# CSAR FOCUS

Edition 5 ,Summer 2000, Editor: Kaukab Jaffri, University of Manchester

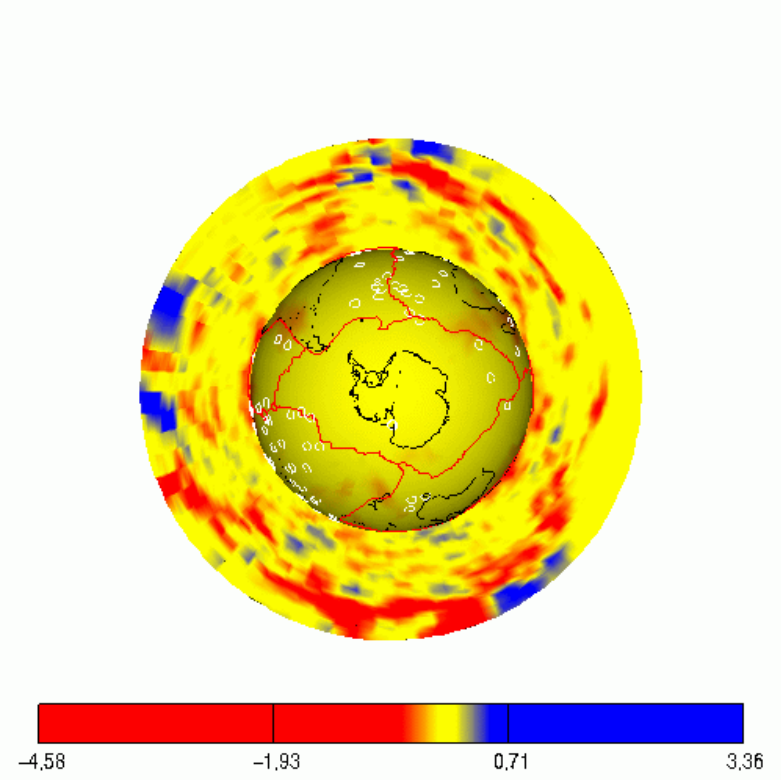


## Terra Consortium uses turing to simulate mantle convection currents

Images produced by CSAR using scientific visualization to show 3D seismic tomography of the Earth' s mantle.

# Content

# Editorial

So we're half way through the first year of the new millennium and, as summer approaches, so does MRCCS's eagerly awaited Summer School on Linux for HPC. For the second year running, the MRCCS Summer School is being held at the University of Manchester on the 6-17 September. The CSAR team look forward to meeting as many of you as possible, so get those applications rolling in - remember it's **free for all CSAR users**!

The beginning of the year brought a Royal visitor to Manchester Computing. HRH The Duke of York came to see the work of CSAR along with other work which has been done at Manchester Computing. I was asked to take part in this event and selected work from a few of CSAR's consortia, which formed part of the CSAR display.

HRH The Duke of York's visit to Manchester Computing earlier this year; being shown work done by CSAR consortia. From right: Keith Cole-Head of MIMAS; Kaukab Jaffri - CSAR frontline; The Lord Mayor of Manchester and HRH the Duke of York

This edition's frontpage features work by The Terra group, lead by Professor Huw Davies, from the University of Liverpool.

We also have an article by Dr Ben Jesson on the newly released CASTEP code as well as updates on other new software, in the **Software Update** feature.

The results of the 1999 CSAR User Survey which was carried out earlier in the year are also included in this edition.

As always, if you would like to contribute any articles to the next edition of CSAR Focus, please contact me.

**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@man.ac.uk
Web: www.csar.cfs.ac.uk

# Visualization Work Of The Earth's Mantle

Dr Huw Davies, University of Liverpool,
Joanna Leng, CSAR Applications Support

**In the UK it is still not common for computational scientists to use 3D visualization to examine their data. Many groups now seem to be at a bottle neck in their work because of this lack of graphics. When the UKHEC formed, it set a goal of improving the use of graphics and to do this they have commissioned three case studies.....we present the first one here.**

This work has been supported by SGI. An octane with 512 Mega Bytes memory, 16 GigaBytes of disk space and with texture memory has been borrowed from SGi, through the CSAR equipment loan scheme (see article on page .Url: www.csar.ac.uk/using/loan.shtml).

*Case Study*:
*Seismic Tomography and Convection Modelling of The Earth's Mantle by The Terra Group*

The UK base of the Terra consortium is at the University of Liverpool and it is led by Dr. Huw Davies. They run computational simulations of the circulation of the Earth's mantle (the layer between the crust and core) on 512 processors of the Cray T3E for up to 12 hours. These results are then compared to seismic tomography images derived by the group.

Currently they can only analyse the results by producing a series of 2D projections, one for each layer of data. All 2D projections are inherently distorted and hence may be misleading. 3D views of the data would be an improvement but it is important that they can be interactively cut, isosurfaced and rotated.

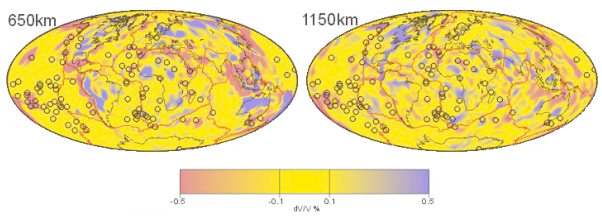


Figure 1: 2D projection of one depth layer - original visual analysis using seismic tomography data

We believe visualisation tools would speed up the analysis, shorten the time between runs and improve acceptance of the results from their peers. Papers have been rejected because of the inadequate nature of the visual presentation.

The mantle is a thick viscous liquid, usually modelled by finite element analysis, the ocean and atmosphere on the other hand are thin and flow easily between computational cells so are modelled by finite differences.

Data produced by finite difference is array data and relatively easy to visualize, but the data produced here was from finite element analysis which is cell data and much more complex to visualize. The simulation data is placed in computational cells in polar coordinate space with its inherent curved surfaces. Graphics renderers deal only with straight lines and are based on a cartesian coordinate system.

To convert the computational cells into visualization cells the coordinates must be converted from polar into cartesian coordinates and then each cell must be resampled to compensate for the curved nature of the original cells and make the new cell set tessellate. The data in the model is strongly related to other information, volcanic "hot spots", plate tectonic boundaries and coastlines. It is vital that this reference data is always visible and unambiguously displayed.

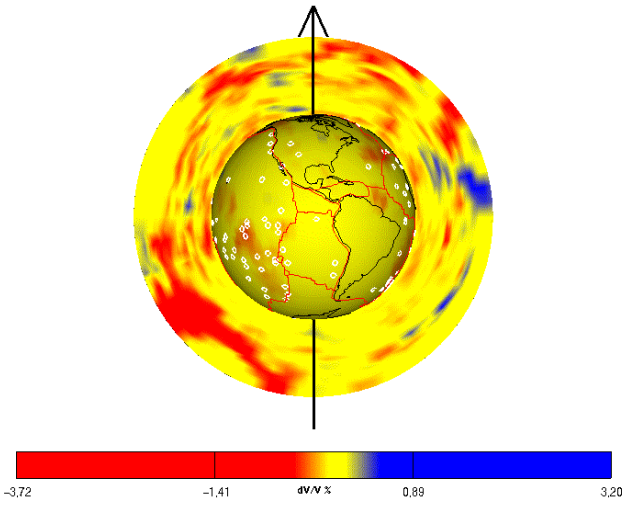


Figure 2: New 3D visualization of seismic tomography data.

