ISBN 1470-5893

CSAR FOCUS

Edition 7 ,Summer 2001, Editor: Kaukab Jaffri, University of Manchester

Euro-Par 2001 Manchester

Manchester hosts Euro-Par 2001

Euro-Par is an annual series of international conferences dedicated to the promotion and advancement of all aspects of parallel computing

28th - 31st August 2001

KEYNOTE SPEAKERS

Jack Dongarra, USA
Ian Foster, USA
Dennis Gannon, USA
Tony Hey, UK
Martin Kersten, The Netherlands,
Thomas Sterling, USA

http://www.man.ac.uk/europar

REGISTRATION NOW OPEN!

www.csar.cfs.ac.uk

CSC

THE UNIVERSITY of MANCHESTER
1851 - 2001
one hundred and fifty years of excellence

sgi

Contents

Editorial

This issue of CSAR Focus brings you details of the annual Euro-par conference, which will be hosted jointly by the University of Manchester and UMIST, 28th August - 31st August 2000. Professor John Gurd, from the Department of Computer Science at the University of Manchester and Chair of Euro-Par 2001, tells us a bit more about what there is on offer and why you should attend.

Among the other conferences being held here at Manchester Computing this year is Eurographics 200,1 running from 4th - 7th September, details of which can be found on page 24.

There are also details of the Workshop on "HPC and SGI NUMAflex" to be held in 8th - 12th October, which will replace the HPC Summer School at the University of Manchester.

We pleased to announce that the University of Manchester is hosting the Cray User Group (CUG) Summit 2002 in May next year, details can be found on page 25.

We have two reports on CSAR Class 3 projects, Professor Peter Coveney's consortium writes about the Class 3 work that they have been invloved in, and Dr Emma Finch from the Department of Earth Sciences here at the University of Manchester talks about her Class 3 work.

As well as an update on the CSAR Technolgy Refresh, we have an article about the life of mathematician, George Green, after whom we have named the latest new CSAR machine, the SGI Origin 3000.

We also feature an article on the software tool VIPAR - Visualization in Parallel.

The results of the 2000 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
Kilburn Building
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

Disclaimer

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

Copyright

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

Euro-Par 2001
28th - 31st August, Manchester Conference Centre, UMIST

Professsor John Gurd,
Department of Computer Science,
University of Manchester

Euro-Par is a leading international conference in the field of parallel computing. It was formed in 1995 when a number of previously distinct events were merged together to form Euro-Par'95, held in Stockholm. Since then Euro-Par has been held in Lyon (1996), Passau (1997), Southampton (1998), Toulouse (1999) and Munich (2000). The conference website is at URL:

http://www.euro-par.org

Euro-Par is held annually. It is dedicated to the promotion and advancement of all aspects of parallel computing. The major themes can be divided into the broad categories of hardware,

Porting the AMBER forcefield to LAMMPS–massively parallel molecular dynamics simulations of DNA

C.R. Grindon*, S.A. Harris*, T.J. Evans$, K.E. Novik$, P.V. Coveney$, C.A. Laughton*.

*Cancer Research Laboratories, School of Pharmaceutical Sciences, University of Nottingham, Nottingham NG7 2Rdleiver a.

$Centre for Computational Science, Department of Chemistry, Queen Mary, University of London, Mile End Road, London E1 4NS.

Molecular dynamics (MD) is proving to be a valuable method for analysing the structure and flexibility of DNA (1). Atomistic MD simulations of DNA are computationally very expensive and, using conventional algorithms, are limited in practice to around 10ns (2). There are important motions within DNA which are predicted to occur over longer timescales than this (Table 1).

Table 1:– Timescales for typical DNA motions.

Time scale	Main types of internal motion
Picosecond	Short living motions and oscillations of atoms.
Nanosecond	Oscillations of small groups of atoms: sugars, phosphates, bases; bending and twisting of the double helix.
Microsecond	Winding and unwinding of the double helix; opening of base pairs.
Millisecond	Dissociation of the double helix; super helicity; overall rotation.
Second	Writhing; isomerisation; division of bacteria.

Porting of the Forcefield

LAMMPS (Large Atomic/Molecular Massively Parallel Simulator) (3) is a parallel MD code with accurate treatment of long-range electrostatic interactions (a particularly important consideration in DNA simulations) based on the PPPM (particle-particle/particle mesh) algorithm and Ewald summation to handle the periodic boundary conditions. The advantage of LAMMPS over other codes is its ability to run calculations on parallel machines with large numbers of processors without great loss of efficiency (4), increasing the size and complexity of systems one is able to study.

We have ported the AMBER forcefield (5), which is well established for the simulation of DNA dynamics, into LAMMPS. The aim of this study has been to test this porting by comparing the dynamics of the DNA dodecamer d(CTTTTGCAAAAG) as predicted by LAMMPS with dynamics data from our previous extensive analysis of this sequence using AMBER (6).

Minor changes to the LAMMPS code were required to cope with AMBER's treatment of 1-4 non-bonded interactions. Static energy calculations then showed excellent agreement between AMBER and LAMMPS (table 2).

Table 2: Static energy analysis of a representative structure of $d(CTTTTGCAAAAG)_2$.

ENERGY TYPE	AMBER (Kcal/mol)	LAMMPS (Kcal/mol)
Bond	0.0239	0.0239
Angle	399.8833	399.8833
Dihedral	438.7989	438.7989
Total VDW	2532.0560	2532.2143
Total Electrostatic	-27167.3825	-27167.0726
Total Energy	-23796.6204	-23796.1522

Temperature Control

AMBER implements the Berendsen algorithm to control temperature but this is not available in LAMMPS. Two alternative temperature controls are available in LAMMPS, Langevin and Nose-Hoover. The latter approach was abandoned due to a hot solute/cold solvent problem. We have adopted the "Fix Style Langevin" temperature control where the solvent and solute temperatures are scaled separately.

Table 3: Average temperatures during 10ps runs using Nose-Hoover and Fix Style Langevin temperature couplings.

	Average Temperature (K)	Standard Deviation
Nose-Hoover Solute	320.4484	11.5720
Nose-Hoover Solvent	296.7024	5.8680
Nose-Hoover Whole system	299.7040	5.0824
Fix Style Langevin Solute	300.7899	8.9964
Fix Style Langevin Solvent	299.4550	3.9176
Fix Style Langevin Whole system	299.6238	3.5603

The hyperflexibilty of the DNA, particularly at the termini of the double helix, when Nose-Hoover temperature coupling was used are apparent from the analysis of atomic fluctuations (Graph 1).

Graph 1 – Atomic fluctuation data for 10ps runs of Nose-Hoover and Fix Style Langevin temperature couplings compared to AMBER.

