you all to participate in this publication, with any articles of interest to other CSAR users, or comments on any aspect of HPC. 」

Copyright

Pictures are reproduced by permission. Any further reproduction will require the owners permission.

Disclaimer

Articles are the responsibility of the individual author: the Editor and the CSAR Service assume no responsibility or liability for articles or information appearing in this document. Opinions expressed in some articles may not reflect those of the CSAR Service.

Kaukab Jaffri
Editor, CSAR Focus

Manchester Computing
University of Manchester
Oxford Road
Manchester
M13 9PL

Tel: (+44) 161 275 6824
Fax: (+44) 161 275 6040
Email: k.jaffri@mcc.ac.uk
Web: www.csar.cfs.ac.uk

Turing Probes the Earth's Core

D.Alfe, M.J.Gillan, G.D. Price
Research School of
Geological and Geophysical
Sciences, Birkbeck College
and University College
London

We have exploited the power of the CSAR service to attack one of the great unsolved questions in the earth sciences: how hot is the Earth's core?

A report on our work just published in the journal *Nature*¹ was picked up by the media, and we were interviewed on Channel 4 News and the Today programme. There was even a half-page item on our work in the Daily Mirror, and on the BBC online web page².

The core of the Earth accounts for approximately 30% of the total mass of our planet (schematically in Figure 1.), the inner solid core is crystallizing from the liquid outer core, and the heat released flows to the surface, driving all the living geological processes of the Earth, including plate tectonics, volcanism and earthquakes.

Moreover, this heat flux is also responsible for the convective motions in the liquid core, which generate the Earth's magnetic field, shielding us from the lethal solar wind. So, a sound knowledge of the core is of fundamental importance, yet, it is one of the most difficult things to study, and its properties are poorly constrained. For example, we know that it is mainly made by iron, but it can't be pure iron, because its density is too low. So it must contain some light element, and the most likely candidates are sulphur, oxygen, carbon and silicon, but the real composition of the Earth's core remains one of the major unsolved problems. The density and the pressure are known quite accurately (within a percent), but the temperature is unknown, with estimates ranging from approx. 4000 to approx. 8000 K.




Figure 1: Schematic representation of the Earth interior

Since the main constituent of the core is iron, it is relevant to study its properties under core pressures. The pressure p at the inner-core boundary (ICB) is 330 GPa, and the melting temperature T_m of Fe is needed up to at least this value. Static compression measurements of T_m with the diamond anvil cell (DAC) have been done up to *ca.* 200 GPa, but even at lower pressures results for T_m disagree by several hundred K. Very recent DAC diffraction experiments demonstrate that in the range 60 - 80 GPa the solid phase from which Fe melts has the hexagonal close-packed (ϵ) structure. Shock experiments at present give the only way of going to higher pressures, but their interpretation is not simple, and there is a scatter of at least 2000 K in T_m at ICB pressures. There have also been attempts to obtain the melting curve from parameterized atomistic models for the energetics of iron, but the reliability of these models is uncertain.

Parameter-free *ab initio* techniques based on density functional theory (DFT) have recently made enormous strides³, and for many materials, including transition metals, reliance on parameterised models is now unnecessary. With *ab initio* molecular-dynamics methods (Car-Parrinello)⁴, this is true for both liquids and solids³. Solid iron in all its known crystal structures has been extensively studied with DFT. The accuracy of DFT depends on the approximation used for the electronic exchange-correlation energy, but with the 'generalised-gradient approximation' (GGA) the experimental properties of Fe are very accurately reproduced⁵.

The calculations we present here have been done with the VASP code⁶, which is extremely robust and stable for simulation of metals, we have also doubled the speed of the code for the molecular dynamics simulations using an efficient charge extrapolation⁷. Combining this with the resources of the CSAR service we have been able to do calculations of unprecedented size and length, which have been essential for the success of this work.

The strategy we follow to calculate the melting curve is very simple: we calculate the Gibbs free energy of solid and liquid iron as a function of pressure and temperature. The melting temperature T_m at any pressure is found by equalizing the free energies. In fact, we calculate the Helmholtz free energy $F(V,T)$ as a function of volume V , and from this we obtain the pressure through the relation $p = -(\delta F/\delta V)_T$ and G through its definition $G = F + pV$.

The free energy of the harmonic vibrating solid is conceptually simplest, and is given (per primitive cell) by the standard high-temperature formula:

$$F = F_{\text{perf}} + N_{\mathbf{k}}^{-1} \sum_{\mathbf{k}s} \ln(\hbar\omega_{\mathbf{k}s}/k_B T)$$

where F_{perf} is the free energy of the rigid perfect lattice and $\omega_{\mathbf{k}s}$ is the frequency of the phonon of branch s at wavevector \mathbf{k} . The sum goes over the $N_{\mathbf{k}}$ wavevectors in the Brillouin zone of the crystal and over the branches s , the frequencies $\omega_{\mathbf{k}s}$ were obtained by diagonalising the force-constant matrix.

To calculate the liquid free energy and the anharmonic contribution to the solid free energy, we use the technique of 'thermodynamic integration', which yields the difference $\Delta F = F - F_{\text{ref}}$ between the free energy of the *ab initio* system and that of a reference system, F_{ref} . It is a standard result that this difference is:

$$\Delta F = \int_0^1 d\lambda \langle U - U_{\text{ref}} \rangle_\lambda,$$

where the thermal average $\langle \cdot \rangle$ is evaluated for the system governed by $U_\lambda = (1-\lambda)U_{\text{ref}} + \lambda U$ where U_{ref} and U are the internal energies of the reference system and the *ab initio* system respectively.


Even with the available CSAR power, the practical feasibility of calculating *ab initio* free energies of liquids and anharmonic solids hinges on finding a reference system for which F_{ref} is readily calculable and the difference $U - U_{\text{ref}}$ is very small.

For the liquid, the crucial discovery that makes our threshold precision achievable is that the *ab initio* energy U is extremely well reproduced by a model U_{ref} consisting of a sum of pair potentials $\phi(r)$, with $\phi(r)$ chosen to be an inverse-power repulsive potential: $\phi(r) = A/r^\alpha$. The parameters A and α are adjusted to minimise the strength of the fluctuations. We have demonstrated the excellence of this U_{ref} by doing *ab initio* dynamical simulations at 18 different thermodynamic states spanning the (p,T) range of interest, and we find that a single choice of A and α is almost equally good over the whole range.

Figure 2: The *ab initio* melting curve of iron compared with experimental results. Chain curve shows *ab initio* results; solid and dotted curves are interpolations of DAC measurements; points are more recent measurements, and represent a lower bound rather than the melting curve itself; open squares and open diamond with error bar are shock data.

The inverse-power model is not a good enough reference system for the anharmonic solid, but we find that an outstandingly good reference system for this case can be obtained by suitably combining the inverse-power and harmonic *ab initio* total energies. We stress that the choice of reference system affects only the efficiency of the calculations, and not the final results.

Figure 2 compares our *ab initio* melting curve with experimental results from both DAC and shock measurements. For pressures $p < 200$ GPa covered by DAC, our curve lies *ca.* 1000 K above the lower set of experimental data and *ca.* 500 K above the more recent values represented by two points in the figure (whose authors stress that their values are only a lower bound to T_m); our agreement with the higher values is close. In comparing with shock data, we recall the well-known difficulties of obtaining temperature in shock experiments. Our curve falls significantly below some of the T_m measurements, in which temperature was deduced by measuring optical emission, but accords quite closely with the other single point, obtained by estimating the shock temperature using thermodynamic calculation.



Our T_m of 6670K at the ICB pressure is close to the values inferred from some previous estimates, and gives evidence against the much higher or lower estimates sometimes proposed.

How good is our melting curve? Based on the analysis of the non cancelling errors in the Gibbs free energies for the liquid and the solid, and on the possible systematic errors due to GGA, we believe that our predictions for T_m are good to ± 600 K. This suggests that they are consistent with some of the recent DAC T_m ; at high pressures, we are consistent with some of the shock T_m .

The *ab initio* techniques for calculating thermodynamic properties under extreme conditions are expected to find application to many other problems concerning the Earth's deep interior.

We conclude by remarking again that this work involved enormous computational effort, which has been possible only thanks to the CSAR service.

References

1. Alfe, D., Gillan, M. J., Price, G. D. The melting curve of iron at the pressures of the Earth's core from *ab initio* calculations. *Nature* **401**, 462-464 (1999).
2. <http://news2.thls.bbc.co.uk/hi/english/sci/tech/newsid>
3. Gillan, M. J. The virtual matter laboratory *Contemp. Phys.* **38**, 115-134 (1997).
4. Car, R. & Parrinello, M. Unified approach for molecular dynamics and density functional theory *Phys. Rev. Lett.* **55**, 2471-2474 (1985).
5. de Wijs, G. A., Kresse, G., Vocadlo, L., Dobson, D., Alfe, D., Gillan, M. J. & Price, G. D. The viscosity of liquid iron at the physical conditions of the Earth's core. *Nature* **392**, 805-807 (1998).
6. Kresse, G. Furthmuller, Efficient iterative schemes for ab-initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B* **54**, 11169-11186 (1996).
7. Alfe, D. Ab-initio molecular dynamics, a simple algorithm for charge extrapolation. *Computer Phys. Commun.*, **118**, 31-33 (1999).

