ISSN 1470-5893

CSAR FOCUS

Edition 6 Winter 2000, Editor: Kaukab Jaffri, University of Manchester

CSAR at the SC2000 Exhibition/Conference
Dallas, Texas

6 - 10th November 2000

Content

Editorial

This issue of *CSAR Focus* brings you the much awaited news about the CSAR Technology Refresh. Our Technology Refresh Team Leader, Dr Stephen Pickles, updates you on all you need to know on this exciting phase in the development of the CSAR Service.

Professor Ron Perrot, chairman of the Technology Watch Panel on High Performance Computing, gives views and proposals on the way forward for High End Computing in the UK.

We also report on CSAR's recent trip to the SuperComputing 2000 Exhibition and Conference, in Dallas, Texas, where we displayed work done by the CSAR community, and for the second year running, took part in metacomputing experiments.

We have reports on two major events which were held here at the University of Manchester during the summer: the MRCCS HPC Summer School as well as the Cray/SGI MPP Workshop. Mike Pettipher also reports on the 9th ECMWF Workshop, which he attended in November, on the Use of HPC in Computing Meteorology - Developments in Teracomputiing.

Prof. Mike Gillan and Dr David Bowler from UCL are both CSAR users who are working with a computational chemistry code called CONQUEST, and their article is featured on page 11.

As usual, please feel free to contact me if you have any comments or suggestions about *CSAR Focus*, which is produced twice a year. The next issue will be published in July 2001, for which articles may be submitted at any time.

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

Front page picture (L - R): Terry Hewitt(CSAR), Kaukab Jaffri(CSAR), Michael Robson(Networking), John Brooke (CSAR).

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.

CSAR'S Technology Refresh: What will it mean for you?

Dr Stephen Pickles,
CSAR Technology Refresh
Team Leader

The CSAR service is due for a mid-life technology refresh. A next-generation NUMAflex system, supplied by SGI and equipped with processors from Intel's Itanium family, is to succeed the existing Cray-T3E 1200E "turing" as CSAR's flagship computing facility. The system, to be known as "napier", will be installed in an initial 32-processor configuration in summer 2001, and scaled up incrementally during the remainder of the year. By the end of December 2001, napier will offer sustained performance of nearly twice that of turing in its original configuration. Napier will also be able to run MPI applications with total memory requirements significantly larger than turing can satisfy today. Napier will run the Linux operating system.

Eventually, turing will be retired from service. The date is yet to be determined, but the writing is on the wall.

Some of us will miss the reliable, old workhorse that we are just beginning to understand after years of acquaintance. Some will be daunted by the prospect of porting codes from turing to napier and optimising them yet again — keeping up with Moore's Law can be hard work! Others will have guessed correctly that there will be no more significant upgrades to turing, and are wondering what CSAR can do to arrest the growth of turing's batch queues between now and the advent of napier. Everybody will want to know a little more about napier and what to expect of it. It is these issues that this article is intended to address.

Napier

Like turing, napier will have modern compilers (fortran, C and C++), optimised scientific libraries (BLAS, LAPACK etc), and message-passing software (MPI and SHMEM). But there will be differences. Some codes will compile cleanly and run correctly first time.

Others will require more work.

Things to look out for include dependence on the more obscure Cray libraries (such as the Cray benchlib routines and certain parallel fourier transforms), implicit assumptions about the size of variables (a fortran REAL is 8 bytes long on Cray systems but 4 bytes on most others) or the behaviour of certain MPI routines (MPI_SEND buffering policies, and "features" in the Cray implementation of MPI groups). Although the Cray-T3E concepts of stream buffers and E-registers will have no meaning on napier, most of the effort you invested in single processor optimisation for the Cray-T3E will not be wasted on Itanium, and the massive on-chip caches are sure to help.

Unlike turing, napier will not lock you in to the message-passing programming paradigm. There will be support for shared-memory parallelisation, including OpenMP. Linux, too, will bring benefits in the form of a wide variety of software, utilities and performance analysis tools, of which many will be open source. The familiar Totalview parallel debugger will be available on napier, along with new profiling utilities that will take over from Cray's Apprentice and PAT.

Space does not permit a detailed exposition here of all the technical issues involved in code porting and optimisation. However, CSAR is developing training courses, workshops and on-line documentation that will explore these topics in depth. You are encouraged to review these materials when they become available, and to seriously consider setting aside a day or two to attend some of the new courses that will appear in CSAR's calendar for 2001.

Smoothing the transition

CSAR has scheduled various new machines and upgrades in order not only to smooth the transition to napier, but also to help alleviate the pressure on turing's batch queues in the interim by providing additional processing power. These measures include

- a new Itanium cluster "fourier",
- a major upgrade to fermat,
- a new Origin 3000 system "green",

not to mention a variety of exciting guest systems which will extend the core CSAR systems to a veritable smorgasbord of alternative architectures.

Fourier

An SGI-supplied cluster based on pre-release Intel Itanium processors and a Myrinet interconnect was installed on September 6, 2000, making it the first system of its kind in Europe. Over the coming months, fourier, as the cluster is known, will increase in size to 32 processors (16 nodes x 2 processors/node). Like napier, fourier runs the Linux operation system. The primary use of fourier will be as a vehicle for porting users' application codes to Itanium and Linux, and single-node optimisation of the same codes on Itanium. A secondary use will be small-scale production work in 2001, for applications that can tolerate the relatively slower communication latencies of Myrinet. To date, it has not been possible to grant unrestricted access to fourier, owing to Intel Non-Disclosure Agreements. Instead, early access has been contingent upon a careful process of screening and selection, and the user's signature on a form that effectively extends the University of Manchester's Non-Disclosure Agreement to encompass the user. The options available now to users wishing to make an early start are therefore either to apply for selection as an early user, or to arrange for CSAR Applications Support staff to do the porting on their behalf, through the CSAR support tokens scheme.

Fermat Upgrade

By the time you read this, the existing Origin 2000 system "fermat" will have been upgraded to 128 MIPS R12000A 400 Mhz processors and 128 GBytes of memory—that's four times more memory per processor than on turing. The use of fermat for small to medium scale production work is now encouraged.

The new policy will be of special benefit for users with applications that map more naturally to shared-memory programming models than to the message-passing paradigm.

Green

In addition to the fermat upgrade, a new 128 processor Origin 3000 system is due in March 2001. The new system will be known as "green" in honour of the English mathematician George Green.