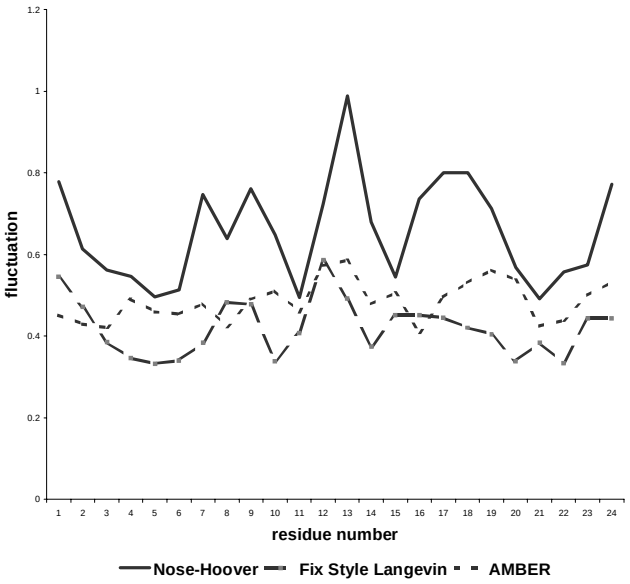

Scaling and Efficiency

We are using the 816 node Cray T3E-1200E supercomputer at CSAR, the UK's national supercomputing facility, to run dynamics. Short (10ps) simulations on up to 64 processors indicate that the code is relatively efficient (equation 1) and scales well, although not linearly, in comparison to AMBER which performs much more poorly in parallel situations (see graph 2).

Efficiency = np*t/l (Equation 1)
Where np = number of processors
t = time taken (s)
l = length of simulation (fs)

Graph 2 - Efficiency of LAMMPS compared to AMBER, normalised to the efficiency of LAMMPS on 1 processor.

ANALYSIS OF LONGER SIMULATIONS

Three longer simulations (3ns+) have been run using 64 processors and analysis has been carried out on 2ns equilibrated portions of these. The three simulations differ only in temperature rescaling parameters, LAMMPS 1 = 0.01, LAMMPS 2 = 0.001 and LAMMPS 3 = 0.0001 (rescaling parameter in inverse time units (fs^{-1}) therefore LAMMPS 3 has most relaxed temperature rescaling).

Three analysis techniques were used to compare the LAMMPS simulations to an AMBER simulation previously carried out:

1. RMSD. RMSD's have been calculated between the time-averaged structure from AMBER and corresponding time-averaged structures from the three LAMMPS simulations.

2. ENTROPY. Configurational entropies (7) have been calculated to obtain an overall representation of the flexibility of these systems compared to AMBER.

3. PRINCIPAL COMPONENT ANALYSIS (PCA). PCA has been used to identify the major modes of motion in each trajectory. The similarity between these modes in the AMBER simulation and in each LAMMPS simulation was characterised by calculating the overlap of the top 10 modes (eigenvectors) (8).

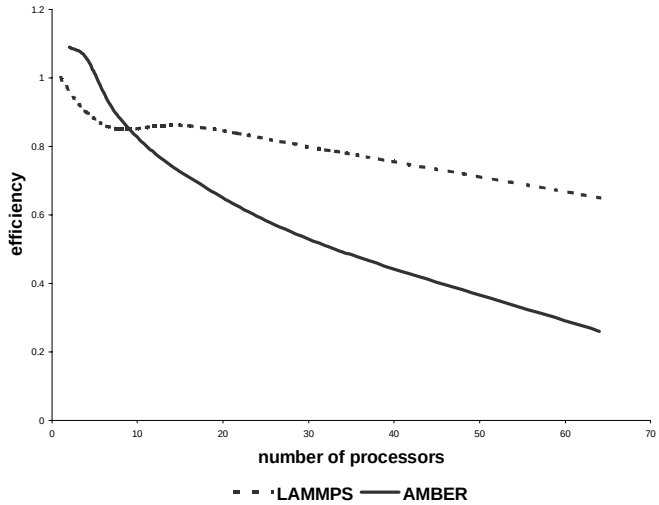

Table 4 – Characteristics of LAMMPS simulations compared to the original AMBER simulation.

SYSTEM	RMSD (Å)	ENTROPY Kcal/mol	PCA OVERLAP
AMBER	n/a	693.67	n/a
LAMMPS 1	0.9483	589.79	0.4214
LAMMPS2	0.6486	637.56	0.6309
LAMMPS3	0.6297	646.17	0.6678

Summary

The AMBER forcefield for DNA has been successfully ported into LAMMPS, this is confirmed by excellent agreement of static energies.

We have determined (some of) the optimal parameters for stable simulations of DNA in LAMMPS. We have shown that, given the same problem, LAMMPS scales better than AMBER on a T3E. We have shown that the time-averaged behaviour of the DNA is well preserved between the two simulation techniques and that dynamical characteristics are also well preserved. It is clear from this data that LAMMPS has great potential in this field. Given large scale access to the T3E at CSAR, we would expect to be able to probe hitherto unprecedented regimes of dynamical behaviour.

Acknowledgements

We would like to thank CSAR for access to the Cray T3E, Steve Plimpton for his assistance and the BBSRC for funding of a Class 3 grant.

References

(1) Sherer E.C., Harris S.A., Soliva R., Orozco M., Laughton C.A, *J.Am.Chem.Soc.,* 1999, **121**, p5981.

(2) Cubero E., Sherer E.C, Luque F.J., Orozco M., Laughton C.A., *J.Am.Chem.Soc.,* 1999, **121**, p8653.

(3) Plimpton S.J., 1997, LAMMPS version5, CRADA Collaboration, Sandia National Laboratory, USA.

(4) Maillet J.-B., Lachet V., Coveney P.V, *Phys.Chem.Chem.Phys,* 1999, **1**, p5277.

(5) Case D.A. *et al.* 1999, AMBER 6, University of California, San Francisco.

(6) Gavathiotis E., Sharman G.J., Searle M.S., *N.A.R.* 2000, **28**, (3) p728.

(7) Schlitter J., *Chem.Phys.Lett.,* 1993, **215**, (6) p617.

(8) Hess B., *Phys.Rev.E,* 2000, **62**, (6) p8438.

CSAR Technology Refesh - Update Summer 2001

Dr. Stephen Pickles,
CSAR Technolgy Refresh
Team Leader

In the last issue of CSAR Focus, I described CSAR's plans for the Technology Refresh. But six months is a long time in HPC, and a lot's happened since then.....

The biggest surprise is the accelerated implementation schedule for the Origin 3800 system "Green". Already at 256 processors (that's *double* the promised capacity), Green will see the addition of yet *another* 256 processors in July. It will be Green, not Napier, that takes over from Turing as CSAR's flagship facility.

At the same time as the schedule for Green has been brought forward, the schedule for Napier has been put back. It is now highly unlikely that any SGI NumaFlex system based on processors from Intel's Itanium family will be going into production service at CSAR in the next twelve months. This article is intended to bring you up to date with the CSAR Technology Refresh.

Fermat

The Origin 2000 system "Fermat" was upgraded to 128 MIPS R12000A 400 Mhz processors and 128 GBytes of memory in December. Despite a few teething problems, the take-up of Fermat for production work was so encouraging that we brought forward the upgrade schedule for Green.

Green

The new SGI Origin 3800 system, "Green", was commissioned in April with 128 MIPS R12000A 400 MHz processors and soon upgraded to 256 processors. Yet *another* 256 processors will be installed during July, and all 512 processors should be available for production work in August.