Huw Davies has been pleased by the results of this visualization work. He works with researchers in the USA and they too are impressed, they have some 3D visualizations but not on data of this complexity or with this flexibility. The Terra group are keen to use the visualization tools on a permanent basis. The next case study will commence soon.

The simulation is high resolution and produces large data sets, which must be converted into its visualization cell set. The conversion resamples the cells and replaces each one with three to four cells. The size and structure make the resulting visualization difficult to display and slow to manipulate unless specialist graphics hardware is used on a machine with a large amount of memory. The Terra group have no specialist graphics hardware and they would prefer to spend any new computational resources on increasing the resolution of their simulation which ironically increases their visualization problems. Their expertise is in simulation not graphics and we are now considering the possibility of developing a visualization server as a part of the CSAR service.

These simulations are computationally expensive and the data is only analysed at the end of a run. If the computation goes down the "wrong" path resources are wasted. It would be useful if they could analyse results while the simulation is progressing and where appropriate alter parameters. This is called "computational steering" and with a visualization server could improve the system for the Terra group.

# Report from HPCN2000

Fumie Costen,
CSAR Applications Support,

**The 8th International Conference, High Performance Computing and Networking Europe 2000 had 4 Plenary sessions, 16 ordinary parallel sessions, 5 Workshops, 4 associated events, and poster session. The sessions focussed on Applications in Information Systems, Applications in Science and Engineering, Industrial User track, General End-User track, Computational Science track, Computer Science track.**

## CSAR Award at HPCN 2000

Fumie Costen of CSAR applications attended HPCN 2000 in Amsterdam (May 8-10) presentation at the conference with a paper titled "The problems and the solutions of the metacomputing experiment in SC99" by Stephen Pickles, Fumie Costen, John Brooke from Manchester Computing, and Edgar Gabriel, Matthias M.uller and Michael Resch from the High Performance Computing Center in Stuttgart, Germany and Stephen Ord from N.R.A.L., Jodrell Bank. The paper was awarded the "Best Conference Paper Award" and Fumie received the award on behalf of all the authors at the Conference Dinner. Here she describes her impressions of he conference as a whole. For details of the award-winning paper please see: www.science.uva.nl/events/HPCN2000/index.html#Award.

The content of the conference covered a wide range of computing and networking issues. The sessions, workshops and the associated events involved up to 6 parallel sessions. I was therefore not able to attend all the presentations. However I describe below some work which seemed to be of particular relevance and importance.

## PC Cluster Pilot Project

CSC in Finland reported on results of the one-year cluster pilot project (www.csc.com/metacomputer/pckluster).

They used RedHat Linux 5.2/6.0 with kernel 2.2.5, MPICH, and PBS using 32 processor PCs. The processors were Pentium II with a 400 MHz frequency, 128 MB central memory and 4GB local SCSI disk.

They found getting the systems to work properly and tuned, such as building programming environment, took the longest time among all the stages to carry out this project. Data communication bottlenecks turned out to be a major performance issue. They concluded PC clusters are not yet sufficiently mature as a good general purpose HPC platform. Most current usage of PC clusters is for specific applications and projects. The final report can be downloaded at www.csc.fi/reports/pc-cluster/Finalreport.pdf

## Message Passing Interface for Windows system

Universidade de Coimbra in Portugal have developed a Windows-based Message Passing Interface which provides a thread-safe, full implementation of the MPI standard for Windows operating systems, which they claim is considered the best freely available implementation. This

architecture frees the user code (which has dependence on both MPI management code and the communication access code) from core management activities and makes the library core totally independent of the communication media. The complete specification of Multiple Device Interface (a generic interface) can be found at the WMPI website:

http://dsg.dei.uc.pt/wmpi.

Another characteristic of WMPI is its ability to support an unlimited number of simultaneous MPI-devices which can be both threaded and non-threaded.

### Message Passing Library for Linux Systems

The University of Pisa in Italy has produced a compact, thread-safe communication library for cluster s of Linux OS computers. The system consists of 11 standard PCs with Intel Pentium II processors at 266 MHz with 128 MBytes of RAM and an IDE disk. These computers were interconnected by a fast Ethernet network using a 3Com SuperStackII switch. This library allows processes running on different processing elements of a cluster to share portions of memory.

This is the simplification of MPI which leads to efficient performance. They have shown experimental results that demonstrate that the library outperforms classical implementations of mpich(MPI-1) in collective operations and achieves comparable performance in point-to-point communications on Beowulf class workstation clusters. The comparison with MPI-2 implementations has not been done.

### Parallelization of codes

There were several presentations on finite-element methods and similar algorithms, from GMD-Institute for Algorithm and Scientific Computing in Germany, from other research Organizations for Information Science and Technology in Japan and from SGI in France. J. Clinckemaillie et al paid attention to improving the scalability of the algorithm, K. Garatani applied the earth simulator to wave propagation problems and paid attention to the parallelization of the code with the comment that the main problem of load balancing was imbalance in communication . Further information can be found at:

http://geofem.tokyo.rist.or.jp.

D.Nicolopoulos et al considered the Navier Strokes equations which can be used to derive acoustic behaviour of complex industrial systems. They parallelized the code on NUMA computer systems.

I found the conference very useful and interesting. However, due to the broad diversity of topics it was impossible for me to fully represent all aspects. In my report above I have concentrated on my own area, which is cluster computing.

# CSAR Applications/Optimisation News

Dr John Brooke,
CSAR Applications Team Leader

**Alongside our normal work of providing training, porting and maintaining software and helping users with their codes there have been several significant developments in the past few months. We describe two of these here, developments towards providing high-end computing via computational grids and the first steps towards what has been described as "visual supercomputing".**

This work is not exclusive to CSAR/MRCCS and it is good to report that a vigorous and exciting program of collaborative work is developing around initiatives launched by UKHEC. For more details of this work, and for a description of the role of UKHEC please see http://www.ukhec.ac.uk. Further, in collaboration with the UKCP Consortium and with staff at Daresbury Laboratory, a new version of the Molecular Simulations package, CASTEP, has been parallelised and released to the UK academic community - see the separate article about this by Ben Jesson for further details.

## Grid Computing

HPC staff at Manchester have been actively involved in the development of what has now come to be called GRID technology, right back to the days when it was described as metacomputing. We were particularly keen to exploit the possibilities opened up by establishment a fast ATM based metropolitan area network in the Greater Manchester area known as GMING. This was a revolutionary development as far as the UK was concerned and as the network has been so successful it has been extended to cover the north-west of England. The ATM technology allowed us to provide HPC and visualization for medical and industrial users via a variety of funded projects.

For details of these, please see the research WWW pages at http://www.man.ac.uk/MVC/research.