Although fermat and green are based on MIPS processors and the Irix operating system (instead of Itanium and Linux), they have a very real relevance to the migration from turing to napier. This arises from commonalities in systems architecture (NUMAflex), interconnect (NUMAlink) and SGI's compiler technology shared between Origin 3000 systems and napier itself. A code which compiles cleanly, runs correctly and scales well on an Origin 3000 will be better prepared for napier than one which has been tested only on a Cray-T3E.

Ways forward

It is entirely possible to wait for the arrival of napier and then port to Itanium family processors, SGI compilers, and the Linux operating system in one fell swoop. But the availability of fourier and green in advance of napier opens up several alternative paths for migration — an early start and a phased, incremental strategy will see you well placed to take full advantage of napier when it becomes available. Compatibility with SGI compilers can be tested on fourier, green or even fermat. Fourier is the vehicle of choice for single-node optimisation work on Itanium processors, but green will be more appropriate for tuning parallel scalability.

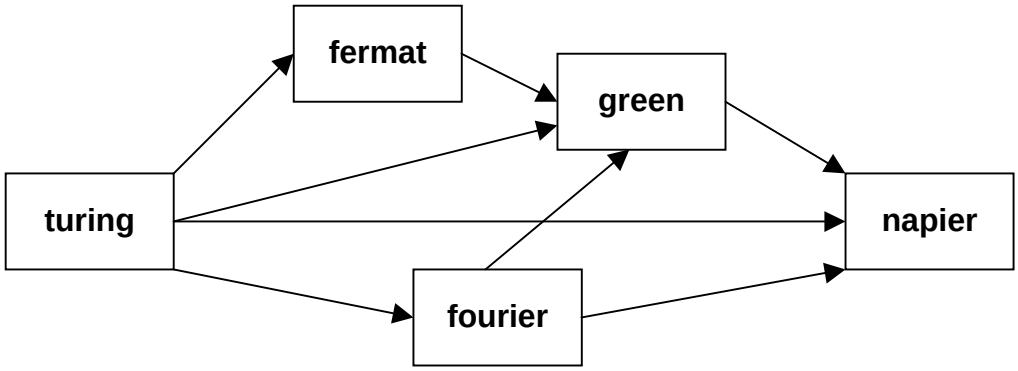

Alternative routes for migration from turing to napier

Planning ahead

So just how much work will the migration from turing to napier involve?

The short answer is that we don't know, and we need your help to find out. You can assist us by participating in the forthcoming Technology Refresh Technical Survey and answering the questions to the best of your ability. Principal Investigators can help further by reviewing their projects' capacity plans, giving special attention to requirements for training and support in the light of napier's imminence — if additional staff will be needed, the sooner we know, the better. CSAR is committed to making the transition to napier as smooth and painless as possible. With your cooperation, we will do just that.

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.

UK Terascale Computing

Prof. Ron Perrott
Technology Watch Panel,
Queens University,
Belfast

Introduction

Three years ago the inter-council High Performance Computing Strategy Committee (HSC) established a Technology Watch Panel (TWP) whose terms of reference include providing advice on high performance computing (HPC) technology requirement for ***internationally competitive research in the UK***, including developments in hardware, software and tools. In addition to talking to vendors and users one of the methods used by TWP to fulfil its terms of reference is to visit other countries on a fact-finding mission. Two years ago TWP visited the Accelerated Strategic Computing Initiative (ASCI) sites in the US and as a result recommended the UK purchase a Terascale computer – preparatory work on that procurement has now started. A second TWP mission to the USA recently took place (October 30 – November 2, 2000). The focus of the current trip was on the recent US Terascale computing procurements and GRID-based developments. The sites visited included the Argonne National Laboratory (ANL), the NSF Partnership for Academic Computational Infrastructure Centres (PACI), (namely, the National Centre for Supercomputing Applications (NCSA) at Urbana-Champagne and the San Diego Supercomputer Centre (SDSC)), the headquarters of the DOE, NSF and DARPA (Washington), and NASA Ames, California.

At present in the US there are several well-funded initiatives in the area of grid computing and HPC. The following summarises some of the observations of the TWP and the recommendations made to HSC to ensure that the UK remains internationally competitive.

Observations

1. Grid

Grid activities and projects are proceeding at a substantial pace; it has caught the imagination of the funding agencies and the research community is responding. The US seems to be well ahead of the UK in practical experience and projects. The level of partnership and co-ordination is high. It seems that Universities can join a NSF PACI Center grid by donating part of the machine time of their particular machine to be used by the grid.

2. UK–US Transatlantic Terabit Testbed

There is on-going interest in the proposal of a UK-US Terabit Transatlantic Testbed project. In order to progress the project we need to identify an application, which needs such a network.

3. Access Grid

This is a collaboration environment, which provides multiple high-resolution projected displays from cameras at multiple sites, using mostly off-the-shelf equipment. It is taking off fast and is becoming the medium of choice for co-working across the PACI Centres. There is a high demand among scientists. Access Grid is the vehicle for Chautauqua's (popular community meetings to tell people about new technologies), distributed training and collaboration meetings. The NCSA currently has 20 sites using a standard multicast backbone and intends to package the software in standard form for anyone to use. ANL is a key player. There are plans for a site in Japan, Australia and a couple of sites in Europe.

4. National Science Foundation (NSF)

In 2000, Pittsburgh Supercomputer Centre (PSC) won NSF's first competition for a Terascale computer, being awarded $36M for equipment plus $9M for support over three years. Five proposals were reduced to a shortlist of three. PSC's proposal was for a 6 Teraflop/s peak Compaq Alpha system using the Quadrics switch (682 Compaq AlphaServer ES41 nodes with 2,728 EV68 2.2 Gigaflop/s CPUs). A 64 x 4-processor EV67 processor system, with a peak performance of 340 Gigaflop/s, is being installed now. The full system, with federated switch, will be delivered in summer 2001. The choice of 4-processor nodes was designed to meet NSF user requirements. The machine will be a programme-wide resource, supported jointly by the two PACI centres and accessed through the single web page interface to all PACI resources.

In response to Presidential Information Technology Advisory Committee (PITAC), which recommended two machines, NSF has allocated a further $45M for a second Terascale machine procurement.

5. Accelerated Strategic Computing Initiative (ASCI)

In general the ASCI programme is on target with both its hardware and software projects, for example, the 100 teraflop/s system is on track for 2004. We were given more details on the organisation of the programme; there is a central directorate of 12 people with a staff of two at each national laboratory, they meet at regular intervals as well as by conference calls.