With a massive 512 gigabytes of memory, a state-of-the-art interconnect, and support for MPI, SHMEM *and* OpenMP, Green can run scientific applications that simply aren't feasible on Turing. Green and Fermat are binary compatible – they even use the same processors. But Green has a better memory sub-system and a faster interconnect, so applications should scale better on Green. Green has already displaced Turing as the system of choice for many capability users.

CSAR is going to try operating Green rather differently to Fermat. There will be no interactive work on Green. In fact, you won't normally be able to log in. At first glance, you might think this is a step backwards, but there are reasons for doing it this way. By making all the processors on Green available for batch work, we win in two ways. Firstly, we can schedule jobs more efficiently, thereby increasing overall throughput to everybody's benefit. Secondly, since more processors can be available to a single batch job, we can run larger jobs than we could otherwise. You still have interactive access on Fermat for program development, testing and debugging; we will monitor demand for the interactive domain on Fermat closely to see if we need to increase the number of processors available for interactive work. Your home directories on Fermat are your home directories on Green, and we've tried to make the software environments of the two systems as similar as possible to minimize the need for changing job scripts. In fact it should be easy to write scripts that work on either system. We hope we've got this right, but we can't be sure without your feedback.

Storage Area Network

As part of the Technology Refresh, CSAR has undertaken an analysis of storage and I/O infrastructure. As a result of this analysis, we will also be introducing a Storage Area Network

which will provide a consistent view of files across the new systems and vastly improved data bandwidth, at the same time as reducing the interdependency of the CSAR systems and the reliance on NFS.

Fourier

Fourier, the SGI-supplied cluster based on pre-release Intel Itanium processors and a Myrinet interconnect has now increased in size to 32 processors (16 nodes x 2 processors/node, each node having 2 Gigabytes of memory). The primary purpose of Fourier has been to facilitate the porting users' application codes to Itanium and Linux, and single-node optimisation of the same codes on Itanium. But these use now seem much less relevant given the increased importance of Green, and any reasonable use, such as small-scale production work or mixed mode programming (combining MPI and OpenMP in a single application) is also permitted. CPU time on Fourier continues to be free, at least for the time being. Intel Non-Disclosure Agreements will continue to apply, and new users will be required to sign a form, until such time as Fourier is upgraded to production quality Itanium processors. This upgrade will not occur until Intel have announced the commercial availability of Itanium processors, but this announcement is imminent.

Migration strategies

The recommended migration strategy- port to Green- is much simpler than it was six months ago.

But not everyone will be excited by the prospect of porting codes and transferring production work to a new system. The good news is that you don't have to do it on your own.

Code Portability Assessment Programme

In addition to the normal application and optimization support mechanisms, CSAR is also offering *free* Code Portability Assessments to every project. We take a copy of your code and a test case that runs on Turing, work on it for five days to port it to the CSAR system of your choice, then give it back to you along with a report showing what we did and why. If there are any outstanding obstacles, we will explain what they are. The offer is strictly "while stocks last" (our resources are not unlimited), and requests will be processed on a first come, first served basis. So act now to avoid disappointment. See http://www.csar.cfs.ac.uk/using/portability.shtml for more information on the Code Portability Assessment Programme.

Training

Because of the importance of the Technology Refresh, we decided to replace the annual MRCCS Summer School by two 5-day residential workshops on relevant topics. The first of these, "HPC on Intel Itanium" was originally scheduled for 2-6 July, 2001, but has since been postponed indefinitely. The second "HPC on SGI NumaFlex" is scheduled for 8-12 October, 2001, and will be focused primarily on topics of immediate relevance to users of Green. The registration fee is waived for all CSAR users. More information is available at http://www.man.ac.uk/mrccs/summer_school/2001/.

We're also providing courses on Origin 2000/3000 and Intel Itanium optimisation as part of our normal training programme. See http://www.csar.cfs.ac.uk/using/courses.shtml for more information.

Technical Survey

Thanks to all those who completed the CSAR Technical Survey. Your feedback is valuable to us, and will help to shape the future of the CSAR service. The survey is now closed.

If you would like expert advice on how the technology refresh affects you, please contact the CSAR helpdesk. A CSAR applications support specialist will be in touch to discuss your particular requirements.

George Green (1793-1841)

Dr Dan Kidger,
CSAR Applications Support

George Green was an English Mathematician who published work in the fields of hydrodynamics, electricity and magnetism, although he is perhaps best known for 'Green's Theorem' which is used in almost all computer codes that solve partial differential equations.

Early Life

George Green was born in 1793 to George Green snr. – a baker in central Nottingham and his wife Sarah. They also had a daughter two years later. At the age of eight, he started school at Robert Goodacre's Academy. Mr Goodacre acquired many scientific instruments for his school, including a barometer and an orrery. George Green's father's business started to prosper, and in 1807 his father established a windmill on the village of Sneinton, 2 miles east of Nottingham. Ten years later, Mr. Green had a large house built next door, and the family moved in.

Green the Miller

In 1829, George Green snr. died and left the mill and business to George. In 1831, there was a riot where angry mobs attacked the mill. George Green fired his musket at the crowd, whilst his sister passed him ammunition. George Green spent his days working at the mill. However every day he continued to study mathematics late at night in the uppermost room of the mill

'Jane'

The mill manager was William Smith who lived in a cottage next to the mill. He had a daughter, Jane Smith (born 1802). Jane and George had a relationship that lasted for sixteen years. They had a daughter in 1824, and a subsequently six more children together, but strangely never married.

His key paper of 1828

In 1823, George joined Bromley House Library in Nottingham, which was the centre for academic study in the town before the University was established. In 1828 he wrote his key *paper 'An Essay on the Application of Mathematical Analysis to the Theories of Electricity and Magnetism'.* This paper although disregarded ant the time, has been regarded as the `beginning of mathematical physics in England'. George had to pay for its publication himself. Sir Edward Broomhead, a local landowner read the paper and encouraged Green to write three more papers which Broomhead got published in Cambridge Philosophical Society. Broomhead had a friends at Cambridge including Charles Babbage.

Cambridge

Green had by now a strong reputation, but felt he was still a miller and not yet part of the academic establishment. So, in 1833, Green let out his mill and moved to Cambridge to study as an undergraduate at Caius College. He was by now forty years old. He passed Exams in Latin, Greek and Ecclesiastical History in spite of the fact that most other undergraduates were only there for the "Parties and Drunkenness" *(some things never change! – ed.)*

He graduated in Mathematics in 1837. He published six more papers before being elected as a fellow of the college in 1839, but unfortunately his health began to fail. He died back in Sneinton on 31st May 1841. He was buried alongside his parents in the churchyard of St Stephen's, close by his windmill.

The famous windmill in Nottinghamshire where Green lived. It is now open the public as a museum and science centre (http://www.innotts.co.uk/greensmill/)

References

1. D M Cannell *George Green , Miller and Mathematician* paperback, illustrated, £3.00

2. http://www-groups.dcs.st-and.ac.uk:80/~history/Mathematicians/Green.html

3. http://www.nottingham.ac.uk/~ppzwww/green/homepage.htm

The SGI NUMAflex Design

– SGI

The new SGI™ Origin™ 3000 servers and SGI Onyx® 3000 graphics systems use a new modular hardware design called NUMAflex™. NUMAflex is the third generation of the SGI nonuniform memory access (NUMA) architecture. In the NUMAflex design, the various internal components of a computer system are separated into modular components, or "bricks," that provide CPU and memory, I/O, interconnection, graphics, and storage.