Feature: CSAR Fujitsu VPP300

— Manjit Boparai,
CSC/CSAR Systems
Administrator

Background

The Fujitsu VPP300 (hostname "fuji") is a vector supercomputer with 8 CPUs, each rated at 2.2Gflops/sec peak performance and each with 2Gbytes of distributed SDRAM memory. The system, which utilises modern CMOS technology and is air cooled, is housed in the machine room at Manchester Computing, University of Manchester (UoM). It is closely coupled with the CSAR Origin2000, fermat, for use of the HSM/tape archive facility.

The original system, a 3 CPU Fujitsu VPP300, was moved from Rutherford Appleton Laboratory and arrived at UoM in January 1999 with minimal disruption (a fully operational system was available to users within two days of the hardware arriving at UoM).

Surge in Usage and Subsequent Upgrade

Usage of the Fujitsu began to increase rapidly from spring until the summer, when the machine capacity proved insufficient to cope with the workload. The system was tuned as far as possible to manage the huge surge in batch workload, e.g. kernel parameters, NQS parameters, use of the jobswap facility to enable short development jobs to have faster turnaround. Essentially, however, we simply did not have enough processors to cope with the workload. (Interestingly, as usage of the Fujitsu grew, so did NERC's usage of the CSAR T3E, turing.)



FUJITSU

VPP300 - "CSAR Fuji"



The remainder of the summer thus saw negotiations between NERC and CfS for possible enhancements to the system, culminating in a decision to upgrade the Fujitsu to an 8-CPU system as soon as practicable. The 'go-live' date for the upgraded Fujitsu was planned to be 1st November.

A new 8 CPU Fujitsu VPP300 arrived at UoM on 19th October 1999 and was temporarily installed alongside the existing 'live' 3 CPU system so this could continue to operate for as long as possible whilst the new system was configured and tested. Over the next few days systems work and electrical testing were carried out. This was completed by 25th October. We were left with one day to spare since the switchover to the new system had been announced to take place on 27th October, with 'live' usage brought forward to start from 28th October.

The switchover from the old to the new system on the 27th October went very well, with the final 8-CPU system operational as of late afternoon on the same day. Finally, it was decided to wait another day before formally opening up the system for general use on Friday 29th October. This allowed us the time to ensure changes made during the upgrade were transparent to users. A few keen users were able to logon and run batch jobs from Thursday 28th October in order to test the system.

In summary, the upgrade was a success, and has provided us with a (much needed!) system with almost treble the performance we had 10 months ago.

Using the Fujitsu VPP300

Access to the Fujitsu is available via ATM and FDDI networks (originally FDDI only) at 'fuji.cfs.ac.uk' and 'fuji-fddi.cfs.ac.uk' respectively. The default route uses ATM.

The Fujitsu runs the UXP/V operating system which is based on standard Unix System V Release 4.1, enhanced for the vector parallel architecture.

Users can logon interactively to the primary processor (PE0) for development work. All 7 secondary processors (PEs 1-7) are dedicated to batch work, which is managed via the Network Queueing System (NQS).

Service tokens can be traded to enable usage of the Fujitsu like any of the other CSAR HPC systems, so that PIs and users can choose the most effective machines for their research.

The CSAR Application Form has been updated to include Fujitsu resources and should be used for all classes of application and all research councils.

Software Libraries available are NAG, BLAS and Fujitsu's Scientific Subroutine Library. Other software currently available includes emacs, nupdate, perl, gzip, etc. Also, vector versions of the cp, mv and cpio commands are available.

Compilers available are Fortran 90 (frt), VPP Fortran (same concept as HPF), ANSI C compiler and a vectorising C compiler (vcc). Additionally, users can compile Fortran and C programs for the VPP300 using the Fujitsu cross compilers which are available on fermat.

Application Analysis tools are available for code profiling and CPU performance.

User disk areas on the Fujitsu VPP300 are located on the /home and /work filesystems (currently 76Gb in total) and on the /tmp filesystems (66Gb in total).

CSAR wins award at SuperComputing 99

Dr John Brooke,
Head of CSAR
Applications Support

The University of Manchester (UoM) presented a Research Exhibit at SuperComputing 99 (SC99). This was the 10th in the series of SC events that started back in Portland, Oregon (USA) in 1989 and was returning there for the 10th anniversary.

HPCGames Challenge at SC99

This was, however, the first time that the University of Manchester has presented a stand at the exhibition. The SC events have become the major focus for manufacturers and research groups to demonstrate the very latest techniques in supercomputing and they are extremely competitive events. When you are exhibiting alongside organisations such as the US and Japanese government funded major research laboratories and initiatives such as the ASCI (Accelerated Supercomputing Initiative), you have to present something a bit special to get noticed. This involves forward planning and cooperation since it is necessary for European sites to collaborate in order to achieve the levels of the US and Japanese sites. CSAR and MRCCS contributed experiments and presentations to the UoM exhibit.

In particular, CSAR and MRCCS have been working with HLRS Stuttgart (a centre of excellence for visualization and high performance computing in Germany) to test the PACX-MPI message passing library (developed at HLRS) over European high-speed networks. This work has also involved cooperation with Japanese sites such as JAERI (Japanese Atomic Energy Research Institute) and ETL (Electrotechnical Laboratory). In particular we are comparing the performance and functionality of PACX with STAMPI, developed at JAERI also for the purpose of supporting grid-based metacomputing.

As a result of this collaboration, a joint presentation from Manchester Computing, HLRS, ETL along with US partners PSC (Pittsburgh Supercomputing Centre) and SNL (Sandia National Laboratory), were joint winners of the top prize for the most challenging scientific application in the HPC Games competition. The awards are presented each year to projects that demonstrate novel uses of supercomputing to stretch the limits of what high performance computing can achieve. The team presented metacomputing experiments to run single parallel programs on four supercomputers connected by links spanning three continents. The ultimate aim of these experiments is to show that global supercomputing is possible, thus enabling problems too big for any single machine to be tackled by a global network of computers. A report of this work has been published in the December 3rd edition of HPCWire. A picture of a rather tired but happy trans-continental team can be seen in the photograph overleaf:



The award-winning team, including Dr John Brooke from CSAR (right of picture)