6. Presidential Information Technology Advisory Committee (PITAC)

The PITAC committee is still operating and developing programmes in line with its initial report. Recently PITAC has produced an open software report for the area of high performance computing; it is available on the web at http://www.ccic.gov.

7. Collaboration between computer and application scientists.

The US DoE within the ASCI programme has established at each of its national laboratories (Sandia, Livermore, and Los Alamos) an Institute of Computer Science. They are funded at the rate $6M per year and they have been operating for about 18 months. The objective is to strengthen computer science as a discipline. In essence the computer scientists are given scope to do whatever research they wish but the hope is that they will become interested in the applications that are currently being undertaken at the laboratories. Academics from outside the laboratories are encouraged to get involved. They are pleased with the results so far.

8. Interagency Co-ordination

There is a high level of co-ordination between agencies. The High End Computing and Communications Co-ordinating Group (or HECWG) fosters co-ordination and co-operation among the US government agencies involved in high-end computing through monthly meetings. Initially they will identify areas of overlap and duplication; areas of strength in an agency will be used as the focal point for that activity. They have held several meetings and are producing a strategy document.

9. New Proposals

i) SciDAC

Scientific Discovery through Advanced Computing (SciDAC) is a new programme to replace the Strategy Simulation Initiative, which Congress did not fund. SciDAC is a programme of the Office of Science. It promotes a co-ordinated approach to applications development to produce scientific community codes. Support for hardware infrastructure includes

a) Flagship computing facilities to provide generic high-end computing;

b) Topical facilities to provide high end dedicated resources for specific communities;

c) Experimental computing support for computer science and experimental technology.

ii) Biosciences

It has been difficult to get the National Institute of Health (NIH) to support computational work, however, Larry Smarr (San Diego) has co-chaired an NIH committee which recommended the creation of 6-8 centres of excellence in e.g. the virtual human, bioinformatics, and real-time medical imaging. However, NIH has not yet accepted the recommendations. Biology is the fastest growing part of the NSF computational programme.

10. Industrial promotion

At the NSF PACI Center in Illinois there is a Partnership with Industry programme known as the Private Sector Program.

The goal is to collaborate with the commercial sector to guarantee that the work of the Center can be applied to real world challenges faced by America's business. Only one industrial partner per industry is accepted into the programme to guarantee proprietary business advantage. Currently there are 9 companies registered in the programme. The companies pay to join the programme approx. $0.5M per year and each company is asked for a 3-year commitment. The Center has approximately 10% of its machine cycles available to be used as part of the Partnership programme.

Recommendations

1. High end computing

A rolling programme of high-end procurements for capability computing should be implemented. This should include Flagship computing facilities – the very high end – as well as Topical and Experimental computing facilities. These latter facilities are linked into particular communities, for example, QCD, computer science, protein folding etc. Hopefully this would be of benefit to all Research Councils, and prepare for the new architectures likely to emerge in the second half of the decade.

The criteria of success should be to provide enabling technology which "makes a difference" in terms of scientific discovery. Modest goals for the Flagship facilities based on the US high-end strategy are:

- 5 Teraflop/s peak in 2001/02;
- 20-25 Teraflop/s peak in 2004/05;
- 50-100 Teraflop/s peak in 2007/08.

2. Clusters

An initiative to develop and evaluate commodity clusters should be implemented. The main objectives would be:

- to inform future high-end procurements and local initiatives;
- to bring on new high-end applications, e.g. in the bio-sciences;
- to work jointly with the US PACI centres.

3. Grid

In order to exploit the full potential of the grid a Partnership between institutions should be implemented. This should involve international collaboration via:

- Access Grid as the primary vehicle for joint working;
- workshops and co-ordinated participation in international grid fora;
- exchanges with US PACI centres and NASA.

The Partnership should bring some order to UK HPC provision through leveraging JIF/JREI/SEI/SRIF facilities by

- enhancing local resources in exchange for joining the grid;
- promoting common software standards;
- supporting local facilities more efficiently than at present.

Thus there should be strong incentives for joining the Partnership, but no compulsion.

My thanks to the other members of the TWP, namely, Professor R Kenway, Dr Martyn Guest and Professor Tony Hey.

CONQUEST: A Quantum Leap in Atomic-Scale Simulation

Dr David Bowler, Prof. Mike Gillan
University College London

Simulations of matter on the atomic scale based on quantum mechanics have made a colossal impact on a wide range of basic sciences over the past 10 years, including physics, chemistry, biology and earth sciences, and the CSAR service and its predecessors have been crucial for this work in the U.K. All kinds of scientific problems have been tackled, ranging all the way from chemical reactions on surfaces to the physics of the Earth's core.

But up to now, the algorithms have suffered from a major bottle-neck, because the computer effort needed increases too fast with the number of atoms being simulated. This has meant that large systems containing thousands of atoms, which are essential for studying complex materials, nanostructures or biological systems, have been out of reach. This bottle-neck has now been well and truly broken. With the CONQUEST code running on the Cray-T3E turing, practical quantum calculations have been achieved on systems of over 16,000 atoms, and the results show that the scaling with numbers of atoms and processors is excellent, even for calculations using nearly the entire machine.

The CONQUEST code has been developed over the past five years, with EPSRC support via the U.K. Car-Parrinello consortium, one of the largest consortia working at CSAR. The methods are based on those used in the well-known CASTEP code, but the key difference is that the computer effort needed by CONQUEST is strictly proportional to the number of atoms, instead of CASTEP's increase as the cube of this number.

Both codes perform fully ab initio quantum calculations, using the density-functional theory of many-electron systems, and they can achieve similar precision. The CONQUEST code relies heavily on the parallel manipulation of sparse matrices, and CSAR support in optimising that part of the code was vital in making the code run efficiently on turing.

As an illustration of the excellent scaling with number of atoms, the graph in figure 1 below shows timings for calculations on silicon systems containing from 4096 to 12288 atoms run on 512 processors of turing. Remarkably, the scaling with number of atoms is actually slightly sub-linear. Other tests show that, in terms of Amdahl's law, the unparallelised part of the code is completely negligible when we run it on turing. With further code improvements now on the way, we are optimistic that ab initio calculations on systems of up to 50,000 atoms will be feasible in the near future - an increase of nearly two orders of magnitude over what was possible before!