Two projects that are particularly relevant to GRID computing are Novice (for medical applications) and RCNET (for engineering). Based on this experience we find the extension to much wider area applications to be timely. An important advance that the concept of a GRID has brought to metacomputing, is the realisation that for ease of use there needs to be infrastructure and middleware to enable the distribution and running of work around the wide area networks a routine rather than a "heroic" task. However, all available evidence indicates that even in the US this situation is some way off. It is not self-evident that computation and information can be treated in the same way as electrical power (the inspiration for the term GRID comes from the development of a national GRID for electricity). Some well-founded demonstration testbeds where applications could be trialed would be very helpful for the computational community. The particle physics community is actively developing a GRID to handle data from the Large Hadron Collider at CERN. It remains to be seen how well the structure of a GRID based on data handling would work for GRIDs based on other requirements, such as large scale computational modelling or integration of virtual reality environments with simulations.

Martyn Foster describes elswhere in this issue how GRID experiments are being developed on a UK-wide scale. On a European scale, it is exciting to report that MRCCS will be a partner in the EUROGRID project that will be funded to establish a European-wide GRID for major large-scale experiments. The CSAR facilities will be part of this

GRID and we will be working to develop applications that can exploit it. The main application areas targeted for development are modelling of biological molecules, climate and weather simulation, finite-element engineering applications and large-scale experiments between major HPC centres. These include processing experimental data being streamed in from major installations (such as the Jodrell Bank data featured in the SC99 experiments described in CSAR Focus Edition 4).

### Visual Supercomputing

Visual supercomputing involves integrating best practice in visualization and virtual reality with large scale computation. It can be used to enable researchers to extract information and insight from very large or complex datasets produced by simulations. It can also be used to monitor or even steer computations so that the insight of the researcher can be actively involved in the modelling process. CSAR specialists have worked with users in both of these areas. Firstly Jo Leng has been working with Huw Davies of the Terra group (involved with simulating convection in the earth's mantle, see article page ) to develop AVS modules suitable for working in spherical geometry. This has applications to both geophysics and astrophysics (since planets and stars are approximately spherical) and the techniques have not been widely explored so far. This work will be made available to other interested group via UKHEC publications and technical reports. Secondly we have assisted Professor Nick Avis group from the University of Salford in staging a demonstration of computational steering on a supercomputer. MRCCS also hosted a very interesting seminar by Peter Love from the University of Oxford describing work done on computational steering in mixed fluids and molecular dynamics. Since MRCCS and CSAR share office accommodation and have very close links with MVC (Manchester Visualization Centre), we will be able to explore the possibilities for visual supercomputing offered by the new Virtual Reality Centre to be built at MVC in May/June 2000. We will report on initial experiments in a future edition of CSAR focus. Close collaboration with SGI as partners in CfS is proving very helpful in developing the visual aspects of supercomputing. In particular, the Terra group is currently benefiting from the loan of a specialist graphics workstation from SGI.

We are interested in hearing from any CSAR users who would like to discuss working with us to adapt their applications to exploit the new technologies described above. Please contact the CSAR Frontline in the first instance.

### MRCCS

The second series of HPC/Visualisation seminars organised by MRCCS were completed just before Easter. This has been a very successful weekly series of seminars, covering a broad range of topics relevant for both HPC and Visualisation. There have been speakers from UK HPC service providers (CSAR/MRCCS, Daresbury Laboratory and EPCC), academic institutions including Minnesota State University in the US, and commercial organisations such as Nag. There will be occasional seminars during the summer, with the main programme starting again in the autumn.

If you have a request for a particular topic or speaker, or if you would like to present one yourself, please contact us. For further details, please see: http://www.man.ac.uk/mrccs/seminars

Forthcoming events organised by MRCCS are the Linux for HPC Summer School and the 6th Cray/SGI MPP workshop, both in September this year, both described elsewhere in this newsletter.

# CASTEP: Quantum Mechanical Atomistic Simulation Code

Dr Ben Jesson,
CSAR Optimisation Support

**May 2000 saw the release to the UK academic community of version 4.2 of the quantum mechanical atomistic simulation computer code, "CASTEP". The release of this code, which has been keenly awaited by many CSAR users, is the result of an extensive collaboration, over several months, between the academic developers of the code, the Daresbury HPCI centre, and the CSAR applications and optimisation support teams.**

"CASTEP" (CAmbridge Serial Total Energy Package), originally written in the 1980's by Mike Payne at Cambridge University, is now developed by the UK Car-Parrinello (UKCP) consortium of academic research groups (1) and a commercial partner, Molecular Simulation Inc (MSI) (2). MSI markets and sells the code to commercial users world-wide, but under the terms of a UKCP-MSI agreement, CASTEP is available for free to all UK universities (CSAR users wanting to use CASTEP should contact Dr. Phil Lindan at the Daresbury Laboratory (3)).

For a fuller description of the purpose of CASTEP, its functionality and the science that underpins it, the reader is referred to the web pages referenced in the footnotes, and the publications listed therein.

In essence, however, CASTEP uses density functional theory (DFT) (specifically, using plane-waves and pseudopotentials) to solve approximately the Schrödinger equation for periodic systems of atoms, yielding the total energy, atomic forces and internal stresses in the system, as well as interesting electronic properties (the electron wavefunction, charge density distribution, density of electronic states, etc).

## CSAR's role

The development and optimisation support provided by CSAR for CASTEP was carried out by Ben Jesson and Stephen Pickles from the CSAR team, working in conjunction with Rob Allan and Ian Bush of the Daresbury HPCI centre, who have previously worked substantially on the CASTEP code.

The first task of this team was the parallelisation of some key parts of the CASTEP code. CASTEP is a large and complex piece of software, consisting of approximately 120,000 lines of Fortran 77/90. It has been developed over more than a decade by a number of people, and during that time much of the older parts of the code have been parallelised using the Message Passing Interface (MPI) library. However, some of the most important new functionality of the previous CASTEP release, version 3.9 (including the 'density mixing' electronic minimiser and the ability to use 'ultrasoft' pseudopotentials), had not been parallelised, dramatically limiting the utility of CASTEP for CSAR users.

Further, even when this parallelisation work was complete, it was still not immediately possible to perform the very large CASTEP calculations that members of the UKCP were hoping to run on the CSAR T3E. This was due to the memory usage of some CASTEP parallel algorithms, which was such that, regardless of the number of processors, the memory required on each one would always exceed that available (ie 256 Mb on the CSAR T3E).

The second phase of CSAR's development work was therefore to address this problem, specifically for the Cray T3E (the other CSAR computers, an Origin2000 and VPP300, have more memory available to each processor and so do not suffer so acutely from these problems).

Finally, the substantial and in-depth knowledge of the CASTEP code that has been gained through this development work has been applied to the optimisation of the code, particularly aimed at the Cray T3E.

### Results

It is rather hard to illustrate the results of the initial parallelisation work performed by CSAR since, for example, the precise parallel speed-ups obtained will depend strongly on the nature of the calculation performed and, in any case, will be affected by the ongoing optimisation work. Nevertheless, good parallel performance was obtained for the test case calculations used in the work.