The brick-and-cable SGI NUMAflex architecture isolates components, so they can be exchanged, serviced, and upgraded independently of each other. Many hardware components can be removed and serviced as needed without interrupting system operation. Individual bricks can rest on a service tray (see Figure 1), which locks to the front of the system rack at various heights, while maintenance is performed. The patented SGI peripheral component interconnect (PCI) carrier sled lets you swap PCI cards without tools while the system is running for all the PCI slots in the PCI expansion brick and I/O brick.

The NUMAflex design also replaces the central bus, which is shared by components in conventional computer architecture, with the SGI crossbar switch technology. The crossbar switch dynamically and directly links any two computer subsystems, giving them a high-bandwidth, low-latency path without interfering or competing with other system activity. The crossbar switch allows multiple, simultaneous, dynamically allocated connections between system components. The modular NUMAflex design and improvements in SGI NUMA architecture provide greater system and memory bandwidth, reduced latencies, and greater application scalability than existing bus-based symmetric multiprocessing (SMP) systems.

The SGI NUMAflex design allows you to purchase and configure a system brick by brick to meet your needs. You can purchase only the computation, memory, I/O, graphics, and storage that you need; then, you can add more bricks or reconfigure your system as your needs change. For maximum fault containment, the IRIX® 6.5 operating system supports partitioning, allowing you the flexibility to divide a single large SSI (single system image) server into smaller SSI systems.

Figure 1: NUMAflex Open Rack with Service Tray

The Bricks

With NUMAflex technology, each drawer-like brick in a system has a specific function and can be linked, through the SGI NUMAlink™ high-speed system interconnect, to other bricks of varying types to create a fully customized configuration. The same bricks, depending on their number or configuration, can be used to upgrade, change, or improve your system.

The following are the types of NUMAflex bricks:

- R-brick (router interconnect)
- C-brick (CPU module)
- I-brick (base I/O module)
- D-brick (disk storage)
- P-brick (PCI expansion)
- X-brick (XIO expansion)
- G-brick (InfiniteReality® graphics)

The following sections describe the various bricks.

R-brick (Router Interconnect).

The R-brick is the structural building block of the system, a high-speed crossbar switch that connects processors and memory. Unlike a traditional system bus, the crossbar switch dynamically and directly links any two computer subsystems, giving them a high-bandwidth, low-latency path without interfering or competing with other system activity. The R-brick lets you add infrastructure as you need it — from routerless deskside systems to an eight-port router multirack configuration that allows for 512 processors in a single shared-memory environment. In addition, the R-brick allows each system component to be serviced or upgraded individually. Four ports on an R-brick connect to a NUMAlink port (1.6 GB/s each direction) on a C-brick; four other ports connect to other R-bricks.

R-brick

C-brick (CPU Module)

The C-brick is the CPU module. It can contain up to four MIPS® R12000A processors and up to 8 GBs of local memory. A NUMAflex crossbar memorycontroller delivers 200% greater CPU-to-memory bandwidth and twice as many CPUs per node than SGI Origin 2000 systems. The C-brick has one NUMAlink port (1.6 GB/s each direction) for connectivity to another C-brick or an R-brick and one Xtown2 connector (1.2 GB/s each direction) for connectivity to an I-brick. (XTown2 is an XIO network interface used for Gigabyte System Networks (GSNs).)

C-brick

I-brick (Base I/O Module)

The I-brick, standard in all systems, provides the base level of I/O. The I-brick contains the fibre channel system disk, a DVD- or CD-ROM drive, and four PCI slots. The I-brick also provides network access through a 10/100Base-T Ethernet port and peripheral access through one IEEE 1394 channel and two USB channels. The I-brick has two Xtown2 ports (1.2 GB/s each direction) that connect to an Xtown2 port on a C-brick or G-brick.

I-brick

D-brick (Disk Storage)

The D-brick is a disk enclosure that supports fibre channel attached disks. A D-brick has standard dual power supplies and supports up to 12 drives that have 18, 36, or 73 GB capacity. The D-brick connects through a fibre channel mPCI card to an I-brick or P-brick.

D-brick

P-brick (PCI Expansion)

The P-brick provides for additional PCI expansion. It supports 12 PCI cards distributed over six 64-bit, 66-MHz buses. The total peak bandwidth is greater than 3 GB/s. The P-brick has two Xtown2 ports (1.2 GB/s each direction) that connect to Xtown2 ports on a C-brick or G-brick.

P- and X- brick

X-brick (XIO Expansion)

The X-brick is the XIO expansion brick. It contains four half-height XIO slots that are fully compatible with the XIO slots in SGI Origin 2000 systems. They provide support for HIPPI, GSN, VME, and digital video. The X-brick has two Xtown2 ports (1.2 GB/s each direction) that connect to an Xtown2 port on a C-brick or G-brick.

G-brick (Graphics Expansion)

The G-brick is the graphics expansion brick. It provides integrated InfiniteReality3™ visualization capabilities that can be scaled from 1 to 16 pipes (eight G-bricks in a system). Each G-brick has one port that supports two raster managers (2RM port) and one 4RM port. Each pipe of a G-brick connects to an Xtown2 channel of an I-, P-, or X-brick through a DNet cable.

G- brick

Power Bay

The power bay contains three to six distributed power supplies (DPSs). The power bay provides industry standard 48 volts DC power to C-, I-, P-, X-. and R-bricks. The power bay is configured in an N+1 redundant configuration so that any individual DPS failure will not affect the operation of the system.

Copyright

The SGI NUMAflex Design Copyright 2000, Silicon Graphics, Inc. All rights reserved. Used by permission of Silicon Graphics, Inc.

Silicon Graphics, InfiniteReality, IRIX, and Onyx are registered trademarks and SGI, InfiniteReality3, NUMAflex, NUMAlink, and Origin are trademarks ofSilicon Graphics, Inc. MIPS is a registered trademark of MIPS Technology,Inc., used under license by Silicon Graphics, Inc.

Segmentation, Growth and Linkage of Normal Fault Arrays: A Three-Dimensional Numerical Analogue of Rift Evolution

Dr Emma Finch,
Department of Earth Sciences,
University of Manchester

Active and ancient normal fault arrays comprise fault zones that are discontinuous along strike and consist of a number of distinct segments at a variety of scales. It is known that the progressive growth of normal faults ultimately results in linkage of originally isolated fault segments to form large (25-50 km) fault zones. Current understanding of the development of such zones, however, is restricted by an inability to resolve the temporal evolution of relationships between the spacing, length and along-strike segmentation of natural faults in three dimensions over geological time scales.