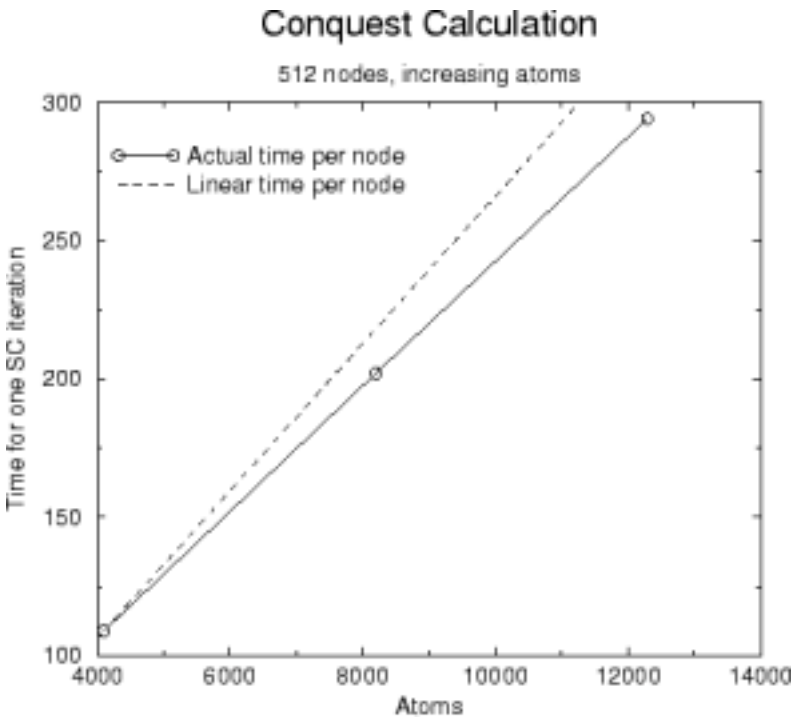

Figure 1: Time for CONQUEST calculation on systems of silicon atoms containing from 4096 to 12288 atoms run on 512 processors of turing, showing the slightly sub-linear increase of time with number of atoms.

Report on SuperComputing 2000 Conference/Exhibition: 6-10th November 2000, Dallas, Texas

Kaukab Jaffri,
CSAR User Services

This year, the SC2000 team consisted of Terry Hewitt, Dr John Brooke, Kaukab Jaffri from CSAR, and Michael Robson from the Networking team at Manchester Computing.

The preparations began two and a half months before the event, with meetings held each week. John Brooke was leading the team as well as co-ordinating the GRID experiment (see following article by John Brooke). Michael was our networking specialist and was working with John to support the metacomputing experiment. I was primarily involved with the overall planning of the stand, poster and publicity material, and making sure the equipment and ourselves got to Dallas on time!

Terry Hewitt, HPC User Services Manager and Director of Manchester Visualisation Centre also attended SC2000. Ian Curington of AVS Inc. and Terry jointly presented a one day tutorial entitled "Frameworks for Visualization of Large Data Sets", which was very well attended.

Our ten hour flight got us into Dallas on Friday afternoon - we had left Manchester around 6am on Friday morning! The most stressful part of our the journey was waiting for our luggage as we had taken the risk of not shipping our posters over, but putting them in big poster tubes and taking them as luggage! The thought of the posters not turning up was enough to make me travel sick, as we would not have had anything to display on our stand! Thankfully, our three tubes were flung out of the aircraft onto the conveyer belt and appeared with the rest of our luggage, and we all breathed a sigh of relief!

Saturday morning came, and we had all managed to get a good nights sleep. The hotel had provided coaches to take all the exhibitors to Dallas Convention Centre. These ran to and from the venue all day. On arrival at the Convention Centre, I was amazed at the size of the venue! It was an enourmous warehouse, with lots of little trucks, massive cranes and forklift trucks cluttering the place, it looked like a building site inside a building!

The larger vendors such as IBM, Cray, Hewlett Packard, ICL and Compaq, as well as our partners, SGI and CSC were building enormous extravagant stands, as it is at exhibitions like this that they are able to show off their new products and gain business, However, we were there as a research exhibit, and therefore our stand was somewhat smaller and less prominent that these! See Figure 1.

Figure 1:
This is what our stand looked like when we first arrived

By Sunday evening, our stand was ready. We has put the posters up, set up the O2 and Octane machines that had been loaned to us by SGI, got our network connections and most importantly found out where the free coffee and biscuits were. See Figure 2.

Figure 2: Our completed stand

On Sunday evening we were invited to a ranch party were there was a buffet and line dancing, as well as armidillo racing. This party was held for all the exhibitors at SC2000.

Monday through to Thursday was when the actual exhibition and conference took place. Mike and my role was quite simple, to stand at our stand and tell people about our services. As well as work done by CSAR, our posters included other national services provided by Manchester Computing such as MIMAS and Web of Science. John Brooke worked very hard thoughout the week, and unlike us, did not take the opportunity to go exploring round Dallas. We were given the opportunity to go shopping on two occasions, but found that the nearest shopping centre was 40 minutes away by train!

Although both our partners, SGI and CSC have branches in the US, our colleagues from the UK also attended SC2000. Paul White and Roger Greatrex from CSC in Farnbourough, and Crispin Keable and John Fleming from SGI in Reading, all of whom are part of the CSAR Service, attended the event.

We had a variety of people coming to our stand, many were interested in the VIPAR work that CSAR has done recently which involves parallel visualization techniques. We had a lot of interest in the the metacomputing and eGRID experiments, and we also managed to run out of all of the CSAR Focus's that we took with us!

There were lots of other research stands like ours, including some from differerent universities around the world, involved in similar work to ours. The opportunities to network and make contacts at this event are vast.

On the whole, SC2000 was an interesting event to be part of. It is the place to be if you are involved in supercomputing as it has everything you ever wanted to know, from hardware to software and compilers to petaflops.....and of course there are the lavish parties in the evenings!

Figure 3: Terry giving a presentation at SGI's stand

CSAR and UKHEC at SC2000

Dr John Brooke,
CSAR Special Projects Team Leader

The UK HPC community was well represented at SC2000, with a substantial CSAR prescence on the University of Manchester research booth and at the EPCC booth. If you add to this participation in live demonstrations in the NetChallenge competition and visualization demos of CSAR work at the SGI booth, the UK measured up well on the European scale of things. However, this should be seen in perspective, although the University of Manchester had one of the largest European Research booths, we were completely dwarfed, and occasionally deafened, by the presentations at the ASCI booth next door to ours.