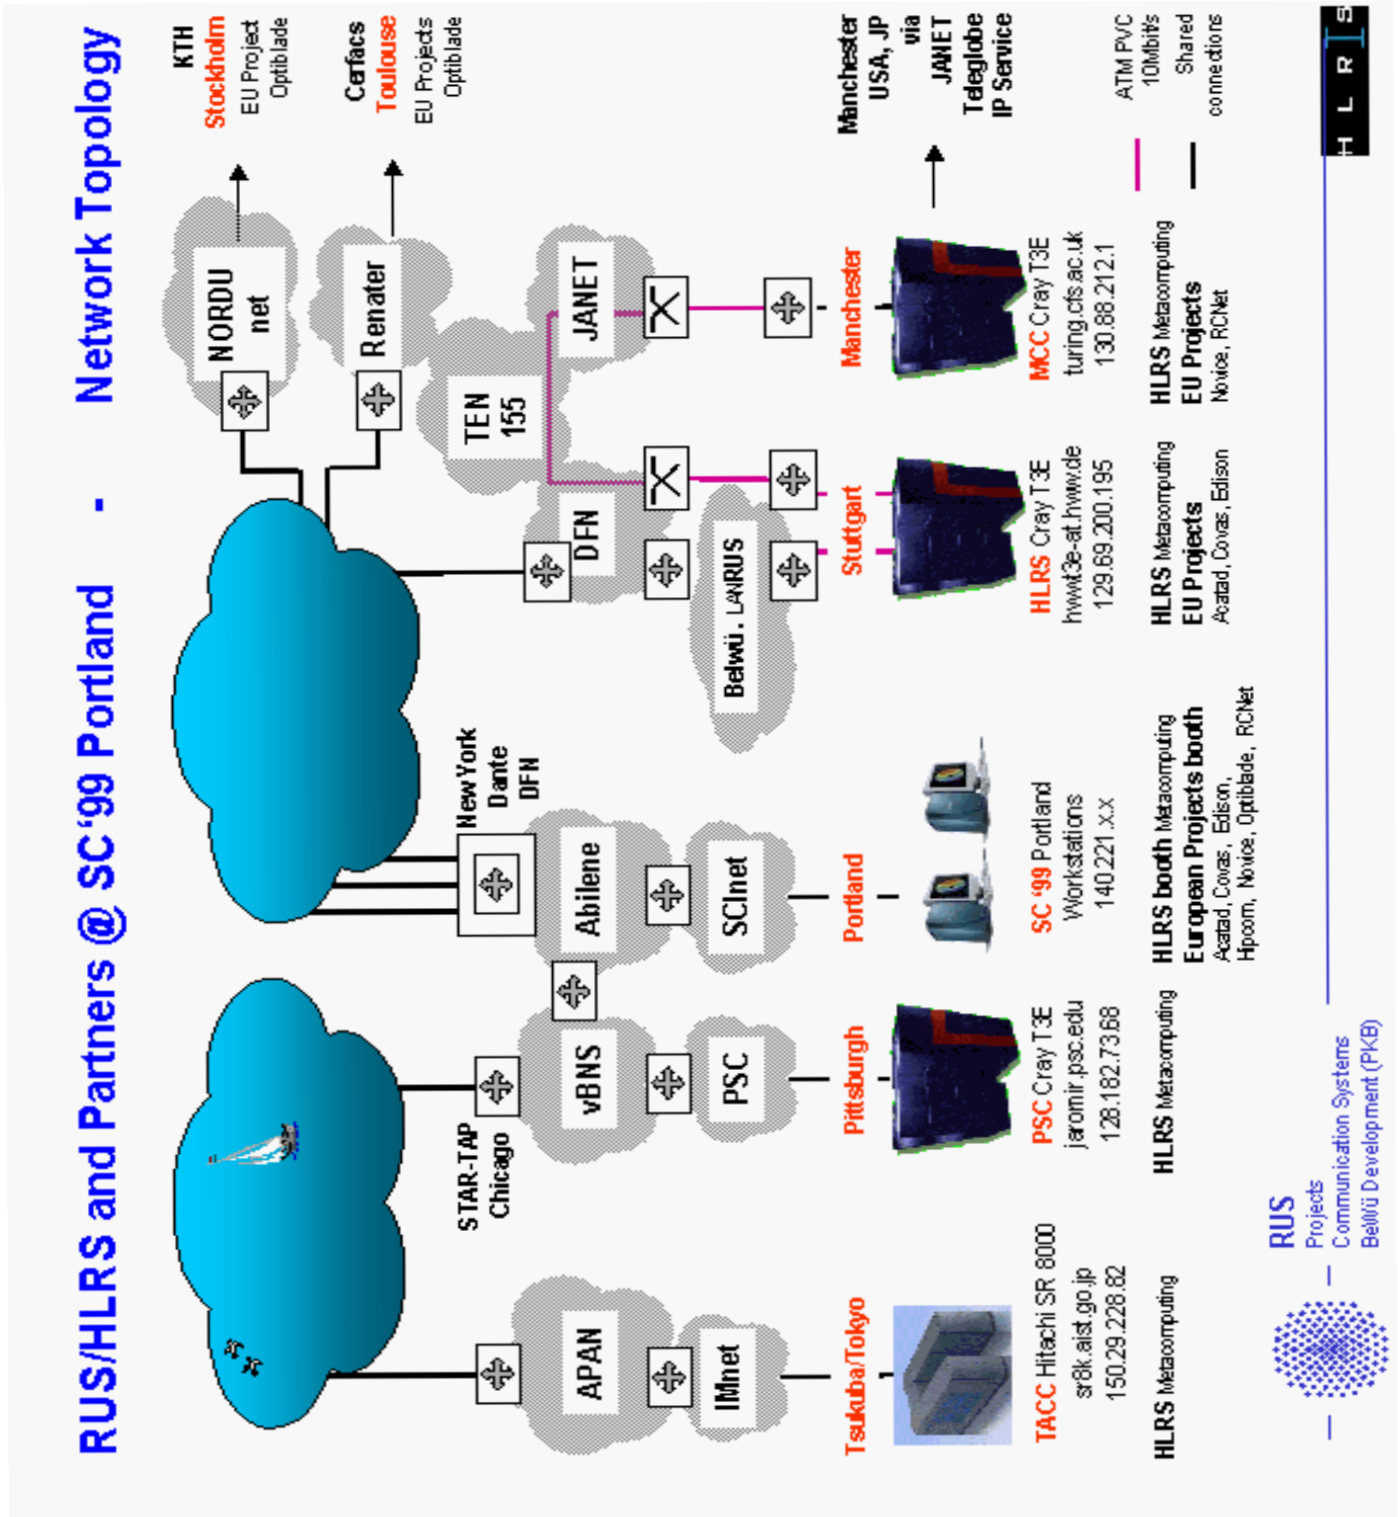
Exploitation of Global Networks for Global Computing

Massively parallel supercomputers are designed so that the processors inside them can communicate via channels that have very high information bandwidth and very low latency. In connecting together such supercomputers it is currently not possible to get close to this level of network performance but it is vital that such characteristics are maintained as much as possible. Thus conventional internet routes are not feasible because the high levels of traffic leads to highly variable bandwidth and latency which would degrade the performance of the cluster to levels which are unacceptable. Hence we need to exploit the possibilities afforded by special high-performance networks, and in so doing demonstrate that the investment in such networks leads to a corresponding increase in scientific productivity.

Today such national high-speed networks exist and in the European Community trans-national networks are now in place and are being tested. On a trans-continental level, however, such networks have not been fully tested. These experiments linked the three major centres of scientific computing activity, in the US, Europe and Japan, and are an important step towards the goal of establishing a genuinely trans-global network for supercomputing applications.

The layout of the network can be seen in Figure 2. The Manchester-Stuttgart link was a dedicated Permanent Virtual Circuit (PVC) using ATM communications. The advantage of a PVC is that it gives a consistent Quality of Service (QoS), in our case up to 10 Megabits per second (communications technology quotes transmission speeds in bits rather than bytes).

Figure 2: The network topology used in the ter(r)acomputer.
(Reproduced by kind permission of HLRS)



We would like to thank the network providers, Janet (UK), DFN (Germany) and the Dante project (operating the European TEN-155 ATM network) for establishing the PVC. We tried to setup a dedicated link between Manchester and the US research networks but this was not possible. Ironically, the day after the final experiment we were unexpectedly offered a special UK-US link. There was a dedicated link between Stuttgart and the US. All routes within the US ran over the Abilene research network but this was shared IP traffic rather than a dedicated point-to-point PVC. A major part of the whole experiment was establishing these links, and the experience gained in this is an important component in the drive to establish global metacomputing grids.

The CSAR/UoM Component of the Experiments

The CSAR/UoM contribution involved a collaboration between the CSAR service and workers at Jodrell Bank searching for binary pulsars in globular clusters. Extensive processing of the data from the radio telescope is necessary to reconstruct the original signal from the effects of the column of interstellar material intervening between earth and the globular cluster (known as dispersion). Once this has been done it is possible to detect the signals of several pulsars in the same cluster since, being a compact object, the effects of the interstellar medium will be very nearly uniform across the cluster.

The intensive work of searching the signal with a variety of trial dispersions was farmed out between the three T3E machines at Manchester, Pittsburgh and Stuttgart. The data was initially stored on Turing and was transmitted to the remote machines by message passing as an integral part of the experiment.

Since the data comprises several Gigabytes worth of observations, this is clearly a very severe test of the intercontinental links in terms of consistency of bandwidth. The problem needed to be very carefully load balanced between the machines, because the speed of data transmission between the machines was over 500 times slower than the speed of transmission between the processors in a single machine. The metacomputing task involved solving this problem by overlapping the data transmission with the processing so that, apart from some time at the start and finish, transmitting and processing of data are concurrent.

For the major experiment on Wednesday 15th November, we ran with a configuration of 512 processors on Turing and 128 processors at both Stuttgart and Pittsburgh. All the machines finished approximately together, with time differences a small fraction of the total time. As the inter-machine links improve it will be possible to make these figures more balanced, in this case we were concerned to show that even given the current state of intercontinental networks such experiments are possible.



M80 - Actual Globular Cluster

An important part of the work involved setting up and testing the trans-continental networks and here the help of the Networking group at Manchester Computing was invaluable, thus showing that the siting of the national CSAR supercomputing facility at a major UK networking centre enables maximum use of the supercomputing resource. A joint programming effort between Stephen Ord from Jodrell and Stephen Pickles from CSAR adapted the search code to use PACX-MPI rather than the native MPI on the T3E, and to take into account the load-balancing requirements of the metacomputer. It is remarkable that they were able to do this, since the Jodrell/CSAR collaboration only started in mid-October, less than four weeks before the live demonstration on November 15th.

The Ter(r)acomputer

One application run by HLRS utilized all four machines, thus showing that PACX-MPI can cope with clusters of heterogeneous architecture. The combined theoretical peak of all four machines was 2.2 Tflops and the Ter(r)acomputer spanned 10,000 miles. The application that utilized all four machines was a molecular dynamics simulation with over 4 million particles. HLRS also ran a CFD code, URANUS, that simulated the reentry into the atmosphere of a space vehicle. This code was developed for the international space station and is linked to a collaborative visualization project with SNL. The collaborative visualization was shown working between two different booths at SC99. The molecular dynamics and CFD applications related to different challenges from the bandwidth issues of the Manchester work, namely the need to mask latency by overlapping communication and computation.

Thus the metacomputing network was thoroughly tested and the results of the experiments will be described via technical papers at conferences and via the WWW. The CSAR WWW pages will have links to all relevant work.

Where next?

The particular application that CSAR/UoM brought to this project was the processing of experimental data via supercomputers. In the case of the binary pulsar, this involves searching through data produced by scanning a globular cluster at a narrow range of radio frequencies. Traditionally this has been done using specialised hardware called a filter bank. However for very high-frequency pulsars (frequency here means the frequency of the pulses from the object) this method is not yet practical. Such very high-frequency pulsars in binary systems, where the objects are close together, provide a testing ground for gravitational physics. First, though, it is necessary to detect the pulsars. This is made difficult by the phenomenon of frequency dispersion, as the radio waves travel through the interstellar medium between ourselves and the cluster the signal gets smeared with the time of arrival being a function of the radio frequency. If we know the dispersion relation we can restore the original sharp pulse by processing the signal.