One of the biggest problems in investigating these large-scale earth systems by numerical methods is the dimension of the media required to accurately simulate crustal processes and the need for a large amount of computational time for simulations to be representative of geological time scales. With the assistance of members of CSAR, a three-dimensional numerical analogue of extension is being developed to investigate the effect of thickness of the brittle crust (seismogenic layer) on the growth and linkage of faults during tectonic extension. The model uses a discrete element technique that represents the crust as spherical particles that interact through physically realistic forces in a two-layer, brittle/ductile system. Terms for gravitation and isostatic flotation are also included. Currently, the models are running using regular particle distributions in small media (~30,000) to ensure that the code is stable. Figure 1 illustrates the surface topographic profile of a medium at 16% extension of the crust in one such numerical simulation. Here, the topography scales from deep (light coloured) to high (dark coloured) terrain. Examples of faults are indicated on the figure, the pale grey surface represents a sea-level added to aid visual acuity. The work is already beginning to throw light on questions from geology regarding the relationship between the growth and linkage of faults and their spatial distribution within the brittle crust.

Figure 1 shows that regular geometries lead to preferred orientations for failure within the simulations. It is planned that further development of this code will permit investigation using media that consist of greater than 10^6 elements in a random distribution. It is hoped that results from these experiments will provide templates for interpretation of structures imaged in 3D subsurface datasets and allow investigation/ prediction of the distribution of sub-seismic fault populations. The code is also applicable to the investigation of earth surface processes and collaboration is currently underway to couple this 3D model with landscape evolution models that are being developed by Dr. Stuart Hardy in the Department of Earth Science, Manchester.

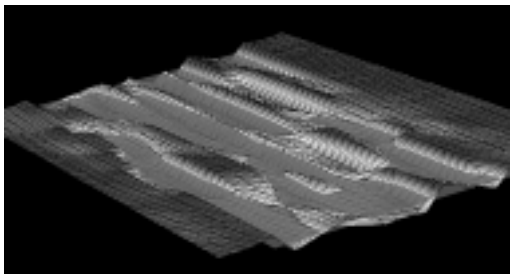

Figure 1; Example of the upper surface topography of the 3D model after 16% extension to the left. The crests of a number of faults are indicated by white circles.

This work has recently been presented at the American Geophysical Union Fall Meeting in December, 2000 and further analysis of the data will be presented at the American Association of Petroleum Geologists in June, 2001. The work was funded for one year by EPSRC.

Acknowledgements

Thanks to Terry Hewitt and Mike Pettipher of CSAR for their help and guidance at the start of this project.

VIPAR: Visualization in Parallel

Jo Leng, Kevin Roy, Dr J.M.Brooke
CSAR Applications Support

VIPAR was originally a 2 year EPSRC research project started in April 1996 and concluded in December 1997. It was carried out at the Manchester Visualization Centre, University of Manchester and had industrial partners AVS Inc. and Meiko Ltd. VIPAR started as proof of concept research but now there is interest in making it into a stable tool that can speed up the visualization pipeline. A preliminary study of VIPAR and the work required was conducted as a UKHEC case study. A report will be available from the UKHEC web site shortly (http://www.ukhec.ac.uk).

The project's original aim was to remove performance bottlenecks from within the visualization pipeline. At that time several other projects had implemented solutions that were tied to specific parallel support libraries, hardware or visualization systems. These all had poor portability, VIPAR was developed to have better portability.

The Structure of VIPAR and It's Libraries

Visualization systems like AVS/Express, Iris Explorer, IBM Data Explorer and Khoros are described as Application Builders or Modular Visualization Environments (MVE). These tools provide access to large visualization libraries via a highly developed GUI which is easy for new or non-technical personnel to operate. The user drags and drops modules from the systems libraries and connects them to produce a network which in turn produces an associated application and visualization.

The VIPAR libraries are a suite of routines that provide an interface between the visualization system and the message passing system. The three libraries are, VPRvsi - a visualization library to provide an interface to the visualization system, VPRidd - an intermediary library that performs the necessary calculations and provides the interface between the two other libraries, and VPRdd - the lower level message passing and parallel process control library. These libraries are designed to sit on top of one another and provide increasing levels of complexity.

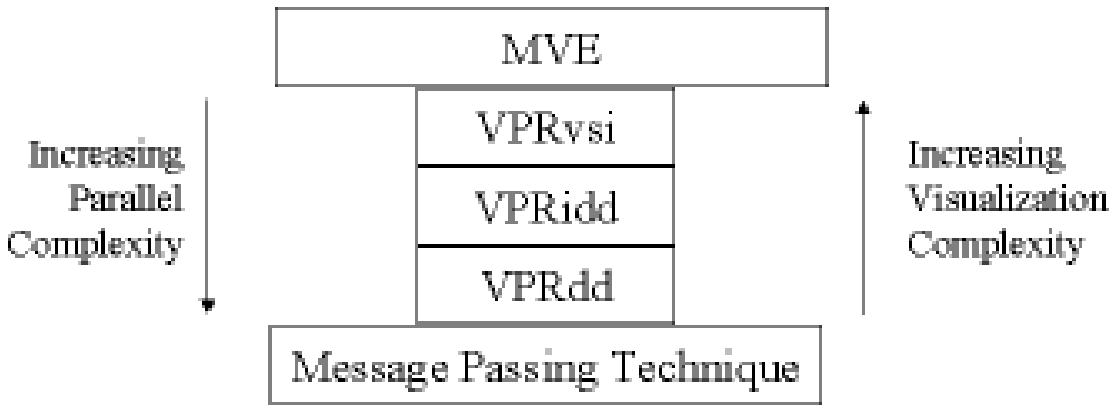

Fig 1: Shows how the VIPAR libraries are layered to achieve portability between MPI libraries and the visualization system.

VIPAR is also a construction kit for producing parallel visualization modules quickly and easily. However this tool, the DDTool, was only ever produced as a prototype.

Current Interest In VIPAR

Since the end of the VIPAR project there have been several changes that have created interest in re-opening the project:

* Development of hardware rendering systems has meant that the major visualization bottlenecks have moved away from the rendering components into other parts of the visualization pipeline. It makes sense to parallelize just the bottlenecks, i.e. the CPU intensive components.

* Now there are more and larger MPP systems allowing users to increase the size of data sets.

* More application areas are using visualization.

* Computational steering/monitoring is becoming more important.

Technical Aspects of Reviving VIPAR

The main technical problem of reviving VIPAR came from the fact that VIPAR originally started as a research project, it was not a software development project. It proved its hypothesis by proof of concept, there was no tested VIPAR distribution produced.

The Parallel Libraries

When the original VIPAR was designed it was assumed that MPI libraries would move to the MPI-2 standard and this would have spawning included, parallelisation within VIPAR currently relies on spawning. If this had happened the complexity of VIPAR could be reduced, however, it is still only LAM MPI that supports spawning. LAM MPI is freeware and not fully implemented/tested on all platforms.

Now VIPAR is being re-evaluated a more robust solution is being considered. If a socket was put into VIPAR it could be used to call codes on the local or remote machine and these codes could use the most efficient parallel library whether it is MPI or Open MP. This new solution would also reduce the complexity of VIPAR and so increase its usability.