The improvement in the memory usage of the parallel code was acheived through the implementation of a number of alternative parallel algorithms for which the memory requirement could be scaled down by using more processors. In the case of some of CASTEP's file input and output routines, this necessitated the use of certain Cray-specific I/O routines. The results of this work are illustrated for one medium-sized test case, a 37-atom simulation of an aluminium impurity in silica, in figure 1.

As can be seen, the memory usage per processor, even when many processors are used, tends to a particular lower limit below which it cannot further be reduced. For large systems, this lower limit was greater than the 256 Mbytes available on the T3E.

However, following the implementation of alternative parallel algorithms, the limiting memory usage has been substantially reduced. For larger systems it is likely that this reduction will be even more dramatic than that shown here, and indeed much larger CASTEP calculations can now be performed on the Cray T3E than were hitherto possible.

The development work already described has been aimed primarily at getting the latest CASTEP functionality working on the T3E for large calculations. Now that this has largely been achieved, we are beginning to turn our attention to the optimisation of the code, so that these calculations can be performed as efficiently as possible. This task is complicated by the fact that optimisation needs of the code depend strongly on the nature of the calculation being performed. However, some initial optimisation results have been obtained, which can be expected to make some significant improvement to the performance of the code for a wide range of calculations.

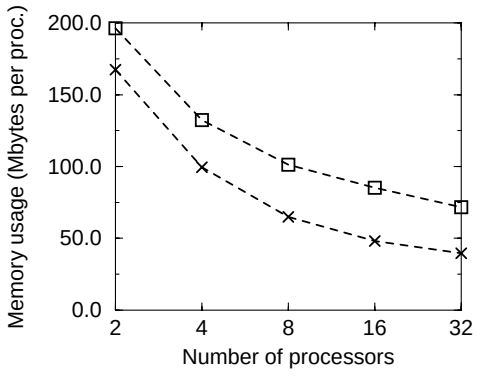


Figure 1 : Per-processor memory usage for a typical medium-sized CASTEP calculation. Squares and crosses show memory usage before and after the CSAR work.

These particular optimisations focus on the implementation of the Fast Fourier Transform (FFT) algorithm in the code, which for many calculations can account for a very significant part of the total execution time. CASTEP uses a three-dimensional FFT on a distributed data grid, which involves three sets of non-distributed one-dimensional FFTs, interspersed with intra-processor and inter-processor data copy operations.

Various optimisations have been applied to this part of the code, the results of which are indicated in figure 2. This plot shows CPU times per FFT, for a range of grid sizes, before and after optimisation. The squares indicate timings taken from the original unoptimised routine, and the crosses show the result of applying a rather general cache-based optimisation to the 1-D FFT part of the code, which should show a s similar benefit an all cache-based computer architectures (including, for example, an SGI Origin2000 such as fermat). The diamonds show timings obtained using a Cray-specific library routine for the 1-D FFT, and the triangles show the effect of also using a Cray-specific routine for the intra-processor data copy operations. These results illustrate rather well the general principle that, while some generic optimisations can make a very significant difference to the performance of a piece of code, often vendor-specific library routines can provide much greater performance for relatively little programmer effort.

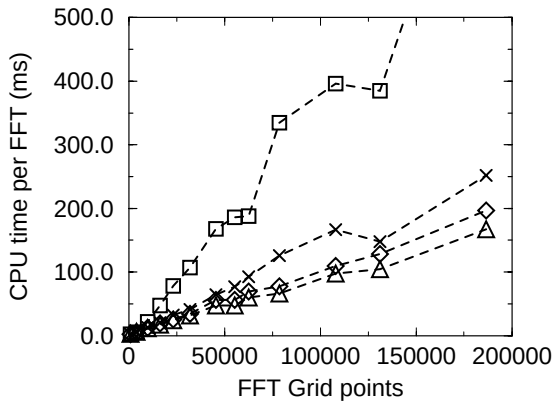


Figure 2 : CPU times per FFT as a function of grid size, before and after various optimisations as described in the text.

The timings shown in figure 2 focus exclusively on the FFT routine, and therefore in a real example the improvement in performance may not be quite as dramatic as it suggests. However, for the medium-sized calculation described above, the overall execution time of the code on 16 processors drops from 312 seconds with the original FFT to just 183 seconds when using the Cray-specific routines, a 70% improvement in performance!

## Conclusions

The parallelisation and optimisation support provided by CSAR to the UKCP consortium in connection with the CASTEP code is clearly of an unusually extensive and in-depth nature. However, it serves as a good illustration of the enormous benefits that can accrue from such a relationship. For example, although CSAR's support for CASTEP has continued over a number of months, the optimisation work described above took place over only a small fraction of that time, and it is not unreasonable to expect that other codes could similarly benefit from the investment of just a few support tokens!

## Footnotes

(1) http://www.cse.clrc.ac.uk/Activity/UKCP

(2) http://www.msi.com/materials/cerius2/castep.html

(3) e-mail: ukcp@dl.ac.uk

# Eat Less Chips Consume More Serial

Dr Keith Taylor,
CSAR Senior Applications Support

**The T3E is a powerful machine. The 1200E at Manchester, for example, has 784 parallel application (APP) nodes, each offering a peak performance of 1200 MFlops/s. Thus, in principle, a single user has access to nearly 1 Tflops/s of computing power. A mouth-watering prospect!**

As we all know, however, your code must first be parallelized using message passing, because the T3E is a distributed memory architecture. HPF will take the drudgery out of this process, but, unfortunately, as at Manchester, it may not be available on your system. So we resort to MPI, say. Some time later, we punch the air. We're producing correct results, acceptably close to those we previously obtained from the serial code. Having paid close attention to load balancing, we can also demonstrate reasonable scalability. Good. Now, we can run much bigger cases, much faster.

Hang on, though. What performance are we actually achieving? Oh dear, less than 10% of peak, not uncommon for practical applications. So, what can we do to improve matters. Well, we've heard of SHMEM, the data passing SHared MEMory access library, and how efficient it is compared to message passing. It's Cray specific, of course. (In fact, SHMEM now also exists on the Compaq AlphaServer.) So we'll lose portability. But, since we don't envisage changing machines in the near future (God forbid!), we'll try it, making sure to encapsulate the communications so that we can swap back, if necessary, without too much pain.

Still no noticeable improvement. OK. Our application is a time-stepping Computational Fluid Dynamics code which (hopefully) converges to a steady state. Upto now, we've been careful to synchronize after each time step. Couldn't we perhaps remove the synchronization, and let each processing element (PE) run free, accepting and sending 'halo' data when it feels like it, thus squeezing the last drop out of load balancing. Fine. Iterations on each PE are now running faster, but we're having to run more iterations, thus cancelling out the fine effort we have put in. Why?

Well, it is my experience that iterative codes have a natural evolutionary path, and the more we deviate from it, the longer the code will take to converge. Basically, at each step, there exist two solutions, one running forwards in time which is stable, and the other running backwards which is unstable. The art of convergence is to suppress the latter in favour of the former. (Think of generating Bessel functions of various orders, using recursion. If you want to compute the Js you go downwards in order; otherwise the Ys dominate.) This explains why the Multi-Grid method, for example, displays a law of diminishing returns. Multi-Grid attempts to

alleviate the limitation on size of time step, imposed by the CFL condition, by solving for physical variable differentials on a system of coarser grids. However, you're only evolving the original equations for the physical variables themselves on the finest grid where you want the solution. Thus, after a certain level of coarseness, 5 say, you obtain no advantage, and the extra work entailed in interpolating and extrapolating between grids, as well as the increasing number of iterations required, more than cancels out the initial gains.