In order to do this we need to trial various dispersion slopes looking for signals well above the background noise. It is this search that involves many simultaneous FFTs, going from the time to the frequency domain and then back again, and is very expensive in processing time. Once the dispersion slope is found this will be (to a reasonable approximation) the same for all objects in the cluster. On theoretical grounds it is expected that there will be several binary pulsars in the one cluster and these can be detected from the reconstructed signal with the dispersion removed.

Work is in place to set up a fast link from Jodrell Bank to the Network Northwest ATM network and this will enable the signal to be fed directly into Turing thus optimising the whole search process. At present the data is copied onto tape at Jodrell and sent to us to be read onto Turing. We would be interested to hear of other institutions with large data-processing needs, since they may also benefit from using supercomputing resources. To enquire further, please email me at: j.m.brooke@mcc.ac.uk or telephone: 0161 275 6814

More details are of our metacomputing work are available on the CSAR WWW pages at:

<http://www.csar.cfs.ac.uk/staff/costen/sc99>

We would be interested in working with other UK HPC facilities to develop this grid-based metacomputing to link HPC facilities. The 3rd European Metacomputing workshop will be held at HLRS on May 3-5 2000, please see:

<http://www.hlr.de/meta2000.html>

This feature is intended to be a regular part of *CSAR Focus*, describing various initiatives by CSAR Applications and Optimisation staff which are of potential interest and benefit to users of the service.

Free Code Evaluation

As you will be aware, Applications and Optimisation support on the CSAR service can be "purchased" via support tokens. Such support needs are applied for on the relevant Class 1, 2 or 3 forms when applications are made to use the service. This is a very new model for UK researchers and so we are thinking of ways to work with CSAR projects to determine their support needs. As part of this initiative we are offering all CSAR projects upto 3 days worth of support effort independently of the tokens scheme.

All projects have been allocated a liaison contact from the Applications and Optimisation teams. If you wish to discuss how your application or code could benefit from our support, please approach your liaison contact to discuss this. (If you do not know who your liaison contact is, please contact the CSAR Helpdesk.) The 3 days of support is intended to allow us to evaluate your application and discuss ways in which it could make fullest uses of resources available. Several projects already have support from the current HPCI centres. In such cases we would liaise with the relevant staff at those centres to discuss collaboration between the CSAR team and themselves for the benefit of your application.

CSAR staff could work with you to profile your application, and to determine where it could benefit from specialised optimisation to exploit the capabilities of the processors on Turing.

By sharing this expertise with computational scientists, who understand the scientific issues behind the organisation of the code, we believe that more could be achieved in terms of performance on Turing. This is because any optimisation must operate within this structure. Optimisation work also needs to be integrated with the application and to be consistent with evolving functionality as code develops.

Another area where we could help is in advice on pre and post-processing. For example the use of visualization to enhance the insight you obtain from your results; issues arising from handling large volumes of data; exploiting the multi-architecture nature of the CSAR service in using software available only on the Origin 2000 to process results obtained on the more specialised T3E. We have staff involved in investigating all of these areas, particularly meta-computing and visualization at present, with other projects planned. It was precisely to offer this range of flexibility that the service was established with a multi-architecture structure and with the provision for adding guest systems as they are requested, or as they become available to the CfS consortium.

For more details, or for further discussion, please contact John Brooke or Mike Pettipher, or contact your project liaison contact directly.

As part of this initiative we are offering all CSAR projects upto 3 days worth of support effort independently of the tokens scheme.

MRCCS News: Report on the MRCCS HPC Summer School

— Dr Mike Pettipher
Head of CSAR Optimization



Summer School on Distributed Memory Programming and Scientific Visualization University of Manchester 6th - 17th September 1999

The Manchester Research Centre for Computational Science (MRCCS) is a focus for high performance computing activities within the University of Manchester, bringing together research groups of international standing in high performance computing, visualization, data mining and the applications of such technology to a range of scientific problems.

This Summer School was the first annual two-week residential event organised by the MRCCS. It was designed to show how to maximize productivity on distributed memory machines, either highly specialized systems such as the Cray T3E or on clusters of workstations - and how to visualize the results. The residential nature and duration of the event provided an ideal environment for the participants to meet and discuss with, international experts in these fields, and to develop contact with the other participants.

The speakers were: Professor Jack Dongarra (University of Tennessee), Dr Len Freeman (University of Manchester), Terry Hewitt (University of Manchester), Dr Lawrence Mulholland (NAG, Oxford), Dr Hans-Georg Pagendarm (National Aerospace Research Centre, Germany) and Michael Resch (University of Stuttgart).

There were 29 registered delegates, most from various parts of the United Kingdom, but also from Holland, Germany, Norway and Brazil. Their backgrounds were equally diverse, including CFD, Computational Chemistry, Mathematics, Computer Science, Earth Sciences, Environmental Modelling and Visualization. This diversity resulted in contact between people who would not normally meet each other when working in their own disciplines.

The format of the Summer School was to alternate between lectures and practical sessions, although flexibility did allow topics to be adapted according to the wishes and expertise of the participants. The main subjects covered were distributed memory programming with MPI, including performance analysis with VAMPIR; visualization with AVS; parallel numerical algorithms and the use of numerical libraries, including the NAG Parallel Library. The practical sessions were performed on a cluster of PCs, an SGI Origin2000 and the CSAR Cray T3E itself. During the second week participants developed a visualization presentation, based on data from a double glazing problem or a tornado (both provided). Members of the winning group were each presented with a bottle of wine on the last afternoon!

Some topics of particular interest:

Automatically Tuned Linear Algebra Software (ATLAS)

Jack Dongarra described this software which automatically generates efficient BLAS3 code for the machine on which it is to be run. This was demonstrated during the practical sessions, when the performance of various different implementations for matrix multiplication were compared (on a Pentium running Linux). The Atlas version was significantly the best.

Comparitive Visualization

Hans-Georg Pagendarm gave examples showing how scientific visualization can help in the design of an aircraft. The use of HPC simulation in conjunction with visualization provides feedback on the theoretical predictions. This results in more efficient use of all resources, especially wind tunnels, and reduces the time spent on development.

Parallel Algorithm Development

One of the practical sessions involved the development of a parallel algorithm for the integration of a simple function. Participants were able to see how refinements in the algorithm affected load balancing and consequently performance. The range of systems available, with different characteristics and different usage (particularly dedicated versus shared access), also demonstrated portability issues.

Three seminars were also given during the Summer School:

Medical Applications in VRML by Dr Nigel John from the University of Manchester

Nigel described a web based 'tool' for teaching surgical techniques, particularly the insertion of instruments into 3D bodies including the brain and the spinal column.

Mantle Circulation Modelling by Professor Huw Davies from the University of Liverpool

This work provided an excellent example of distributed memory programming (on the Cray T3E) in conjunction with visualization.

Scientific Visualization on Virtual Reality Responsive Workbench by Michal Koutek from the University of Delft. Michal described application of a virtual reality workbench including a case study on a flood control system.

In addition to the technical sessions, a number of other events were organised. One such event was a talk by Professor Tom Kilburn (now retired from the Department of Computer Science at the University of Manchester) on the development of 'The Baby' - the world's first stored memory computer, built at the University of Manchester in 1948. For further information, see the separate article about Professor Kilburn within this newsletter). This was followed by a visit to the Museum of Science and Industry, which now houses a replica of the original Baby, after it had been rebuilt in the same computer room that now houses the CSAR Cray T3E.

A few comments from the participants:

"We could learn the up-to-date information and look back to the beginning of the computer with Professor Kilburn."

(Dr Ming Cai, University of Loughborough)

"The summer school program was well-designed and informative. I like in particular the lecture-then-practical format. It gives a good feel for what has just been taught in the lecture."

(Bobby Cheng, University of Manchester)

"On the social aspects of the Summer School, I enjoyed meeting and talking to so many people from different disciplines of science. The social events and especially the Saturday trip made my one and a half weeks there enjoyable as well as educating."

(Savas Parastatidis, University of Newcastle)

"... the whole Summer School was very interesting. And of course meeting other people working with HPC was very nice."

(Vegard Eide, University of Trondheim)

"I think the course was successful for me, as indicated by the fact that just one day after the end of the summer school I started parallelizing my code using MPI."

(Dr Tanja van Mourik, University College London)

In conclusion, the Summer School was a great success, with much being learnt, and all participants sent home with very large folders of materials, as well as a CD, for further study. Preparations are now being made for the next Summer School, in the year 2000. Please see http://www.man.ac.uk/mrccs/summer_school/ for further details.

Finally it is recognised that an event such as this can happen only as a result of the hard work of many people. Particular thanks should go to the speakers, to the participants themselves and to staff at Manchester Computing.

MRCCS Seminars

The Manchester Research Centre for Computational Science runs regular seminars on all aspects of High Performance Computing and Data Visualisation. The seminars are usually held on Friday afternoons in the Department of Computer Science.

To see the latest timetable of seminars, see <http://www.man.ac.uk/mrccs/seminars/>

21 Jan

Terrence Fernando, CVE, University of Salford
Virtual Environments for Assessing Maintainability Issues of Products

28 Jan

Rupert Ford, Centre for Novel Computing,
Dept. of Computer Science, University of Manchester
Load Balancing the UKMO Weather Forecasting model

25 Feb

Keith Taylor, High Performance Computing Group, Manchester Computing
Running Distributed Applications between Three Heterogeneous
Supercomputers

24 March

Andrew Coward, Southampton Oceanography Centre
Running the OCCAM code on the T3E

MRCSS Summer School Lecture : Professor Tom Kilburn and the Making of the "Baby"

— Mohammad Khalil,
Manchester Computing

One of the highlights of the MRCSS Summer School was the lecture given to the summer students by Professor Emeritus Tom Kilburn.