The Visualization System

The original VIPAR modules were designed for academic testing and did not have the sort of robustness required for routine work. Much effort has been put into transforming VIPAR's structure into a sturdy project structure (the International AVS Centre's project structure), making the linking between AVS/Express and the VIPAR libraries non-specific (connected via environment variables rather than absolute paths) and re-implementing components so that 2 or more VIPAR modules could be used in one session. At the same time the user interface has been re-developed to be more robust and hopefully idiot proof.

The Current Status of VIPAR

VIPAR has currently been ported so that it will run the isosurfacing example with the 32 bit version of AVS/Express and LAM-MPI on Fermat (a SGI Origin). The user interface, error handling and structure have all been improved to make the it more robust to use and install.

The isosurfacing module produces a surface of a given value within a scalar volume of data. The surface is in effect a 3D contour. You would use an isosurface if you wished to see the graphical depiction of a particular data value within a volume of data. The figures below show an isosurface and its associated network in AVS/Express. When the parallel isosurface module, PMiso is expanded you can see that it is made up of a distributor, a harness and a user interface.

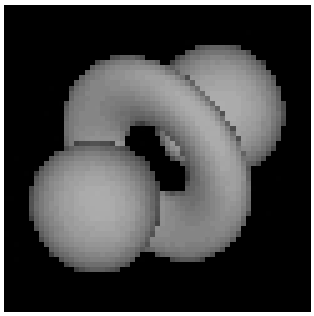

Fig 2: Isosurface of the electron potential of a hydrogen atom.

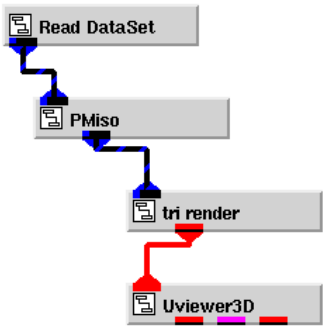

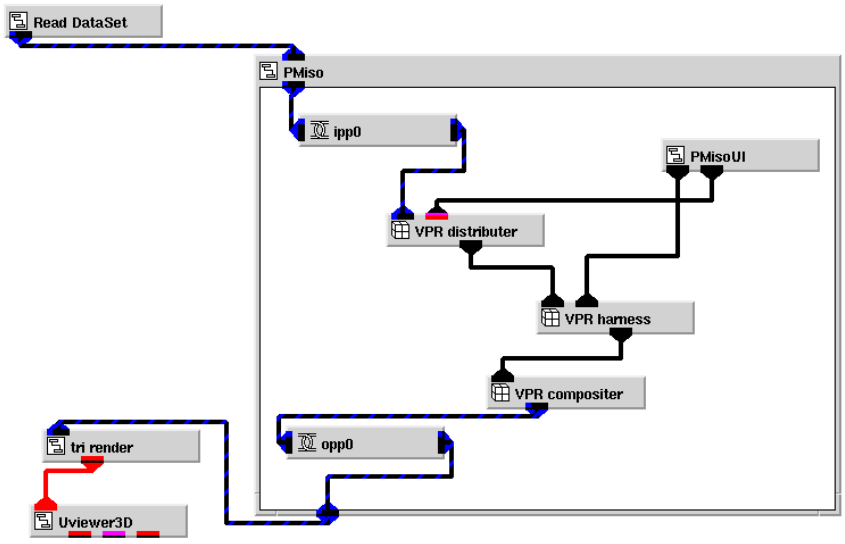

Fig 3: The network containing the VIPAR isosurfacer. The image on the left has the VIPAR module expanded to show the internal structure.

In the future we aim to take VIPAR out of research and to make it into a stable product that can communicate through a portal to other machines and parallel libraries. Not only would this improve the performance of VIPAR but would enable computational steering and reuse of computational functions within the visualization system.

Acknowledgements

The original research project was developed by S. Larkin, A. Grant and W.T. Hewitt at Manchester Visualisation Centre, Manchester Computing, University of Manchester.

People Focus

Jon McClaren

I started work for CSAR on 5th March 2001, working as a Software Engineer on the EUROGRID project, reporting to John Brooke. I work on developing GRID technologies; specifically, I will be constructing a GRID Resource Broker, with the help of John Brooke, and people from FECIT. I am currently completing a test installation of UNICORE which can run jobs on our Cray T3E.

Previously, I did a MPhil and PhD here at the University of Manchester, in the Department of Computer Science, where I was supervised by John Gurd. My MPhil compared different parallelisation paradigms, including directive-based parallelisation and MPI. My Ph.D was in parallel computing - I developed a novel method for automatic parallelisation, based on a hypothesis about data structures in scientific programs. The PhD was only recently submitted, and I'm still waiting for the Viva. I did my BSc over in York, and inbetween my two educational stints, I worked for a couple of years as a Software Engineer with a company in Altrincham. At this company, I was in a development team of around 20 people, working on the company's core product, a LIMS (Laboratory Information Management System). At this time, I worked mainly on UNIX systems (with some NT and VMS work too), writing in C; I also became very familiar with the ORACLE database system.

During my PhD, I became a film addict - usually I go to the cinema twice a week, or even more. Other interests include listening to various forms of music, including upcoming Manchester bands like "Elbow" and "I am Kloot", as well as classical music (especially when its at the Bridgewater Hall). Now that the PhD is done, and I have more time, I'm getting back into cycling and swimming.

Mike Daw

I've been working with computers for the last 5 years. My last job was with Logica - a company that was described by its CEO as a "British Microsoft". I think I'm still bound by a confidentiality agreement, so I won't comment on this.

My other IT jobs were for eccentric little software houses populated by an eclectic, interesting and sometimes talented bunch of people who have been my guides and mentors.

Specifically, my background has been in all aspects of software development in a variety of programming languages - COBOL, Visual Basic and C being the main ones - but with a common thread of a lot of database work (in administration and SQL).

In the more distant past, I worked as a Maths Teacher in Hyde.

I am currently working on a project to establish the UK's first Access Grid node. This will enable us to take part in SC2001 Global, Denver, Colerado, USA in November. The organisers of this event call it the "first truly global technical conference on the Grid: a multi-national and multi-cultural meeting place for communication and discussion of ideas relating to high-end computing and communications and its impacts on science and society".

I am also taking over responsibility for the "C++ for Scientific Applications" course. In the longer term, as my experience of our supercomputers grows, I will be doing more direct CSAR user support.

I'm looking forward to the challenges of adjusting to the academic culture. I even grew a beard to fit in, but can't afford the quality of pedicurist to wear sandals!

Outside of work, I am captain of a softball team called the Piranhas who play in Manchester Softball League. Unfortunately, our "snap" is worse than our bite. (It's not for nothing that we're known as the Manchester City of the 1st division because of our bid to avoid being relegated.)

My other leisure activities include song-writing, playing and table tennis.

Jon Gibson

I have been working as a member of the CSAR application support team since the last day of February 2001. I have recently returned from Stuttgart, where I gave a presentation at the 4th HLRS Metacomputing Workshop, entitled "The Jodrell Bank pulsar search code and the move to real-time processing". This followed up work in which I have started to investigate the parallelisation of such a system. I have found the work at CSAR to be both challenging and enjoyable and am looking forward to the future here.