Our contention is that what we lacked in size, we made up in quality. CSAR presented two ground-breaking collaborations at SC2000, the largest metacomputing experiment ever (more of this later) and one of the most imaginative GRID applications, the eGRID demonstration of the infamous "Cactus worm". We also showed one of the most imaginative applications of visualization, a trip through the earth's mantle based on a collaboration between CSAR visualization consultant Jo Leng and Huw Davies of the Terra group. The ability to show fully the three dimensional nature of the structures revealed by geoseismology, in a setting that reproduced the curvature of the geometry of the mantle and an ability to zoom in and out and reorient about any given axis, impressed even the SGI and AVS personnel who are used to seeing the latest advances in visualization technology.

We were also able to describe this work in the context of the UKHEC consortium, since this was one of the visualization case studies carried out in UKHEC's first year. CSAR and EPCC took along the latest copies of UKHEC news which featured seminars in OpenMP, computational GRIDs and Beowulf clusters and seminars in Object-Oriented programming and Java.

At the UoM/CSAR stand, we noticed that this year (compared to our first visit in 1999) we were much more involved in discussions with our research peers. The metacomputing experiments performed last year have put us on the map, and our involvement with European GRID projects lead to many interesting conversations with similar projects in the US and Japan. The SCx series of annual events are by far the largest gathering of those involved in High End Computing and many technical and research breakthroughs are presented for the first time at these events. In Figure 1 you can see me in the process of giving the UoM/CSAR contribution to the SC2000 Webcast. The inane grin can be attributed to the fact that I was so nervous I could hardly hold the microphone and since I had not had time to prepare (I had just been giving the global metacomputing demonstration) I was extemporising my presentation. I do not recommend this procedure. The results will be available on the CSAR WWW site as soon as they are sent to us.

Figure 1: John Brooke being filmed for the SC2000 Webcast

Global Metacomputing

The successful presentation of a global metacomputer at SC99 (see CSAR Focus 4) was followed up by an even more ambitious presentation at SC2000. Two new sites were added, JAERI (the Japanese Atomic Energy Research Institute) and NCHC from Taiwan. We were now able to run 6 high-end computers, comprising 5 different architectures over a network which spanned the earth and the total power of the metacomputer was 2.4 Teraflops. The sites involved and peak Gflops rates were:

CSAR (T3E)	691 Gflops
HLRS (T3E)	460 Gflops
PSC (T3E)	460 Gflops
SNL (C-PLANT)	32 Gflops
ETL (SR8000)	512 Gflops
NCHC (SP3)	252 Gflops
JAERI(SX5)	8 Gflops

Figure 2: Presentation of experiment to the judges

We entered the Network Challenge which invited participants to compete for the largest amount of data transferred to the SC2000 exhibition floor during a period of one hour. We would have preferred to have the largest use of bandwidth criterion changed to most intelligent use of bandwidth since the intercontinental links are not designed to squander bandwidth but to use a scarce resource as efficiently as possible. Still, it would be unsporting to pass up a competion against the likes of Ian Foster and Carl Kesserlman (for example) and we aroused a great deal of interest in our work.

In Figure 2 you can see the presentation of live distributed collaborative visualization over the metacomputer in front of the panel of judges (the ones wearing stereo spectacles). Figure 3 is somewhat more cryptic. It shows the progress of the running of the Jodrell Bank DeDispersion code (developed by Stephen Pickles of CSAR and Stephen Ord of Jodrell Bank). The windows show grmap running on the T3Es so that the judges could see the code running simultaneously on all three machines. The smaller windows show the output of the program, as each processor receives a new work unit from the server and then signs off when finished. During the 4 days in which the code was run, timings of the performance of the data-processing and network speeds observed were recorded. This work will be reported back to the UK community via the Global Metacomputing project, see http://www.man.ac.uk/mrccs/global_supercomputing/

Figure 3 (right): Shows the progress of the running of the Jodrell Bank DeDispersion code

One of the representatives of the US Supercomputing Centres described the Jodrell De-Dispersion code as a "golden application" because it intrinsically demands metacomputing techniques and is an acute test of the status of wide area networks.

eGRID and the "Cactus worm"

Those involved in GRID computing will know that there have existed two organisations promoting GRID cooperation on a global scale. One was based in the US and was called the Grid Forum and another based in Europe was called eGRID. These have now joined together along with representatives from Asia and the Pacific Rim to form the Global Grid forum. One of the last independent acts of the eGRID forum was to prepare a demonstration of migrating GRID computing for SC2000. They gave demonstrations at various booths including Sun and HLRS. On Thursday morning they gave a demo at the UoM/ CSAR booth.

The demonstration was of the "Cactus worm". Cactus is an application portal for the solution of partial differential equations on distributed systems. See

http://www.cactuscode.org/

for details. A code solving the wave equation was run on sites across Europe using Globus as the GRID middleware to link the computers together. The "Cactus worm" migrated across this network while continually calculating and maintaining its state as it migrated across machines. This is a powerful demonstration of the use of the GRID. Imagine running a large computational simulation which fires off components for analysis or post-processing in an asynchronous manner as it runs. It is a waste of resources to demand that this is run on a tightly coupled system, due to the inherent asynchronicity, but such spin-offs could use spare capacity on machines across the GRID. The key point is that they are able to maintain coherency while permitting migration. The demonstration is that GRID middleware is essential to such operation as are GRID aware application launchers.

Ed Seidel of the eGRID forum gave a very interesting and coherent presentation despite being nearly drowned out by the presentation on the ASCI booth next door (see Figure 4). The Cactus WWW site gives more details of the application. The next meeting of the Global Grid forum is from the 4th - 7th March 2000 in Amsterdam, details from: http://www.gridforum.org.

Figure 4: Ed Seidel from the eGRID Forum giving a presentation on the "Cactus worm" application portal

Acknowledgements and conclusion

I would like to thank CfS and other sites for donating the computing time on the CSAR service to enable the global metacomputing to take place. Until the GRID technology advances to the point of being able to accomodate accounting and billing for on-demand use of supercomputing resources, these coupled experiments will rely on the generosity of organisations running large HPC Centres. We also acknowledge the loan of an Octane and O2 workstation from SGI as part of their value-added commitment to UK High-End Computing. The importance of good live demos can hardly be over emphasised.

It is these that attract visitors to the Research Booths and many useful contacts were made, which will be fed into the CSAR and UKHEC support for UK researchers next year.

I should also thank all of those who helped to make our Reseach Booth a success. Many people in CfS worked on the project and produced posters, publicity material and specialist systems support (thanks again to Paul White of CSC).