"....The most exciting time was June 1948 when the first machine worked. Without question, nothing could ever compare with that." Tom Kilburn, Autumn 1992.


Professor Kilburn described how it all began in the days when Britain was still trying to get back to normal after the Second World War and when many essentials were strictly rationed.

He described how he teamed up with Freddie Williams, a well-known circuit engineer, at that time in Telecommunication Research Establishment (TRE). After the war, Tom was asked by TRE to join Freddie Williams at the University of Manchester and it was here in January 1947 that the 'Baby' computer was to be born.

Professor Kilburn explained how the project to continue digital storage of information on a Cathode Ray Tube (CRT) stemmed from one of the Radar Stream projects during the war. He clearly described a world where no software and hardware existed. At that time, the Mercury Acoustic Delay Line was considered the only true electronic store for future stored program computers. By Late 1946, Professor Freddie Williams had managed to store a single bit on a CRT.

It took Professor Kilburn several months to suggest an improved method that would use parallel digital store and by the end of that year they were able to store 2048 bits on the CRT which was truly Random Access Memory. In his report in December 1947 to TRE, Professor Kilburn wanted to build a machine that would test this storage and get his second year in Manchester.

"....The most exciting time was June 1948 when the first machine worked. Without question, nothing could ever compare with that."
- Tom Kilburn,
Autumn 1992.



MRCCS Summer School Lecture : Professor Tom Kilburn and the Making of the "Baby"

- Cont.

During 1947 and early 1948, Professor Kilburn led a team to build a small Scale Experimental Machine, the 'Baby'. Although there was a race on both sides of the Atlantic to produce the first true computer, on June 21st 1948, at the University of Manchester, the world's first stored-program electronic digital computer successfully executed its first time. The race was won by Professor Tom Kilburn and his team; the world's first stored program had been invented. Professor Kilburn then showed the workings of the instruction set for the 'Baby' and explained the function of the world's first stored computer program.

Professor Kilburn went on to explain the developments of many machines after the 'Baby'. The first of these was the 'Manchester Mark 1', this was the full-scale computer based on the 'Baby' and was to use a magnetic drum as a secondary storage and was completed at the end of 1949. In 1951 the 'Manchester Mark 1' was replaced by the 'Ferranti Mark 1', a commercial development of the Manchester machine. The 'MEG' and 'Transistor Computer' were the next to be developed by Professor Kilburn and his team. The 'MEG' was the world's first floating point machine and worked successfully in 1954. The transistor technology machine was duly completed in 1955. Professor Kilburn told the students how the main use for these computers was pure mathematics and that no one could have predicted the uses that exist today.

The 'MUSE' was next to be built; it was 1000 times faster than the 'Manchester Mark1' and in 1962 was developed into the commercial machine 'ATLAS'. The 'ATLAS' was believed to be the most powerful and sophisticated computer in the world at that time. Professor Kilburn's team provided large amounts of software as well as the hardware for this 'supercomputer'. This computer provided the ability to handle multiple number of programs. It also introduced virtual input/output and had a large operating system. Many of the facilities we take for granted on computers and operating systems today were invented here at the University of Manchester for the ATLAS.

The 'MU5' was started in 1965 in close collaboration with ICL and was aimed at making use of High-Level Programming Languages, as these were becoming increasingly popular in the world of computing. This was fully operational in 1974. The commercial ICL '2900 Series' was based on the design of the 'MU5'. Professor Tom Kilburn pointed out how the Chief Architect of ICL spoke of the MU5 last year at the Computer 50 Celebrations in Manchester, "The 'MU5' influenced the '2900 Series' and will continue to do so for 30 years into the next millennium".



M.Khalil and Professor Tom Kilburn

┌ MRCCS Summer School Lecture :
Professor Tom Kilburn and the
Making of the "Baby" └

- Cont.

Tom spoke of the University as a place where an excellent computing service has been provided since 1948. In the early days, service was provided for people who came and did the programming themselves on the machines.

It is true that without the inspirational ideas of Professor Tom Kilburn, Professor Sir Freddie Williams and their team, and their pioneering work on development of their many machines, that the computing world could not have been the same as we see it. Maybe the high profile services like CSAR may not have existed for some time to come.

As Professor Freddie William said " But one day it stopped, and there, shining brightly in the expected place, was the expected answer. It was a moment to remember. This was 21st of June 1948, and nothing was ever the same again."

Thank-you Professor Tom Kilburn for your contribution to the world as we see it today.

In this issue's People Focus, Joanna Leng and Ben Jesson, who both joined the CSAR Applications team during the summer, tell us a little bit about their backgrounds.....

Joanna Leng

I completed a degree in Biophysics at Leeds University in 1989. After this I worked in medical laboratories where I trained and qualified as a Medical Laboratory Scientific Officer and helped develop their work/computer systems.

In 1995 I came to the University of Manchester to do a MSc in Computer Science. My thesis was completed in the Manchester Visualisation Centre, it was in the field of medical visualization. Straight after this I worked on a JISC funded project at MVC where I gave and developed training courses in both medical and general visualisation, gave advice and support to academics using visualization for medical purposes and developed code for several medical visualisation projects. Users of the CSAR service may recognize some of my work from the CSAR 1999 calendar, the images of the Egyptian Mummy were produced as a part of this project.

「



」

My role in CSAR is to encourage the use of visualisation and support researchers who wish to do this. As a starting point I will be visualising seismic data produced by academics using the T3E. I will also be developing web based tools to help users set jobs on the CSAR machines and in the end to aid computational steering.

People Focus: Ben Jesson

I came to Manchester to join CSAR's optimisation group in July of this year, having come straight from Oxford University where I spent four years studying for a BA in chemistry, and then another four working for my doctorate. The subject of the latter was the application of linear-scaling density functional theory (DFT) molecular dynamics methods to the large-scale ab initio simulation of metallic systems. My research was entirely computational in nature, and involved writing a DFT molecular dynamics code, which I also parallelised for Oxford's 84-processor Origin2000 computer. Somewhere along the way I also managed to pick up an MA (Oxon) in nothing in particular.

Before going to Oxford I grew up in Sussex, spending far too much of my time playing with computers. Studying chemistry provided something of a break from this, during which time I managed to develop some more healthy pursuits. These included representing my university in the Oxford-Cambridge Varsity Juggling Matches of '93 and '94 (won one, lost one)! I also enjoy walking, photography and boats, occasionally attempt to play tennis and in May 1998 managed to find time to get married.

Despite these distractions, the computational nature of my doctorate reaffirmed that what I really want to do is play with computers, preferably big ones, and so here I am at CSAR! So far I am still learning the ropes, but I have already become involved with the organisation of the weekly seminar series organised under the banner of the Manchester Research Centre for Computational Science (see <http://www.man.ac.uk/mrccs/seminars/>) of which CSAR is a part.

I have also done quite a bit of code parallelisation and optimisation on the CSAR T3E for the UKCP consortium, with which I was associated as a student, and which is one of the biggest users of CSAR's computing facilities. So far this has been the aspect of my job that I enjoy the most, so I look forward to providing some optimisation support for a few more of the CSAR users out there!



┌ Paul White,
CSC/CSAR Systems
Administrator

No one can doubt that the Cray T3E has been an outstanding success in the field of massively parallel computation. A look at the latest (Nov. 1999) "Top 500" list of Supercomputers shows that 24 out of the top 50 machines in the world are T3Es (the CSAR machine being rated 14th overall and 7th in the list of T3Es).

In October, around 70 of us met for a 2-day workshop, organised by the Cray User Group in Princeton, NJ. The workshop covered all aspects of the Cray T3E - applications, software updates, and plans for the future. In addition to systems administrators (such as myself), there were developers from Cray/SGI and end users from a variety of educational, commercial and government sites in Europe and the US.

The close liaison between the developers and the users of the T3E, encouraged by the Cray User Group, has undoubtedly been a key factor in its success. Presentations from developers and user sites looked at the experiences that have been learned over the past 2 or 3 years with the T3E, and how many of these experiences have been fed-back to the developers in order to enhance the machine. An example of this is the "service limits" feature of the global resource manager, which has considerably reduced the wait time for interactive usage and was installed on the CSAR machine months ahead of its official release.

Papers were also presented on a wide range of applications that have been successfully run on the T3E. These included subjects as diverse as Geo-physical Processing, Atmosphere Modelling, Plasma Turbulence and Molecular Dynamics Simulations. A group working on Complex Magnetic Systems had achieved a sustained 1.02 Teraflops on a T3E with 1480 processors, using full 64-bit arithmetic throughout. There was even a project to deploy T3E systems on naval vessels, for which the criterion for success was measured in gigaflops per cubic foot!

It was difficult to find anyone with a bad word to say about the T3E - it may well turn out to be the most successful product that Cray have produced, by whatever method is chosen to measure success. The only dark cloud on the horizon was the uncertainty over the future of Cray Research (SGI having formed it into a separate business unit). By the time this article reaches the press, it may well be that the issue has been resolved.