I was previously at the University of Manchester between 1989 and 1992, studying electronic and electrical engineering. I went on to get an M.Sc. in advanced scientific computation at Liverpool University, followed by a six-month contract as a research assistant at Shell Research Ltd. This work involved mathematical modelling and then designing and developing software to simulate, and hence predict, the extent of explosions within confined spaces, such as are found on an oil rig. I then chose to go back into education and did a PhD at Liverpool University in collaboration with Shell Research Ltd. My specialist field was the "Application and Analysis of Dissipative Particle Dynamics", a technique for performing fluid dynamics simulations of complex fluids. After completing this, I worked as a software engineer at NASoftware. Although I enjoyed the work, there came a time when I had to move on, so I came here.

I live in the fantastic city of Liverpool, home to the mighty Liverpool FC. One of my favourite ways of wasting time is going to gigs and when I'm feeling adventurous, promoting them. For the record, the best Liverpool bands of the last few years, for whom fame and fortune have so far remained stubbornly elusive, are Ricky Spontane and Flamingo 50. World domination beckons and you heard it first in CSAR Focus. I have recently completed a course in Comedy Scripting but am still waiting for Hollywood, the BBC or Little and Large to start begging for my services. Any day now, I'm sure.

Forthcoming Events

July 2001

15-18

2nd International Global Grid Forum Meeting, Vienna, Virginia, USA
http://www.gridforum.org/Meetings/GGF2/

25-27

International Conference on Information Visualization (IV2001) London, UK
http://www.graphicslink.demon.co.uk

Aug 2001

28-31

Euro-Par 2001, University of Manchester, Manchester, UK. www.man.ac.uk/europar

9-10

The fifth Globus Retreat, San Francisco, USA http://www.globus.org/about/events/retreat2001.html

Sept 2001

4-7

Eurographics 2001, Manchester, UK http://www.eg.org/EG2001

4-7

Parallel Computing 2001 (ParCo2001), Naples, Italy. http://139.174.100.167/

10-11

Second UKHEC Annual Seminar "Next Generation HPC Systems and the Grid" Edinburgh, UK
http://www.ukhec.ac.uk/events/annual2001/

24-28

5th International Conference and Exhibition on HPC in the Asia-Pacific Region, Queensleand, Australia http://www.qpsf.edu.au/hpcasia2001

Oct. 2001

8-11

Cluster Computing 2001 - Third IEEE International Conference on Cluster Computing, Newport Beach, California, USA.

Nov 2001

10-16

SC 2001 - Supercomputing 2001, Denver, Colorado, USA. http://www.sc2001.org

12

GRID 2001 - 2nd International Workshop on Grid Computing, Denver, Colorado, USA.
http://www.gridcomputing.org/grid2001

Announcement

High Performance Computing Workshop on HPC and SGI NUMAflex

8 - 12 October 2001

MANCHESTER

www.man.ac.uk/mrccs/summer_school/2001

Manchester Research Centre for Computational Science (MRCCS) is pleased to announce a 5-day residential workshop on High Performance Computing and SGI's NUMAflex architecture.

Programme

The workshop will cover the following main topics:

* SGI NUMAflex architecture
* SGI compiler technology
* Porting Cray-T3E codes to the Origin 3000
* Origin 3000 optimisation
* Parallel performance on Origin 3000 systems
* IO performance tuning
* Scientific libraries
* Development tools
* SN-IA architecture and roadmap

The practical sessions will make use of a 512-processor Origin 3800 system located at the University of Manchester. World-class experts will be on hand to advise you how to get the best performance out of Origin 3000 systems, so bring your own code!

Speakers

The workshop will feature speakers from SGI (California) and a major overseas Origin 3000 installation.

Registration

Registration is via the on-line application form: http://www.man.ac.uk/mrccs/summer_school/2001/appform.html.

Course Fees and Accommodation Costs

The course is free for all registered users of the CSAR service. For attendees not associated with CSAR, the standard course fee of 750 pounds sterling applies to employees of commercial organisations, and a reduced course fee of 200 pounds sterling applies to academics. CfS has made available a number of grants for academic attendees to cover course fees. If you wish to claim one of these, please apply as early as possible as places are limited. Please note that the course fee does not include accomodation costs. For further information, please contact the CSAR Helpdesk by email (csar-advice@cfs.ac.uk), or telephone +44 (0)161 275 5997/6824, fax on +44 (0)161 275 6040.

EUROGRAPHICS 2001

Manchester
United Kingdom
virtual environments
computer games
visualization

Eurographics 2001

"Challenges in Computer Graphics for the 21st Century"

Manchester, UK

September 4th – 7th 2001

http://www.eg.org/EG2001

KEYNOTE SPEAKERS

Chris Hecker

- **Why Games Will Be the Pre-eminent Art Form of the 21st Century**

Markus Gross

- **Are Points the better Graphics Primitives?**

Holly Rushmeier

- **Rendering: Input and Output**

As we progress forward into the 21st century, computer graphics is becoming an integral part of people's visual thinking skills. Displays range from personal digital assistants to huge on-to-the-wall projection systems. Rapid advances in computer games are driving forces for bringing high-end graphics technology to consumer platforms. These applications and technologies add new content and perspectives to the future of computer graphics and related interactive techniques. With this viewpoint in mind, "Challenges in Computer Graphics for the 21st Century" chosen as the theme for Eurographics 2001.

Within this overreaching theme, there are three sub-themes:

- Visualization,
- Virtual Reality, and
- Computer Games.

All present challenges and opportunities for researchers and practitioners in our field.

Download Conference Timetable (pdf file)

To register, please complete the online registration form (one for each delegate).

Alternatively, you can download the registration form in Adobe Acrobat (PDF) format, print and complete it, and return it with your payment to:

Conference Secretariat,
Manchester Conference Centre,
Trading Services, UMIST,
P O Box 88, Manchester, M60 1QD, UK
Tel: +44 (0)161 200 4068
Fax: +44 (0)161 200 4070
Email: mcc.reg@umist.ac.uk

Announcement

The University of Manchester is delighted

to host the 44th annual

CUG Summit 2002

20 - 24 May 2002

MANCHESTER, UK

http://www.cug.org

The Cray User Group (CUG) is an independent, volunteer-organized, international corporation of member organizations that own or use Cray or SGI computer systems with emphasis on high-performance, technical computing, and visualization.

CUG's mission is to provide the high performance computing community with leadership and information exchange to enable the development and effective use of Cray and SGI computational tools in achieving the business and research objectives of CUG Members.

Each CUG SUMMIT has speakers from Cray, SGI and the user community on a variety of topics:

- General Sessions - giving the latest corporate information and directions.
- Interactive sessions - which provide a forum between attendees and SGI or Cray.
- Parallel Technical Sessions - provide the opportunity to focus on specific knowledge domains of the Special Interest Groups (SIG's).
- Tutorial Sessions - the opportunity to update your technical skills with the help of experts from SGI, Cray and CUG sites.

There are also a series of informal receptions and luncheons where you will have the opportunity to exchange information with your colleagues from other CUG sites and to collaborate with representatives from SGI or Cray Inc.

Attendance

Attendance is restricted to CUG sites and their users. For full details please see the CUG website.