Special thanks are due to three people. Terry Hewitt, who helped to ensure that we got the resources we needed and came out of high level meetings to help with the assembly and disassembly of the booth. Kaukab Jaffri of CSAR supervised the organisation of the booth (a formidable task) and it all ran very smoothly. As part of the Global Metacomputing project funded by the UK Joint Information Systems Committee (JISC) we were able to take a networking expert, Mike Robson, of the Manchester Computing COS group. Mike was able to ensure that we were able to connect our loan machines to the SciNet network and provide a focal contact point which was used by ourselves and our colleagues from eGRID and UKHEC to connect to the global internet.

In this way we feel we justified the theme phrase of our Research booth:

"Manchester Computing: the UK's gateway to the world"

Report from HPC Summer School, 4-15th Septmember 2000

Dr Dan Kidger,
CSAR Application Support,

Manchester Research Centre for Computational Science

This year, a two-week 'Summer School' was hosted by MRCCS at the University of Manchester. Because of the way it is funded, it is free to all CSAR users. The themes are all related to aspects of High End Computing. Last year, the theme was 'Distributed Memory Programming and Scientific Visualization'. This year the topic was 'Linux for High Performance Computing'

Why 'Linux for HPC' ?

It is now well recognised that Linux offers a route to 'do-it-yourself' high-end computing, by means of building your own 'Beowulf'-class machine. A Beowulf is defined as a computer assembled from commodity, usually PC, components that can be bought from more than one vendor and assembled together yourself. There are many choices of hardware and networking components, and these are advancing at an extremely rapid rate. However since this is a build-yourself solution, it is relative easy to upgrade any part of the system. This is unlike a typical mainframe machine, which is tied to a single supplier. Our summer school however was on *Linux* rather than Beowulfs and so covered the bigger picture. Computer vendors like SGI and Fujitsu are well aware of the fact that Linux has many times the number of people developing it than they could ever commit of their own staff to their own brand of Unix. Companies like these now sell pre-built Linux based systems. Users of the CSAR machines should already be aware that the machine that will succeed turing next year will almost certainly be running Linux.

The People who Participated

Figure 1 shows some of the speakers, organisers and participants who attended the summer school. There were thirteen speakers. These were led by Dr. Doug Johnson from the Ohio Supercomputing Center and Niels Walet from UMIST. Both gave presentations and tutorials over several days. There were also lectures and presentation from: Bill Sparks (SGI, California), Werner Krotz-Vogel (PALLAS, Germany), Robin Harker (Workstations, UK), Jonathan Nash (Sychron), John Merlin (Portland Compilers), Richard Watkins (Platform Computing), Nick Davis (IBM), Martyn Guest (Daresbury), John Taylor (Quadrics), Rob Allan (Daresbury), and Einar Rustad (Scali/Dolphin, Norway).

Figure 1: Group photo

There were 18 people who registered as students for the summer school. As well as CSAR users, there were two delegates from Italy (CINECA), one from Holland and even one from Australia (CSIRO). There is also a strong commercial interest in Linux for Supercomputing - this year we also had delegates from Xinit, Atos and SGI.

CINECA (www.cineca.it/HPSystems) is in many ways the Italian equivalent of CSAR. They have a 256 processor Cray T3E, a 64 processor Origin 2000 and a StorageTek Tape Silo for HSM. Also, like Manchester they have a Virtual Reality theatre and an IBM SP. CSAR users may wish to compare some of their online documentation with CSAR's. Like CSAR they are due for a refresh of their technology next year and like ourselves expect a Linux-based system to replace their T3E; hence their interest in our Summer School.

The Lectures

There were 37 hours of lectures delivered over the two weeks. Some of the key topics were:

Why use Linux for HPC?: Linux is more portable than any other operating system, and has more developers world-wide than all of the commercial Unix vendors put together. Since it is free, it makes home-built 'beowulf' systems very cost-effective. Support from the big players such as SGI, Fujitsu and IBM (who have donated compilers and tools for free) has given Linux much credibility for High Performance Computing.

Building the 'MadDog' cluster: This talk gave a personal recollection by Neils Walet of the Physics department at UMIST of the highs and lows of building a high-performance cluster single-handed.

Captools: Developed at the University of Greenwich, this tool helps convert serial programs into parallel ones. It examines the existing source to look for opportunities to parallelise the code. This is particularly useful for those users seeking to use Linux clusters to speed up their serial codes.

Performance of Linux systems versus machines like the T3E. Martyn Guest showed results of the many benchmarks that he has run on Linux and other systems. Of the Linux machines, those with Alpha chips performed best followed by the AMD Athlon and then the Intel Pentium III. On several tests, these machines outperformed the T3E (running on 32 processors of each).

There were also talks by various computer manufacturers that sell complete Linux systems (IBM, SGI and WS-UK), make network hardware (Quadrics, Scali/Dolphin) or develop software for Linux (Platform Computing, Portland Compilers and PALLAS).

The Practicals

As well as the main practical project-based exercise (see below), there were 'hands-on' sessions on:

Vampir: A tool for profiling and tuning parallel programs, particularly MPI ones.

Totalview: An interactive parallel debugger with a full graphical 'point and click' user interface.

PBS: Installing and configuring the Portable Batch Queuing System.

PVM: A forerunner of MPI, but with better support for the heterogeneous systems often found in clusters of workstations.

CapTools: The students used this tool to identify and evaluate opportunities for parallelism in Fortran codes and then compared this with a hand-modified MPI version of the code.

Portland Fortran 90: using this compiler to produce fast executing code under Linux.

Figure 2 : Some real 'hands-on' experience

Building your own Supercomputer

This year we took on the rather ambitious project of getting the attendees to build their own 'supercomputer' from commodity parts. This acted as an ongoing exercise that lasted the full two weeks of the summer school. Each time a lecture was due to start, we had to go around and try and round-up the students who were always busy on their 'beowulfs'. We gave each student a pair of little rubber penguins in their conference bags. As an ice-breaker on the first day, the students had to find the other people who shared penguins with the same colours as theirs. These became the groups of three or four people who then could start to assemble their cluster, once they had 'claimed' their flag to put on top of their stack of PCs.

The PC's were from a student cluster from the School of Engineering that was being upgraded. A total of 34 Pentium100 PC's were lent to us. Each had 32Mb memory, 0.5Gb Hard Disk and a 10Mbit ethernet card, but no cd-rom drives. There were only enough keyboards, mice and monitors for one set per group of 4 machines.