So, what's left. Well, have you looked at the performance you're getting on each PE? No. It's not uncommon. In the heat of the battle to get a parallel code running reasonably, serial performance is often neglected. Let me tell you a story.

CSAR recently ran a scheme whereby users of the Manchester machines (we also have a 16-node Origin2000 and an 8-node Fujitsu VPP-300) were provided with 3 days of free advice with the aim of assessing codes and identifying where their performance could be improved. I was involved in such an exercise. The user had nicely encapsulated the problem: one subroutine, which was called many times, was responsible for almost all of the runtime, over 2000 seconds in the case I was given.

Before doing anything fancy, like trying out different combinations of compiler options and directives, of which there are a myriad available, or thinking about how cache use might be optimized, I decided to take a closer look at the subroutine source. After all, it was only about 80 lines long. And what did I find. There were 7 do loops, two of which accessed elements of a 3-dimensional array across its slowest moving index.

Now, the T3E PE just loves accessing long runs of contiguous values from memory, in sequence. Deviate from this by introducing a stride, for example, and you pay a heavy penalty. Thus,the answer was obvious: permute the order of indices so that the last became the first. But, this could have wide ranging implications throughout the rest of the code. So I decided, in the interests of minimizing the code changes and containing them within the offending subroutine, to extract the required elements first into temporary 1-dimensional arrays. This was carried out by using a SHMEM procedure. (SHMEM is not only capable of performing data transfers between different PEs, but can carry out copying within 1 element's memory, very efficiently.) The result: runtime was halved!

Now, if the last paragraph made little sense to you, what should you do. Well, you could attend a course. For example, CSAR runs "Porting/ Optimising for the Cray T3E-1200E", lasting 3 days, which provides a thorough explanation of the T3E architecture, enabling you to build up a sound programming model of the machine, in order to get the best out of it. More importantly, though, you should avail yourself of the vast amount of wide ranging experience, in developing and running large scientific applications, offered by CSAR's members. You do the science; let us help you to exploit the computer resourcesefficiently.

Above all, however, improve your diet!

# Plug into EScience

Dr Martyn Foster,
CSAR Applications Support

**"If electricity came in boxes..." starts the commercial, portraying a strange view of the world in which our sustaining resource is delivered in a rather inconvenient format. We are reminded of the transparency of something that we use daily.**

In the last procurement round for computational services, EPSRC settled on the technologies provided by Cray Research Inc. - an MPP Class machine with an exceptionally high quality network and a multitude of workstation class processors. This machine maps well onto many scientific applications and delivers a rather significant chunk of computing resource for the UK academic community.

There will have been those that were less than happy with this decision, perhaps the per node memory was not quite as great as their applications would have liked, or their code had been developed for multi threading environments, vector machines, or tuned for some other CPU or environment. Their options in this country do not extend far with the other major UK computing resource also being a T3E, and the question must beg, why can't they trade those CSAR tokens/ research grant for something that meets their needs exactly, something along the lines of the CSAR guest services, but more flexible, more general, and without the pain of new accounts, administration, and alien environments. Perhaps next year they can...

What stops you from running somewhere else today? At a rather fundamental level it is your system administrator, for some reason, s/he simply refuses to let you use the Berkeley r-tools to fetch the needed files across the network (mutters something about security when asked), s/he won't tell you when your batch job would run locally so that you could make that important decision as to whether you want run it somewhere else, and s/he won't sort out the dedicated bandwidth you need in order ship out your data to the remote site in a sensible amount of time. All of that is before you even consider the other site. But, before you fire up your email client, remember this - s/he doesn't have the tools at his disposal to fix these things - not without a great deal of time consuming manual intervention. There are two features described above that are key to the development of transparent access to resources. The first is the insidious fact that all these computing resources are run by different individuals – a problem of heterogeneity of both local policies and machines. The second is one of quality of service.

A grid solution to computing is one that allows individual sites their quirks whilst providing a homogenous interface to the machine. This problem is being addressed by a few research groups and is unfolding into usable systems, notable players are Globus (http://www.globus.org), Unicore (http://www.fz-juelich.de/unicore/) and Legion (http://www.cs.virginia.edu/~legion). Participating sites run this software, tune the local components to their systems and advertise their services within the virtual machine. The existing test-beds are already quite large – GUSTO, the Globus test-bed has well over 100 HPC sites worldwide, including both Manchester and CLRC. Within these environments many of the problems associated with 'borders' disappear, files can be moved around, machines advertise their capabilities and limitations, and passwords/rhost files seem a thing of the past. Is the future here already?

Quality of service, alas, is not so well developed. To be able to make the right decision either explicitly or via some intelligent application (called a broker), one must be able to constrain all of the variables right from your terminal through to the display hardware that lets you view the results – every network link, every batch system, every disk that we will need space on. Would I have the bottle to walk into the conference hall depending on my steerable simulation beginning to start on the Cray to coincide with the slot for my talk? This technology is for tomorrow, but undoubtedly it will arrive.

And what does this mean for science? For current applications it may mean faster turnaround times, the resource you wanted (rather than that which was provided), the cheapest solution or the one that one that can work to a deadline. In the longer term there will be a new breed of grid aware applications – they can access files anywhere on the grid, they will span machines (and adjust themselves based on the quality of the networks joining them) and allow telescopes to talk to spacecraft in real time (with a little interpretation from the Cray). All of this with an ease of use and interface that means you may never have to leave your workstation and can use your own applications to edit and manage files anywhere on the grid. Contracts are brokered by software automatically and the costs translated and deducted from your resource allocation automatically.

CfS is currently investigating both Unicore and the Globus Metacomputing Toolkit. If you would like more information or are interested in participating in the experimental grid services on the CSAR machines please email csar-advice@cfs.ac.uk.