2000 CSAR User Survey Results

Abdul Ali Rasheed,
CSAR Frontline Support

The second annual CSAR User Survey was conducted between 1st and 20th December 2000 and subsequently extended to 19th January 2001 to allow maximum participation. The online survey form comprised of 15 carefully designed questions to cover a range of aspects of the service, including systems, helpdesk and support services, training, feedback mechanisms etc.

In total, 34 completed forms were received, representing around 7% of Class 1/2/3 users. While this figure is low, it is a higher rate than for any other feedback mechanism such as ULF and service Quality Tokens. Survey submissions were entirely anonymous, although users were given the opportunity to provide their name on the form. 16 people did this.

The views expressed in this year's survey are very similar to those for 1999. In particular they show that the CSAR user community is generally satisfied with the service provided, as is evident from fig. 1, which shows the levels of satisfaction for the Overall service.

In summary, over 74% replied in the top two categories (Good or Very Good) while only 3% view it as poor. A full report of the survey was published in May and is available at: http://www.csar.cfs.ac.uk/admin/reports/user_surveys/.

Here we highlight just a few of the results. Figure 2 shows the satisfaction with various systems aspects (all the CSAR systems). For all aspects of the systems, over 83% of those who expressed an opinion are very or fairly satisfied. The most satisfaction is with service availability (more than 94% very or fairly satisfied) and archive (HSM/tape) facility (100% very or fairly satisfied).

The least satisfaction (68% very or fairly satisfied) is with job turn around time. However, the majority (more than 93%) are satisfied (very or fairly) with the temporary disk and more than 79% are satisfied with the job time limits. The response shown in the survey is quite satisfying given the fact that the turing has been very busy in the recent times. A number of steps have been taken to keep up and further improve the level of the service. This includes upgrade of fermat and introduction of green (256 processors SGI Origin 3000) under the *Technology Refresh Scheme*. Green has been configured purely to run as a batch engine with no interactive use. These measures should not only share the job load on turing but also help to improve the job turn around.

91% of the total respondents said that CSAR Service has contributed to the advancements in their research, out of which 71% would not have carried out their research without using the CSAR Service.

We would like to thank all those who participated in the survey. We will continue to improve the service in the light of the user's views. The next user survey will be in December 2001. In the meantime, please feel free to contact us on any service related issues. Your feedback is the best source to ensure that we are providing the best service.

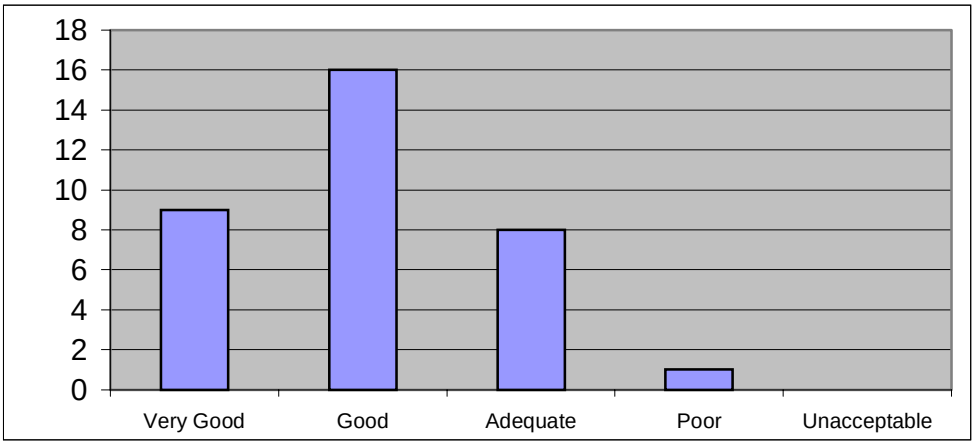

Figure 1: View on the CSAR Service

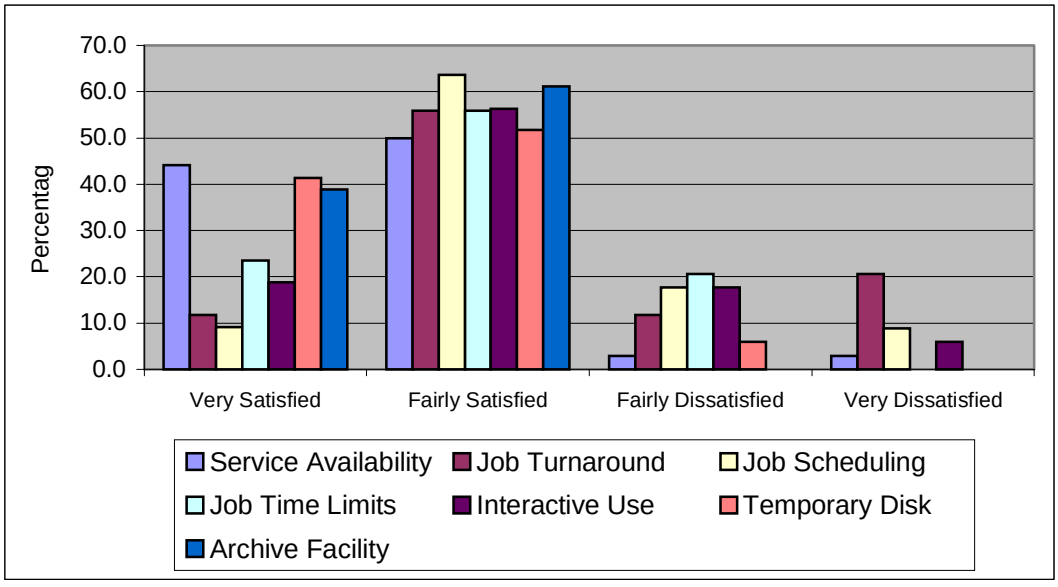

Figure 2: Satisfaction with various systems aspects (as percentage of those expressing a view).

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 published twice a year and is also available to view on the Web: http://www.csar.cfs.ac.uk/general/newsletter.shtml.

CSAR News Pages

All aspects of the service are described via these web pages (http://www.csar.cfs.ac.uk/general/news.shtml), and we also provide additional information via a monthly bulletin by email each month, particularly on new developments and other events associated with the service.

Getting Help

If you require help on any aspect of CSAR, you can contact the CSAR Helpdesk team who will deal with your query promptly and efficiently.

CSAR Helpdesk 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 automatically be logged using the latest call logging system (Remedy Action Request System).

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

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

Contacts

University of Manchester:

Terry Hewitt, *Head of CSAR Scientific Support/Director of Manchester Visualisation Centre*
email: w.t.hewitt@man.ac.uk tel: 0161 275 6095

Dr Mike Pettipher, *HPC User Services Team Leader*
email: m.pettipher@man.ac.uk tel: 0161 275 6063

Dr John Brooke, *CSAR Special Projects Team Leader*
email: j.brooke@man.ac.uk tel: 0161 275 6814

Dr Stephen Pickles, *CSAR Technology Reresh Team Leader*
email:s.pickles@man.ac.uk tel: 0161 275 5974

CSC:

John Rawlins, *Director of Supercomputing Operations, CSC*

Gerry Todd, *HPC Services Manager, CSC*