Looking to the future, there were presentations on 2 successor products to the T3E - the SV2 from Cray Research and the SN-IA from SGI. Both products are unlikely to be generally available for at least 2 years, however. In the meantime, planned enhancements to the T3E, in both software and hardware, should ensure that the T3E remains at the forefront of research the world over. The CSAR Service will ensure that a world-class service will continue to be available for scientists and engineers in the UK.

┌ A look at the latest (November 1999) "Top 500" list of Supercomputers shows that 24 out of the top 50 machines in the world are T3Es. ┘

Software Update: FFTW – the Fastest Fourier Transform In the West!

Dr Keith Taylor,
CSAR Application Support

FFTW is a package of Discrete Fourier Transform routines which have been developed by Matteo Frigo and Steven G. Johnson of the Massachusetts Institute of Technology.

Introduction

FFTW has been steadily acquiring a growing reputation which culminated in the authors being awarded the 1999 J. H. Wilkinson Prize for Numerical Software. I can only provide you with a taste in this short article, but I hope your appetite will have been sufficiently whetted for you to investigate further on FFTW's home web-page at:

<http://www.fftw.org/>

FFTW's Pros And Cons

Amongst its pros are

- It's free!
- It's in the public domain.
- It's portable across a reasonably wide range of machine types.
- It yields reasonable performance with little effort.
- It includes parallel routines based on MPI.

It should also be noted that FFTW may contain more functionality than a particular machine's native routines. For example, there is no 1-dimensional Cray-specific parallel DFT for the T3E, but FFTW has one!

One obvious con is that it may not be as efficient as routines native to a particular machine, but it quite often outperforms these. There is also the slight annoyance that the MPI-based routines are only available as C functions, although it should be straightforward to interface them to Fortran. Fortran versions already exist for the non-parallel routines.

FFTW's Methodology In Brief

FFTW manages to achieve its remarkable combination of portability, flexibility and performance by a methodology which combines knowledge of your particular application, for example the length of transform required, with that of the computer architecture, to assemble an optimum collection from a set of 'codelets' designed to meet varying circumstances. This information is summarised in the form of a 'plan'.

Plans may be created by using a crude estimate of the machine's characteristics, or by a process which tests various combinations of codelets until the best one is found. The latter method takes significantly longer, but once created, the resulting plan can be used repeatedly in your code, and may even be filed away for future use! More details can be found in the various papers listed under 'Documentation' on the above quoted web page.

Some Results

Now for some evidence of FFTW's efficacy. In the following, I shall be comparing its performance on the Cray T3E-1200E here at Manchester, with results kindly provided by Robert Allan of Daresbury Laboratory (<http://www.cse.clrc.ac.uk/Person/r.j.allan>) using Cray-specific routines. The latter results are taken from a comprehensive survey, which is currently being prepared:

R. J. Allan and K. Taylor; "Serial and Parallel FFT Routines", edition 2, Parallel Applications Software on High-Performance Computers; CLRC Daresbury Laboratory, Daresbury, 1999.

In assessing the comparison, it should be born in mind that I made no effort to optimize the FFTW libraries when I installed them, preferring to see how the routines performed "straight off the web". In fact, only the compiler flag '-O3' seems to be employed during the default installation. Presumably, the Cray routines have been subjected to considerable performance enhancement by the company's optimizers.

Firstly, we examine

1-dimensional complex-complex results on 1 processor of a Cray T3E-1200E

Length of transform	Mflops/s using FFTW	Mflops/s using ccfft
12	8.292E-01	1.533E+02
25	2.195E+00	1.377E+02
32	3.043E+00	2.929E+02
50	5.221E+00	2.017E+02
64	7.057E+00	4.048E+02
100	1.178E+01	2.938E+02
200	2.521E+01	3.313E+02
400	4.904E+01	3.893E+02
800	9.022E+01	3.867E+02
1600	1.297E+02	2.762E+02
3200	1.364E+02	1.019E+02
6400	1.258E+02	7.759E+01
10000	1.253E+02	7.375E+01

Not too flattering at the shorter lengths! However, FFTW is gaining by a significant margin, once we have reached 3200 and beyond, in the region of more practical interest. After all, why use a T3E to crack a nut!

Mflops/s performance of complex 2-dimensional routines on T3E-1200E								
No. of PEs	4	8	16	32	64	128	256	512
Size, routine								
322, FFTW	202	211	125	82	---	---	---	---
pccfft2d	176	233	309	372	*	*	*	
642, FFTW	298	447	535	290	186	---	---	---
pccfft2d	317	510	812	629	768	*	*	
1282, FFTW	266	579	920	1114	603	429	---	---
	298	715	1070	1236	1510	1985	*	
2562, FFTW	215	470	1159	1755	2412	1258	988	---
pccfft2d	368	609	1440	2751	3015	2558	2234	
4802, FFTW	318	630	1149	1968	3079	4460	2668	1997
pccfft2d		588	1107	2038	*	*	*	
5122, FFTW	235	459	865	1677	3462	4966	3245	2282
pccfft2d	288	569	1094	2285	3787	6522	6118	
9602, FFTW	253	535	1181	2049	3586	5455	9077	5310
pccfft2d		620	1222	2206	3235	*	*	
10242, FFTW	244	471	911	1807	3059	5628	9917	5763
pccfft2d	305	586	1132	2220	3803	7803	11440	
19202, FFTW	249	486	960	1894	3357	6162	10115	17400
pccfft2d		571	1142	2250	3484	7064	*	
20482, FFTW	206	470	907	1713	3075	5994	10728	18928
pccfft2d	221	527	1020	2027	3464	7350	12210	
38402, FFTW	233	463	900	1723	3370	6318	11381	19505
pccfft2d		416	957	1883	3325	6746	10380	
40962, FFTW	187	377	851	1670	3064	5865	11076	19695
pccfft2d		329	730	1446	2515	5381	9525	

— indicates limit of benefit from parallelization reached

* indicates N not divisible by "number of PEs"

Again, FFTW is none too impressive for the smaller sizes, although it is holding its own to a reasonable degree. However, from 3840 onwards, FFTW is already outshining its native 'pccfft2d' competitor!

The Effects of Cluster Management on Parallel Applications

Fumie Costen,
CSAR Applications Support

As part of CSAR's research into metacomputing issues (see the article by John Brooke in this issue), I will present a paper at the 1st IEEE International Workshop on Cluster Computing

<http://www.dgs.monash.edu.au/~raj कुमार/tfcc/IWCC99/index.html>

Networks of commodity workstations can offer an effective platform for large scale parallel processing. However such systems do not come equipped with specialist operating systems and customized batch job software. We have investigated the running of parallel programs in a heavily loaded cluster. We are seeking to determine whether performance of such codes can be improved by efficient cluster management software.

In the paper to be presented at the workshop, we focus upon the widely used Load Sharing Facility (LSF) from Platform Computing and investigate its ability to run both parallel and serial simulations and increase the through-put of jobs using a cluster of SG workstations. This is of considerable interest to CSAR users as SGI have announced that LSF is to eventually replace NQE as the batch job management system on SGI hardware. Yet another sign that the distinction between commodity computing and High Performance computing is starting to blur.

Our experiments showed that LSF does a good job in terms of balancing the load for serial jobs and avoiding the machines with high processor utilization rates.

Using the result of the experiment, we created a useful diagnostic tool for assessing the impact of load balancing software.

We are continuing the research on this, looking at the influence of individual scheduling parameters on the elapsed time of the job to find the set of the optimum parameters under particular conditions of loading of the cluster. We are testing to see if we can extend the simple performance metric we have developed, to more general use. To do this, it needs to be tested under a wide variety of conditions and using codes which have differing message-passing and load-balancing characteristics.

For further interest, a postscript copy of the paper can be obtained from my page on the CSAR website:

<http://www.csar.cfs.ac.uk/staff/costen/>

A first meeting-ground with laboratory experiment enabled by turing

— Professor Kenneth Taylor
Queen's University Belfast

Turing has now made possible the first comparison between results from a full-dimensional numerical model of multiphoton-double-ionization of helium and those from laboratory experiment.

The comparison is displayed in Figure 1 which plots the ratio of He^{++} yield to He^+ yield as a function of laser intensity at a wavelength of 390 nm. The experimental results were obtained by Dr DiMauro and his group at the Brookhaven National Laboratory in the USA.

After discussions with Dr DiMauro at the 8th International Laser Physics Workshop in Budapest in early July 1999, the Queen's University Belfast group got their complementary calculations pushed through turing by the end of September in time to be presented in an Invited Talk by Professor Taylor at the 8th International Conference on Multiphoton Processes held in Monterey, California in early October 1999.

The importance of the theoretical and complementary laboratory work is that reliable quantitative knowledge can thus be gained on the dynamics of a fundamental three-body charged particle quantum mechanical system (i.e. two electrons and a bare nucleus) as it is driven far from equilibrium by a strongly time-dependent force. The light electrons of the system respond dramatically to this force -supplied by the very high-intensity femto-second laser pulse- whereas the much heavier nucleus makes oscillations of negligible amplitude. A full-dimensionality treatment of the electronic motion is essential to properly account for both the angular momentum absorbed from the laser pulse and the repulsive interaction between the two electrons.