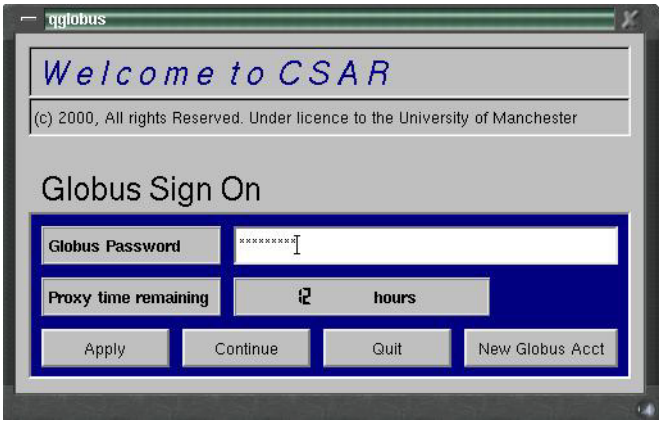


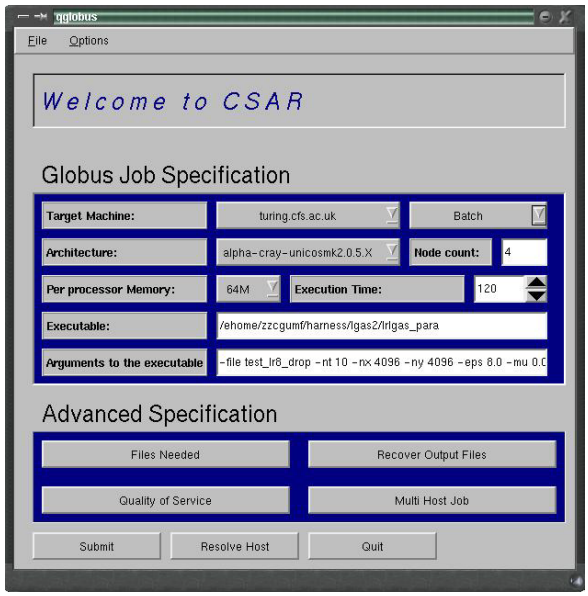

# CSAR Software Update

Kevin Roy,
CSAR Optimisation
Support

## Latest Nag Fortran 77 Libraries

The latest numerical libraries from Nag have been installed on our systems.

The latest release, Mark 19, contains more than 60 new routines that extend and improve the functionality available in the areas of fast Fourier transforms, optimization, eigen value problems, sparse linear algebra, statistics and operations research.

The addition of a sparse nonlinear programming routine to the existing Optimization chapter will be of particular interest to people wishing to solve large scale optimization problems, such as those arising in financial, engineering and aerospace applications.

The Operations Research chapter has been broadened to include solvers for Quadratic Programming problems where selected variables are constrained to be integer.

The rapidly expanding Sparse Linear Algebra chapter now contains iterative methods and preconditioners for complex symmetric and non-Hermitian linear systems of equations, which are particularly suitable for the solution of PDE problems.

Information on using these new libraries can be found from the machine specific links at the top of the web page:

http://www.csar.cfs.ac.uk/software/nag

Also note that the documentation on individual routines is available from NAg (http://www.nag.co.uk) and from Manchester Computing's Applications' website:

http://applications.mcc.ac.uk/info/nag/nag.asp

and select "Mark 19 FORTRAN 77 library on-line manual" from the menu.

# CSAR Software Update

Dr Robin Pinning,
HPC Support

## Latest Chemistry Software Developments

What follows is a short description on the latest developments in Chemistry software available on Turing. The last few weeks has seen the followng developments in the provision of Chemistry software on CSAR's Cray T3E-1200E.

### DL_POLY 2.12

DL_POLY is a parallel molecular dynamics simulation package developed at Daresbury Laboratory by Bill Smith and T.R. Forester for the EPSRC funded Collaborative Computational Project for the Computer Simulation of Condensed Phases (CCP5) and the Molecular Simulation Group (MSG) at Daresbury Laboratory. The package is the property of the Central Laboratory of the Research Councils (CLRC).

Due to our close links with Daresbury Laboratories central copies of all the CLRC Chemistry packages that are suitable for the Cray architecture will be available soon. DL_POLY was the first of these packages to be centrally maintained and installed on the machine and is available to any academic user with an academic license to use the package. The advantage this gives the user is that any bug fixes or updates to the package are implemented as soon as they are available.

For more information see http://www.dl.ac.uk/TCS/Software/DL_POLY/main.html.

### Gaussian 98

After many months of waiting Gaussian 98 is finally available for the Cray T3E.

For the reader that doesn't know what the Gaussian package does:

"Gaussian 98 is the latest in the Gaussian series of electronic structure programs. Designed to model a broad range of molecular systems under a variety of conditions, it performs its computations starting from the basic laws of quantum mechanics. Gaussian 98 is used by chemists, physicists and engineers for research in established and emerging areas of chemical interest, studying molecules and reactions of definite or potential interest, including both stable species and compounds which are difficult or impossible to observe experimentally: short-lived intermediates, transition structures and the like."

"Gaussian 98 can predict energies, molecular structures, vibrational frequencies-along with the numerous molecular properties that are dervied from these three basic computation types-for systems in the gas phase and in solution, and it can model them in both their ground state and excited states. Chemists apply these fundamental results to their own investigations, using Gaussian 98 to explore chemical phenomena like substituent effects, reaction mechanisms, and electronic transitions."

Gaussian 98 contains many improvements and addition functionality over Gaussian 94. These improvements extend the package's ability to model large molecular systems, allow for calculation of a greater range of molecular properties and increase the code's capacity for calculation of excited states.

For more information on Gaussian 98 on the Cray T3E see Carlos Sosa's site at http://outside.cray.com/~cpsosa/.

### NWChem 3.3.1

The NWChem 3.3.1 software package has recently been installed on Turing.

This is the first time this package has been made available on the machine and extends CSAR's already impressive selection of Chemistry packages.

"NWChem is a computational chemistry package that is designed to run on high-performance parallel supercomputers as well as conventional workstation clusters. It aims to be scalable both in its ability to treat large problemsefficiently, and in its usage of available parallel computing resources.

NWChem has been developed by the High-performance Computational Chemistry group of the Environmental Molecular Sciences Laboratory (EMSL) at the Pacific Northwest National Laboratory (PNNL). Most of the implementation has been funded by the EMSL Construction Project."

For further information and details of all the tools used in NWChem please see http://www.emsl.pnl.gov:2080/docs/nwchem/.

# People Focus

**SInce the beginning of the new year, CSAR has had two additions to the team. John Brookes, (not to be confused with CSAR Applications Team Leader John Brooke!) joined the optimisation team in March, and Ali Abdul Rashid who joined a month earlier, replacing Desirae Hill on the Frontline team.**

## Dr John Brookes

During my PhD I was lucky enough to have class 3 access to kilburn (SG Origin 2000), which I used to perform Monte Carlo simulations of extensive air showers. Having been impressed with the staff, facilities and general atmosphere of MC - and having planned to return to Manchester anyway - I was pleased to learn of an opening here. I started work on the 1st of March 2000.

My main responsibility is to aid the parallelisation of finite element codes for the Departments of Engineering and (Applied) Mathematics. The aim is to produce an optimised library of Fortran90/MPI subroutines to ease F.E. program development within those departments. I also have an interest in the soon-to-be-upgraded davinci machine: http://davinci.man.ac.uk/.

## Career Summary

- 1992-96 - II(i) BSc(Hons) Physics with Astrophysics at the University of Leeds
- 1994-95 - Exchange year at Universität Ulm, Germany
- Mar/Apr 1995 - 'Wissenschaftliche Hilfskraft' (research asst.) at Uni-Ulm
- 1996 -to date - PhD in Cosmic Ray Physics at the University of Nottingham
- Mar. 2000 - to date - Parallelisation support for MRCCS at the University of Manchester

## Interests

As you might be able to tell from the picture, I'm quite into juggling and the typical Manc pastimes of supporting Manchester United (I am from Salford, before you say anything :-), shaving my head and growing goatees. I participate (or have done, with the intention to resume) in several sports, most notably karate, football and swimming. I play the violin occasionally (and badly) and have been known to give the odd drum a slap. I also enjoy going to clubs, doing dangerous things (eg jumping out of planes, being a United fan in Leeds etc.) and littering conversations with unfeasibly weak puns.