These machines were 'raw' - most had no operating system at all, since they used to boot off the network. The first step was to remove the boot EPROMs from the Network cards - the first time many of the attendees had attacked a computer with a screwdriver. There then followed a sequence of a 49-point plan which Chris Muller had devised. This was a detailed set of instructions of how to install Linux on a PC and configure for use in a cluster with compilers, MPI, cross-mounted filesystems, etc. The machines were all connected via three joined ethernet hubs. By Friday all groups had got as far as running a parallel job across their four machines. In the second week, they joined the machines into a 'super-cluster' and so could test running jobs across the whole cluster of 24 nodes. This introduced the issue of the need for a batch queuing system to each group's jobs colliding with other users'.!

Computer facility Tour

We arranged a tour of the main machine room of Manchester Computing to se the *turing, fermat, fuji* and the new *fourier* 'in the flesh' . We also organised a demonstration of the new £750K Virtual Reality Theatre which will allow computational steering of parallel programs (http://www.man.ac.uk/MVC/research/VIP_Lab/).

Food and other Social Events

As well as filling the minds of the students we also aimed to fill their stomachs. There were breakfasts and three-course dinners included every day at the Hall of Residence. Lunches were taken in the adjacent Manchester Business School Eatery.

We organised a trip to the theatre at the new £98m Lowry Centre in Salford (www.thelowry.com) to see 'Cirque Eloise' - an animal-free circus performance, which proved very popular. The Hall bar also gave the attendees chance to meet informally with the speakers.

The Organisation

It took over six months of preparation to organise this summer school. I must thank all those people that helped with the organisation: Rob Allan and Chris Muller from Daresbury, Robin Pinning, Kaukab Jaffri and Caroline Evans from CSAR, plus of course all the other people not mentioned above.

What the Delegates thought about the Summer School

Enrico Luparini (Netherlands) writes:

" Excellent approach , excellent support and the pleasure to build a cluster from the scratch. The only negative side was that it wasn't possible to get Myrinet working but it wasn't your fault."

Alistairs Mills (SGI) writes:

"The summer school was most enjoyable. The program was ambitious. In the beginning there were six groups of three or four people, each group having four old PCs, and there was darkness all around. By the end of day 5, the PCs were configured as six Beowulfs each with four nodes..... By the end of day 8, the PCs were configured as a single beowulf with twenty four nodes. By breathing compilers and other software life forms into the system we had life to run a real application which ran for 24 hours....."

Stefano Martinelli (CINECA) writes:

"I think the Summer School was very useful, because it gave me the opportunity to see a wide range of aspects regarding Linux clusters. We could explore extensively hardware aspects and software aspects, both with our hands with our basic and very interactive practical experience, and listening to the technical experience of other experienced users."

Availability of the materials

The speakers notes have been printed in four bound volumes. Additional copies of these are available at a small cost to cover the printing and distribution. A cd-rom of the materials is currently under preparation. Contact csar-advice@cfs.ac.uk is you wish to receive a copy.

The website at www.man.ac.uk/mrccs/summer_school has some more information including various photographs that were taken.

Next Year's Event

Next year, the subject is likely to be "Programming on Napier: the new flagship CSAR machine" and will cover topics on Linux, OpenMP and the Intel Itanium processor. There will be two one week workshops, the title still to be confirmed.

e-hpc.com

Gerry Todd, CSC

We all understand the benefits of supercomputing. The ability to accurately predict and better understand how things react and change in different configurations and environments, to varying stimuli, is an essential activity for many businesses. Supercomputing delivers high quality results in short time-scales - whether forecasting climate changes, analysing financial investments or designing aircraft.

However, the complex problems and simulations required in today's business and research environments pose many challenges:

- How do organisations predict when, and what size capability they will need?
- How do scientists and researchers keep up to date with the latest technology?
- How can organisations gain access to the range of technologies which will enable them to gain a competitive edge without significant financial risk?
- How can organisations exploit the Internet as a tool for business advantage without compromising their commercial security?
- How can industry be kept better informed about the world leading research funded through the UK Research Councils and of the possibilities for exploitation?
- How can industry's problems be best brought forward for the attention of the UK academic/scientific research teams?

Amongst others, CfS has the challenge of providing "bridges to industry" for improved technology/research exchanges and liaison, training and education. CSC, in particular, is charged with providing facilities to enable commercial companies to share the CSAR systems with the academic community, to encourage the use of such leading edge HPC technologies, enable greater technology transfer and encourage industry to give better feedback and stimulate/sponsor more competitive and applied UK academic research.

Formally launched on 6th November 2000, and presented at SC'2000 in Dallas, e-HPC.com is CSC's vehicle (portal) to help address the challenges outlined above - accessible over the Internet on a pay-as-you-use basis. Here the services available in CSAR (SGI, CRAY & Fujitsu) have been extended through joint ventures with several leading HPC vendors (Compaq, Hewlett-Packard and NEC). This is in part to broaden the base of HPC capabilities available for industry/commercial usage - but also to add "Guest System" capabilities and generally extend the facilities for consortia to use within CSAR.

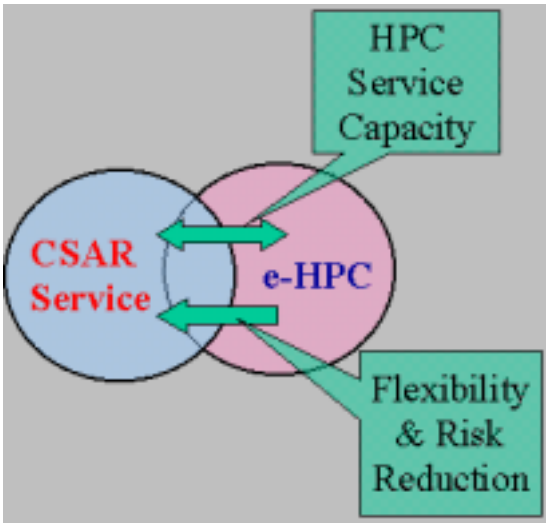

e-HPC.com thus makes high performance computing services available to any organisation that needs regular access to supercomputing facilities, or occasional access to specialist systems in periods of peak demand. The future may see such ventures linked into UK e-Science, data and other compute grids, and perhaps into visualisation grids to give interactive facilities at the end user workstations. If the e-HPC.com service can help any of the organisations you work with, please give CSAR a call…

Report: The 6th European SGI/Cray MPP Workshop

Dr Ben Jesson,
CSAR Applications Support

The Sixth European SGI/Cray MPP Workshop, an event aimed at application developers and end users of massively parallel processor (MPP) computer systems such as the SGI Origin 2000 and Cray T3E, was held in Manchester on the 7th and 8th September 2000. This conference, kindly sponsored by SGI and Cray Inc, was organised by CSAR staff (Ben Jesson, Mike Pettipher, Terry Hewitt and Kaukab Jaffri) working under the banner of the Manchester Research Centre for Computational Science (MRCCS), in which CSAR is a very active partner. Previous workshops in this series have been held in Lausanne, Edinburgh, Paris, Munich and Bologna.