The Time-dependent Schrodinger Equation (TDSE) governs the dynamics of the system. This is a partial differential equation (PDE), first order in the time and, when written in full-dimensionality for laser-driven helium, involves dependence on no less than 5 spatial variables. The Belfast group has over the past six years been developing, and exploiting, state-of-the-art algorithms to solve such a multi-dimensionality PDE on massively parallel supercomputers.

Each of the 5 calculated points on Figure 1 is obtained by solving the TDSE at the given laser intensity and frequency. The general rule is as follows: the longer the laser wavelength and/or the higher the laser intensity, the more computationally demanding are the calculations. A wavelength of 390 nm proves an excellent meeting-ground between laboratory experiment and calculation now possible with turing. Most experiments on helium are carried out at the Ti-Sapphire fundamental wavelength of 780 nm which unfortunately is still completely inaccessible for full-dimensionality calculations - even with the power of turing!- at sufficiently high intensities for double ionization to be possible. The work at 390 nm brought almost the entire power of turing into play especially for the two highest intensity points plotted in Figure 1. Thus the highest intensity point required runs accumululating to 48 wall-clock hours using simultaneously 455 of the 576 processors.




Figure 1. The ratio of the probability for doubly ionizing the helium atom to that for singly ionizing it plotted against laser intensity. The laboratory experiment is that reported by B Sheehy, R Lafon, M Widmer, B Walker and L F DiMauro *Phys Rev A* 58, 3942 (1998) and the turing calculation is by Jonathan Parker, Laura Moore and Ken Taylor of Queen's University Belfast. All experimental points have been moved up in intensity by a global 23% which is within the stated uncertainty in the laboratory determination of laser intensity.

The laser pulse after an initial ramp-on over 4 cycles is allowed to oscillate at a constant amplitude of electric field for at least a further 20 cycles to ensure no transient character in the electronic response remains.

The electrons of helium start out in their ground state but, controlled by the TDSE, evolve rapidly into one- and two-electron ionizing wavepackets leaving the atom. For the middle point plotted on Figure 1

(at intensity $6.5 \times 10^{14} \text{ W/cm}^2$ Figures 2(a) and (b) display respectively the time-dependence of the single and double ionization yield these wavepacket motions give rise to. Note that on each half-cycle of the laser pulse there is an additional step-wise contribution to each yield corresponding to the liberation of wavepackets. What is plotted in Figure 1 is the essentially constant value of the ratio between these yield rates once the initial laser ramp-on is over.




Figure 2. The cumulative ionization yields as a function of time for single ionization (a) and for double ionization(b) of helium for the turing calculation in Figure 1 at a laser intensity of $6.5 \times 10^{14} \text{ W/cm}^2$.

The analysis of the data produced in the turing calculations brings out much richer information on the dynamics of this three-body system than emerges from the Stoneybrook experiment. In this analysis, scientific visualisation methods have played a crucial role and these have also been developed for this problem in recent years by the Belfast group, most notably by Dr Daniel Dundas. For laser pulses of wavelength between 248 and 390 nm it has been possible from this scientific visualisation analysis to determine, for example, that in the two-electron ionizing wavepackets both electrons emerge on the *same* side of the nucleus and with near zero angle initially between their directions of motion. This result is *contrary* to that predicted by many other, less complete, theoretical approaches.

However new results from two German experiments reported at the Monterey meeting have provided the first experimental confirmation of the two-electron wavepacket dynamics found in the turing calculations.

Although these German experiments operated at 780 nm, there is no physical reason for the electronic dynamics to be qualitatively different from those explored in the turing calculations at 390 nm.

The electronic response of helium can however be expected to be very different for substantially shorter laser wavelengths e.g. in the vicinity of 20 nm. Free-electron laser sources in Germany and elsewhere are now making this new shorter wavelength range available to experimentalists. At traditionally used wavelengths (e.g. 390 nm) at intensities around $1.0 \times 10^{15} \text{ W/cm}^2$ the atom ionizes overwhelmingly in a sequential manner. At the much shorter wavelength of 20 nm, physical considerations borne out by exploratory calculations on turing indicate the contrary. Indeed the response of the less-tightly-bound outer two electrons of alkaline earth atoms can be expected to be much less sequential when exposed to laser pulses at standard wavelengths operating at lower intensities around $1.0 \times 10^{15} \text{ W/cm}^2$

「 A first meeting-ground with laboratory experiment enabled by turing 」 — Article cont.

The Belfast group, through the full-dimensionality approach developed on massively parallel machines, is now uniquely well-placed to investigate the qualitatively new physics arising in these cases.

The work at Belfast described above forms part of the activity of the Multiphoton and Electron Collision Processes HPC Consortium operating at CSAR.

HPC/CSAR Course Calendar

2nd Semester 99/00

The course calendar for the 2nd semester, is now available. Bookings for courses must be made in advance via the CSAR Helpdesk: csar-advice@cfs.ac.uk. All the courses will be held at the University of Manchester.

Courses are available to everyone. For up to date details, please see: <http://www.csar.cfs.ac.uk/using/courses.shtml>

February 2000

- | | | |
|-------------------------|---|--------------------------------|
| 1 | Using the CSAR Service | - D. Kidger |
| 17 - 18 (2 days) | Parallel Vector Processing on Fujitsu VPP 300 | - D.Kidger, M.Pettipher |
| 24 - 25 (2 days) | C++ for Scientific Applications | - J. Brooke |
| | Porting/Optimising for the Cray T3E | - D.Kidger |

March 2000

- | | | |
|-----------|--|----------------------|
| 13 | Utilising Clustering Systems | - M.Pettipher |
| 21 | Using the CSAR Service | - D.Kidger |
| 22 | Introduction to Programming with MPI | - S.Pickles |
| 27 | Introduction to High Performance Computing | - F.Costen |

HPC/CSAR Course Calendar

2nd Semester 99/00

April 2000

- 3** Shared Memory Parallelisation with Open/MP - **M.Bane**
- 4** Shared Memory Parallelisation with Threads - **J.Brooke**
- 5** Introduction to LSF for Clustered Systems - **M.Pettipher**
- 10** Using the CSAR Service - **D.Kidger**
- 11** Java for High Performance Computing - **A.Dodd, D.Kidger**
- 12** Introduction to High Performance Fortran - **K.Roy, Pettipher**

May 2000

- 15-17** (3 days) Porting/Optimising for the Cray T3E - **D.Kidger, M.Foster**
- 22** Introduction to Programming with MPI - **S.Pickles**
- 23** Advanced MPI - **S. Pickles**
- 24** Shared Memory Parallelisation with Open/MP - **M.Bane**

Introduction to High Performance Computing

Prerequisites: Fortran or C programming experience.

This is a one day seminar for new and potential users of the CSAR service, and for other people who are interest in the topic of High Performance Computing. The following topics will be covered:

- Need for HPC
- Vector and parallel computing
- Parallel decomposition
- Parallel programming paradigms

Introduction to Programming with MPI

Prerequisites: The ability to program in Fortran and C and familiarity with Unix. Some acquaintance with the terminology of High Performance Computing, such as that covered by the Introduction to High Performance Computing course, is also desirable.

A one day workshop with lectures and practical sessions.

MPI (Message Passing Interface) has emerged as the de-facto standard for writing portable, message-passing programs in parallel processing environments. This course introduces the basic essentials of MPI, yet the material covered is sufficient to parallelize a wide range of serial codes. Topics covered include:

- The message-passing paradigm
- Point to point communication
- Synchronisation
- Basic collective communication (reduction operations, broadcast)

Advanced MPI

Prerequisites: Attendance at the Introduction to Programming with MPI, or a good working knowledge of the topics covered. The ability to program in Fortran or C and familiarity with Unix.

A one day workshop with lectures and practical sessions.

This course covers advanced features of MPI (Message Passing Interface), including the following topics:

- MPI derived types
- Groups, Contexts, Communicators
- Process Topologies
- More collective communication
- Future developments (MPI-2)

Introduction to High Performance Fortran

Prerequisites: The ability to program in Fortran or programming experience in other high level languages.

High Performance Fortran (HPF) is a superset of Fortran 90 containing extensions to Fortran 90 to support data parallel programming. HPF allows programmers to produce portable code for parallel machines, without the difficulty of providing explicit message passing calls. This one day 'hands-on' course provides an overview of the HPF language. Major topics include: relevant features of Fortran 90, the HPF model, HPF directives, data distribution, data parallelism, intrinsic and library procedures, extrinsic procedures and subset HPF.

Shared Memory Parallelisation with Threads

Prerequisites: The ability to program in C and familiarity with Unix.

Multithreading is a modern programming paradigm for implementing parallel applications on shared memory multiprocessors. This 1 day course is targeted at individuals with not previous MT programming experience. Topics covered include:

- Parallel computing and multithreading
- Multithreading basics
- Creating and managing threads
- Synchronising threads

On line access with examples and exercises is an integral part of the course.

Shared Memory Parallelisation with Open/MP

Prerequisites: You should have used either Fortran 77 or Fortran 90 to write programs. Familiarity with Unix, emacs or vi, and the X Window System is highly desired.

OpenMP can be used to run Fortran and C programs in parallel on a shared memory computer, such as the SGI Origin2000s kilburn.mcc.ac.uk and fermat.cfs.ac.uk. The OpenMP specification is now widely accepted and suitable compilers are now provided by all major hardware vendors, including SGI, Sun, Compaq/DEC, Intel, HP and IBM.