## Affiliations

- Graduate member of the Institute of Physics
- Member of Mensa

John Brookes juggling

## Abdul Ali Rashid

I originate from Pakistan, but now live in Manchester. My friends and colleagues know me as Ali. Before joining the CSAR Frontline team I was working part time for Manchester Computing on a PC support helpdesk.

I achieved a B.Sc. in mathematics and physics from the University of Punjab, Pakistan. My initial Masters degree was in applied mathematics at the Quaid-i-Azam University, Pakistan.

I joined the Computer Science Department at the University of Manchester to pursue a Masters degree in Computer Science. I completed this in 1997.

Most people are relieved to hear that I do not solely live for Computing! My hobbies include watching and attempting to play cricket and football. I enjoy voluntary and social work. I have been working for organisations such as the Blood Donor society. I also enjoy travelling.

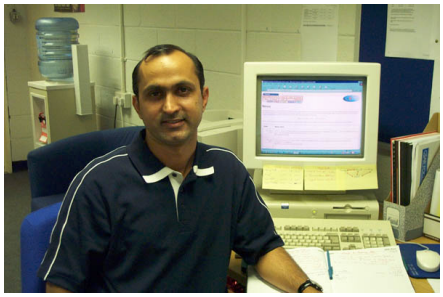

Ali in the CSAR Frontline office

# Loan of High Performance Graphics Workstations

## Would your project benefit from the loan of an SGI Workstation?

CfS has a pool of equipment for temporary use by research groups funded to use CSAR, for the following main purposes:

- To foster collaboration in key areas
- To enable research groups to accelerate research ahead of some deadline (e.g., conferences) when the post-processing phase becomes a bottleneck because the equipment available to them can not cope with the temporary peak.
- For research groups whose productivity could be enhanced through the development of new post-processing and/or code development tools. This facility will enable users to evaluate new methods of working very cost-effectively.
- To support training and education programmes.

### Equipment available

Workstations available under the loan program are subject to change, to ensure that they are kept up to date. At present, the loan pool consists of the following machines:

- 3 x Silicon Graphics 320 visual workstations for Windows NT. Featuring the Integrated Visual Computing architecture, they deliver unprecedented graphics performance with Windows NT applications.
- 3 x Silicon Graphics Octane visual workstations for IRIX. The Octane gives you the processing power and visualisation needed to develop innovative solutions by integrating tasks, combining steps and shortening the time needed to achieve your goals.

### Conditions of Loan

- **Hardware**

CfS will be responsible for delivery, installation, collection and maintenance of the equipment. You will be required to insure the equipment against theft and accidental damage, either by paying a small fee to CfS or by producing evidence that you have the equipment insured.

- **Software**

CfS will ensure that you have a license and the basic software installed. Any third party software needed you will be responsible for the license, and installation, though we will do our best to help you.

- **Report**

When you have finished with the workstation we would like a report on the work you have been able to complete, and a copy of any publications that you generate, and your help in creating some publicity material.

# Loan of High Performance Graphics Workstations

## Loan period

Systems loans will normally be limited to a period of 12 weeks, or up to 16 weeks in exceptional circumstances.

## How to Apply

If you would like to loan some equipment, then please contact the CSAR helpdesk. Please have the following information available:

- Your contact details
- What equipment you would like, and where it is to be based, and how long you will need it for (up to a maximum of 12 weeks)
- A description of the work you will undertake with the equipment

We will consider your request and let you know within two weeks.
For further information, see http://www.csar.cfs.ac.uk/using/loan.shtml.

# Forthcoming Events

**July 2000**

9-21

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

**Aug 2000**

1-4

The Ninth IEEE International Symposium on High Performance Distributed Computing (HPDC-9), Westin William Penn, Pittsburgh, Pennsylvania - http://www.cs.cmu.edu/~cmcl/hpdc2000/

16-18

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

18-20

International Workshop on Parallel Matrix Algorithms and Applications, Neuchatel, Switzerland - http://iiun.unine.ch/Research/matrix/seminars/pmaa.html

**Sept 2000**

7-8

6th European SGI/Cray MPP Workshop, Manchester Computing, University of Manchester, Manchester, UK - http://www.man.ac.uk/mrccs/mpp-workshop6/

4-15

High Performance Computing with Linux. 2nd Annual MRCCS HPC Summer School, Manchester Computing, University of Manchester, UK - http://www.man.ac.uk/mrccs/summer_school/2000.shtml

10-14

EuroPVM/MPI2000, Balatonfured, Lake Balaton, Hungary. Final date for submissions: 26th April 2000 - http://www.lpds.sztaki.hu/EuroPVM_MPI2000/

14-15

Second European Workshop on OpenMP: EWOMP 2000, University of Edinburgh - http://www.epcc.ed.ac.uk/ewomp2000

**Oct 2000**

11-14

SGI2000, Krakow, Poland - http://www.cyf-kr.edu.pl/sgi2k/


# The Sixth European SGI/Cray MPP Workshop

## Manchester, UK

## 7-8 September 2000

ANNOUNCEMENT

The Sixth European SGI/Cray MPP Workshop will be held in Manchester, UK, on the 7th and 8th of September, 2000. The workshop is organised by the Manchester Research Centre for Computational Science (MRCCS), a focus for high performance computing activities at the University of Manchester, and is generously supported by Silicon Graphics, Inc. and Cray, Inc.

There is no registration fee for this workshop.

Important Dates:

9th June (extended) - Deadline for receipt of extended abstracts
23rd June - Notification of acceptance of papers.
11th August - Deadline for receipt of final papers.
25rd August - Deadline for registration of delegates.
7-8 September - Sixth European SGI/Cray MPP Workshop.

This event is distinguished from national Cray User Group meetings by a strong emphasis on applications. It is hoped toinclude talks on MPP topics such as:

- Programming models and methods;
- Parallel performance prediction, evaluation and optimisation;
- Parallel numerical algorithms;
- Application of 3rd-party parallel solvers, libraries and support tools;
- Data management and visualisation;
- MPP production applications (in computational physics, chemistry, engineering, biology, meteo-climatology, earth sciences etc.).

For further information, please see: http://www.man.ac.uk/mrccs/mpp-workshop6/

The Organising Committee can be contacted c/o Dr. Ben Jesson, Manchester Computing, University of Manchester, Manchester M13 9PL, UK (or e-mail b.jesson@man.ac.uk).

# 1999 CSAR User Survey Results

Dr Victoria Pennington,
CSAR Project Manager

The first annual CSAR User Survey was conducted in December 1999. There were 13 questions covering a range of aspects of the service, including systems, helpdesk and support services, training, feedback mechanisms... 60 completed forms were received, representing around 20% of active Class 1/2/3 users. This was somewhat higher than anticipated, based on the levels of feedback via other mechanisms - the recent email from the User Liaison Forum chair (Dr Matthew Foulkes) inviting comments/complaints/suggestions received just 2 replies. Survey submissions were entirely anonymous, although users were given the opportunity to give their name on the form - 27 people did this.