To use OpenMP you simply add directives to your code. This 1 day course will cover all the important directives and include plenty of hands-on exercises for you to learn how to parallelise your code. Where appropriate, emphasis will be given to the practical use of OpenMP.

The day will also include a very brief overview of concepts of parallel programming and an introduction to the performance analysis tools perfex and SpeedShop.

About a third of the day will involve practical exercises. Note that the examples will be in Fortran only, and will be run on an Origin2000. C examples will be available on the C++ for Scientific Applications course.

Utilising Clustered Systems

Prerequisites: Familiarity with a high level language such as Fortran or C.

Utilising clustered systems (whether workstations or PCs) for parallel computing has emerged as a key technology over the last few years. This course will present an overview of the current developments in hardware and software and covers system and network configuration, languages, libraries and tools which can be used to develop and utilise these clusters for parallel computing.

An Introduction to LSF for Clustered Systems

Prerequisites: The ability to program in Fortran or C and familiarity with Unix.

When individual systems are combined into a single clustered resource, it is possible to increase the overall throughput of work, and also to run parallel programs which may require more memory than is available on a single system. This can be effectively achieved only via an appropriate Cluster Management System. LSF (Load Sharing Facility) is a widely used and highly regarded system. This one-day course will describe the basic features of LSF showing how it may be used to make best use of a clustered system.

Introduction to Using the SGI Origin2000

Prerequisites: You should have used Unix before and be familiar with Fortran.

This new 1 day course acts as an introduction to using the SGI Origin2000 high performance computer. Topics covered include: overview of the system architecture, making good use of the compilers, performance analysis tools, and running programs interactively and in batch under NQS. There will be a brief discussion of the use of the autoperalleliser.

This is not an introduction to Unix course. For in-depth optimisation and more discussion and practicals concerning shared memory parallelisation, please attend the 3 day SGI Origin2000 Optimisation and Parallelisation course.

SGI Origin2000 Optimisation and Parallelisation

Prerequisites: Practical experience of programming in either Fortran 77 or Fortran 90 is essential. Familiarity with Unix, X Window System, and make/makefiles is required.

This intensive 3 day course will help users understand the key features of the SGI Origin2000 and how to get the most from codes run on this machine.

Issues covered include: system architecture, program development tools, RISC optimisation techniques, performance analysis tools, compiler options, concepts of parallelisation, data dependency analysis and shared memory programming using OpenMP

The course will be very intensive and include several practical sessions (Fortran only) covering the salient points of the course. C examples will be available on the *C++ for Scientific Applications* course.

Java for High Performance Computing

Java is now firmly established as the language of choice for many application areas, particularly for situations where the code must be extremely robust and portable. It naturally supports multi-threading which allows a code to run simultaneously over several processors. However, currently Java code is slow to run compared to Fortran90 and C.

This course looks at why Java is a good language for 'application codes' such as in computational chemistry or fluid mechanics, and explores how these codes can be made to run almost as fast as if in native Fortran whilst maintaining Java's flexibility and portability, and how the same code can be run on single and parallel architectures.

C++ for Scientific Applications

Prerequisites: Familiarity with the topics covered in the Programming in C++ course or considerable programming experience in C or C++. Familiarity with the topics covered in the Introduction to Unix course.

This course examines the use of C++ in high performance computing applications. The course concentrates on the implementation of numerical methods in C++, and how object-oriented programming can create general classes for scientific applications. The course will be given on a Unix platform.

Using the CSAR Service

Prerequisites: Experience of programming in a Unix environment (Fortran, C, Java etc.).

This is a one day seminar for new users of the CSAR service. It provides basic information on how to use the systems and resources provided by this service. The following topics will be covered:

- Basic hardware configuration
- Files, file systems and tape archiving
- Available software packages (including shells, editors, compilers and applications)
- Compiling a sample program
- Submitting batch jobs on turing, ferat and fuji
- Basic profiling and debugging
- Self registration for new users
- Support (WWW, email, the helpdesk system, specialist applications staff, service quality tokens)

Parallel Vector Processing on Fujitsu VPP 300

Prerequisites: Fortran or C programmer with message passing experience, Unix.

This is a two day course with lectures and practical sessions which is designed to cover the fundamentals of vector processing, with specific reference to efficient usage of the Fujitsu VPP 300.

The following topics will be covered:

- System architecture
- Vectorisation
- Parallel issues specific to VPP 300
- Optimising tools
- Debugging tools

Porting/Optimising for the Cray T3E-1200E

Prerequisites: Fortran or C programmer with message passing experience, Unix.

This is a 3 day workshop which includes lectures and practical sessions.

The Cray T3E, which provides the major computing resource of the CSAR service, is a very powerful, but also a limited resource. This course is designed to provide information and instruction for programmers who wish to use the specific hardware features of the system to maximise the efficiency of their codes.

- System architecture
- Single node optimisation
- Use of libraries
- SHMEM programming
- Benchlib
- I/O
- Optimising tools
- Debugging tool

Forthcoming Events

Check out our list of conferences. This includes not only CSAR related events but details of other events related to High Performance Computing:

February
2000

23-25

Workshop on Advanced Data Storage/Management and Techniques for HPC, CLRC Daresbury Laboratory, UK
<http://www.dl.ac.uk/TCSC/datamanagement/conf2.html>

23-25

Second International Workshop on Next Generation Climate Models for Advanced HPC Facilities, Toulouse, France.
<http://www.tokyo.rist.or.jp/workshop/>

March
2000

30-1 April

Meeting on Large-Scale Computations in the Simulation of Materials Center for Nonlinear Analysis, Carnegie Mellon University, Pittsburg, USA. <http://www.math.cmu.edu/cna/conferences.html>

April
2000

3-4

AWE International Conference on Exploiting Leading-Edge High Performance Computing, Keble College, Oxford
<http://www.awe.co.uk/hpc>

May
2000

3-5

3rd European Metacomputing Workshop, HLRS Stuttgart, Germany
<http://www.hlrs.de/meta2000.html>

10-12

The 8th International Conference on High Performance Computing and Networking Europe, Amsterdam, The Netherlands.
<http://www.wins.uva.nl/events/HPCN2000/>

22-25

2000 International Parallel CFD Conference, Trondheim, Norway
<http://www.parcfd.org>

Forthcoming Events

June
2000

21-23

4th International Meeting on Vector and Parallel Processing (VECPAR'2000), Porto, Portugal.
<http://www.fe.up.pt/vecpar2000>

July
2000

9-21

9th Annual EPSRC Summer School in Numerical Analysis, University of Durham, UK.
<http://fourier.dur.ac.uk:8000/nass>

August
2000

16-18

ERCOFTAC Symposium: From Tetraflops to Petaflops - Supercomputing in the New Millennium, HLRS Stuttgart, Germany
<http://www.hlr.de/news/events/1999/peta.html>

29-1 September

Euro-Par 2000, the Technical University of Munich, Munich, Germany.
<http://www.bode.in.tum.de/~europar/>

September
2000

7-8 (to be confirmed)

Sixth European SGI/Cray MPP Workshop, Manchester Computing, University of Manchester, Manchester, UK

Early September (date to be confirmed)

Supercomputing with Linux, Second MRCCS High Performance Computing Summer School, Manchester Computing, University of Manchester, Manchester, UK

For the timetable of weekly MRCCS seminars held at The University of Manchester, please see the section "MRCCS Seminars" (page 20).

CSAR FOCUS

Edition 4 January 2000, Editor: Kaukab Jaffri



The CSAR Web site (<http://www.csar.cfs.ac.uk/>) has help and information on all aspects of the service, and includes sections on Software, Training, Courses, Registration & Project Management, and links to other HPC sites.

CSAR Focus is also available to view on the Web: <http://www.csar.cfs.ac.uk/general/newsletter.shtml>.

CSAR News Pages

All news items related to CSAR are described via these web pages: <http://www.csar.cfs.ac.uk/general/news.shtml>.

CSAR Helpdesk

If you require help on any aspect of the CSAR service, you can contact the CSAR Helpdesk Team. Kaukab Jaffri or Caroline Evans will deal with your query promptly and efficiently.

CSAR Helpdesk by phone: 0161 275 6824/5997

Alternatively, you may contact the CSAR Helpdesk via email, which is just as prompt to respond, as your call will be automatically be logged using the latest call logging system (Action Request System).

CSAR Helpdesk by email: csar-advice@cfs.ac.uk

The CSAR Helpdesk is open from 8.30am - 6pm Monday to Friday, except on Public Holidays.

University of Manchester Contacts

Head of National Services:			
Julia Chruszcz	0161 275 6100	j.chruszcz@mcc.ac.uk	
Executive Director:			
Louis Blanchet	0161 275 6062	p.l.blanchet@mcc.ac.uk	
Head of CSAR Scientific Support:			
Terry Hewitt	0161 275 6095	w.t.hewitt@mcc.ac.uk	
CSAR Project Manager:			
Dr Victoria Pennington	0161 275 6830	v.pennington@mcc.ac.uk	
Head of CSAR Applications Support:			
Dr John Brooke	0161 275 6814	j.brooke@mcc.ac.uk	
Head of CSAR Optimization Support:			
Dr Mike Pettipher	0161 275 6063	m.pettipher@mcc.ac.uk	