The response showed a good level of satisfaction with the service, with no major issues coming to light. A full report was published in February and is available on the web at http://www.csar.cfs.ac.uk/admin/reports/user_surveys/.

Here we highlight just a few of the results.

Figure 1 overleaf, shows the satisfaction with various systems aspects (all three CSAR systems). For all aspects, over 84% were *very satisfied* or *fairly satisfied*. The most dissatisfaction was with job time limits and the archive (HSM/tape) facility. However, the majority (84%) are satisfied (*very* or *fairly*) with job time limits. It is always possible to run occasional longer jobs by arrangement via the Helpdesk, but an increase in the default job time limit would almost certainly have a detrimental effect on job turnaround for the majority. The archive facility has seen significant improvements over recent months, with new file management software (TMF) installed at the end of last year, and new hardware recently installed. Provision for interactive use has been an issue with some, and in response we have implemented a fully automatic batch job migration mechanism to assist interactive users during core hours.

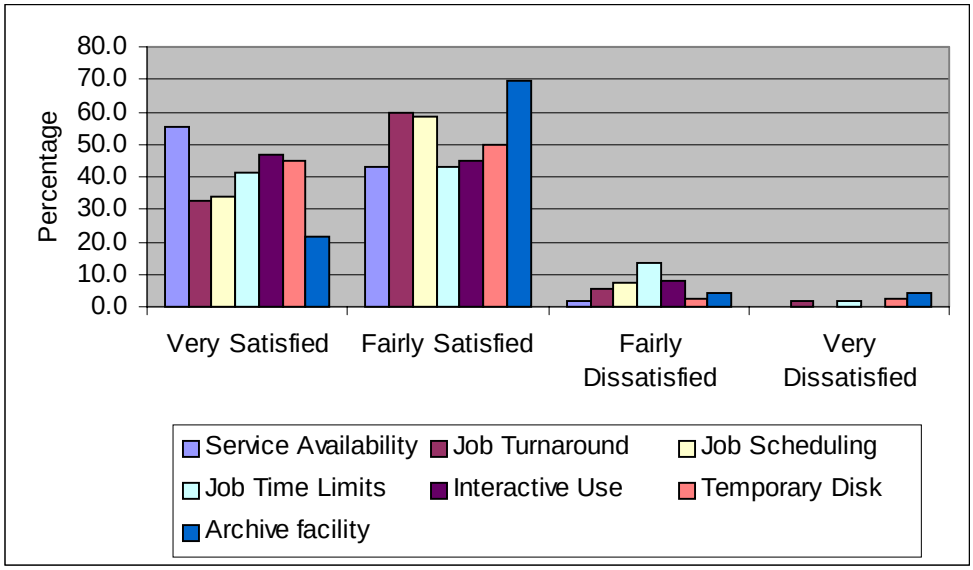


Figure 1: Satisfaction with various systems aspects (as a percentage of those expressing a view).

Users were asked for their view on the overall level of HPC Service provided by CSAR, ranging from *very good* to *unacceptable*.

Figure 2 shows the number giving each of the five responses. In summary, over 88% replied in the top two categories - *very good* or *good*.

A big thank you to all users who contributed to the 1999 survey. The next one will be in December 2000. In the meantime, please feel free to contact us on any issue - we rely on your feedback to ensure we are providing the best possible service.

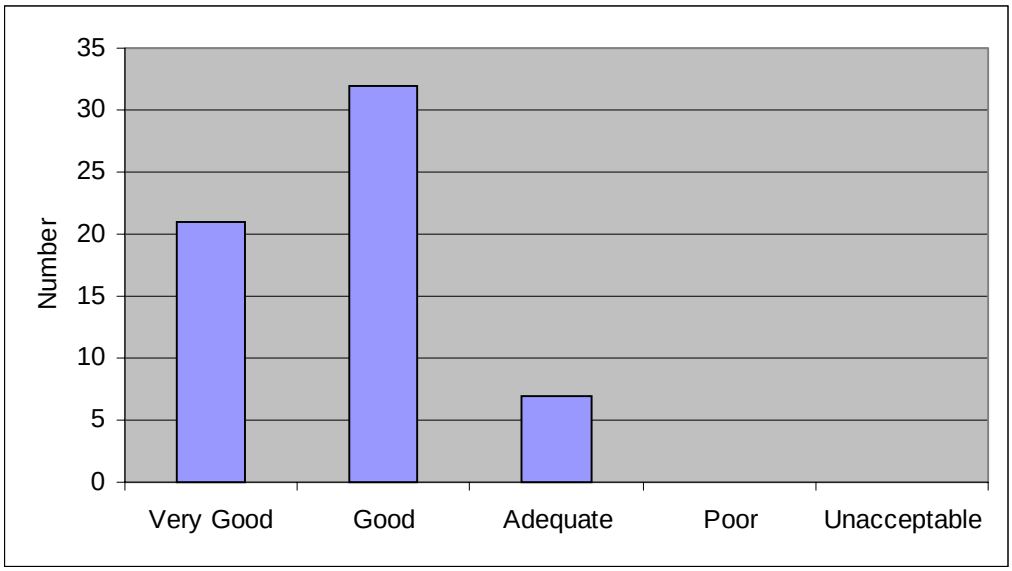


Figure 2 caption - View on the CSAR Service overall.


## Announcement :

# High Performance Computing Summer School on Linux for HPC

## 4th - 15th September 2000, Manchester

Manchester Computing and Manchester Research Centre for Computational Science (MRCCS) are pleased to announce a two-week Summer School in high performance computing using the Linux operating system and commodity hardware.

The School aims to provide the skills needed to benefit from the next generation of MPP HPC solutions, to administer and tune a Linux cluster or to purchase or build Linux-based HPC systems oneself.

This event will bring together people from all areas of the Linux/HPC world. The aim is to provide UK researchers with thelatest developments and tools for running scientific applications in Linux clusters and supercomputers. among those who have been invited are:

Dr Martyn Guest of CLRC Daresbury to discuss the running of real-world scientific applications on Linux (accepted). The Portland Group to discuss the latest compilers for Linux (accepted) SGI developers working with the SGI Linux Supercomputer at Ohio State University Intel developers to talk about performance issues related to the IA64 architecture.

There will also be lab sessions and tutorials on topics such as

- Building a Linux cluster from scratch, which will give participants the opportunity to do exactly this
- Systems administration of a Linux cluster
- Tools for parallelisation of Linux clusters, message passing libraries in a clustered environment
- Handling a multi-user environment.

We will provide access to a variety of types of cluster including clusters using shared-memory nodes and incorporating thetechnology being used in the ASCI program centres in the US. There will be opportunity to try out your own codes in these environments.

The course is designed for UK scientific researchers wishing to learn about how Linux could benefit their research.

The course is free for all registered users of the CSAR Service. For those not associated with CSAR, CfS has made available a number of grants for academic attendees to cover the £750 courses fees. If you wish to claim one of these, please apply as early as possible as places are limited. The course fee of £750 is applicable to commercial organizations.

For more information, see: http://www.man.ac.uk/mrccs/summer_school/2000/ or contact the CSAR helpdesk: email: csar-advice@cfs.ac.uk, tel: 0161 275 6824/5997.