The workshop attracted around 80 delegates from the UK, mainland Europe and the USA, including a number of users of the CSAR Cray T3E and SGI Origin 2000 systems. Thirty contributed talks and posters were presented at the workshop, covering a wide range of topics relevant to MPP applications programming:

- Applications of MPP systems in fields such as Chemistry, Physics;
- Astronomy, Meteorology and Management Science;
- Parallel programming languages, techniques and tools;
- Performance optimisation;
- Data management;
- MPP hardware and systems;
- Parallel numerical algorithms

These talks were refereed by the Programme Committee, a body of high-profile and well respected scientists from the fields of high-performance computing and computational science. In addition to reviewing contributed abstracts, the Programme Committee were also extremely helpful to the organisers, for example in attracting speakers and promoting the workshop.

Those papers accepted by the Programme Committee have been collected into the `Proceedings of the Sixth SGI/Cray MPP Workshop', and published on the MRCCS WWW pages, along with other details of the conference, at:

http://www.man.ac.uk/mrccs/mpp-workshop6/.

Papers or, in some cases, presentation slides can be downloaded from this site in postscript and PDF formats.

In addition to the contributed talks, two invited key-note presentations were given by representatives of the two sponsoring companies. These were "Developing Next-Generation MPP Systems from Application Needs", given by Steve Reinhardt, SGI's SN-IA Product Team Leader, and "The Challenges Ahead: Choosing the Right Architecture for High Performance Computing", by Steve Scott, Cray Inc.'s Chief Architect. These talks illustrated the directions being taken by either company towards new MPP systems, and how those directions relate to the needs of users such as those of the CSAR machines. Presentation slides from both talks are available from the conference WWW site referenced above.

The workshop sessions were held at the Weston Conference Centre in Manchester, in which also were provided the inter-session refreshments, lunches and the conference dinner. These occasions, also entirely sponsored by SGI and Cray Inc, and heavily attended by representatives of both companies, provided just as much a forum for in-depth discussions of MPP production systems and applications as the talks themselves!

From an organiser's perspective, this workshop went pretty smoothly (for the most part!), and seems to have been thoroughly enjoyed by all who attended. Also, from the perspective of one with a personal and professional interest in the applications of MPP computer systems, the great variety of talks from both MPP providers and users made for a highly stimulating event. Overall, the Sixth European SGI/Cray MPP Workshop has been a great success, and, along with many of this year's delegates, I'm sure, I look forward to attending the next in the series, wherever that may be!

Report: 9th ECMWF Workshop - Use of HPC in Meteorology- Developments in TeraComputing, 13-17th November

Dr Mike Pettipher,
Team Leader,
HPC User Services

This was the ninth in a very successful series of workshops held at ECMWF (European Centre for Medium-Range Weather Forecasts) in Shinfield Park, Reading. There were over 140 delegates from all over the world and a full program (see http://www.ecmwf.org/services/training/workshop00.html) of talks from Monday to Thursday, followed by a discussion forum on the Friday morning.

While the emphasis of the workshop was related to High Performance Computing in a weather/climate context, many of the talks were of much broader interest in discussing new architectures and suitability of programming models. It is a very good forum in which to see how the most powerful systems available are being used in real applications, which are of particular relevance and interest to most people around the world - not something that can be said of many other conferences.

I will make no attempt to list or describe all of the talks, but report on some of the themes, and topics of particular interest to me.

One of the most striking points is the vast range of computing facilities that are dedicated to this subject area, both for production work, of weather forecasts for example, and for research, particularly climate research. However most of these facilites are funded in some way by the US Government - a fact not lost by some UK researchers! It is even more impressive when researchers from such institutions say that these facilities are woefully inadequate to perform the calculations needed in understanding climate issues, and global warming in particular.

An example which demonstrates the need for more computing power is that of code or model coupling. This was a recurring topic during the week. The idea of coupling models has been around for some time, notably coupling ocean and atmospheric models, but the prospect of more power gives greater opportunity for coupling other environmental factors which may well have some influence on local or global events. Mike Ashworth from Daresbury Laboratory gave an interesting talk on some work he has been doing on POLCOMS - the Proudman Oceanic-Laboratory Modelling System, using the CSAR service. In this work, a hydrodynamic model was coupled with an ecosystem model (ERSEM - European Regional Seas Ecosystem Model) which includes biological and chemical processes both at the surface and on the sea bed. This produced simulations showing the location and movement of algal blooms in the North Sea resulting from nutrient flow from rivers and the circulation of the currents. In this case, variables from different models were coupled explicitly, but there were a few talks emphasising the benefit of using a 'code coupler', thus minimising the coding changes required, and simplifying the coupling of a variety of different models. Confirmation that this is being taken seriously in a commercial context was given by Karl Solchenbach of Pallas (a vendor developing tools for high performance computing), in his description of MpCCI - a Mesh based Code Coupling Interface, now available from Pallas. (Currently a free copy of MpCCI including full documentation and examples can be downloaded from the Pallas website - http://www.pallas.com.) MpCCI is designed to couple

both commercial codes such as ANSYS and STAR-CD, and also user written codes in Fortran or C, with communication based on MPI. If anybody is interested in using MpCCI, you are of course at liberty to download the material yourself from the Pallas website, but we would be interested to hear from you and of any experiences you have with it.

Another theme within the meeting was the assessment of OpenMP or mixed memory programs in comparison with MPI versions. Ilene Carpenter from SGI described experiences with numerical weather prediction models on the SGI Origin 3000. She showed that excellent performance can be obtained using OpenMP with these codes - comparable to the MPI versions. However, with one exception (in which the author advised that the results should not be taken too seriously yet) nobody showed better performance with OpenMP than with MPI. Even John Levesque from IBM in his enthusiastic talk about the new Power 4 architecture, who fully expected OpenMP to work better within a node, still found that the MPI version won. In many cases this may be because an enormous amount of effort has gone into the development of the MPI version, so it is perhaps not yet a fair comparison. Further, absolute performance may not be the only criteria that should be used to judge the value of a code - the effort to generate and maintain it may be more important in some circumstances. However it does suggest, as did Rupert Ford from the Centre for Novel Computing at Manchester, that at present, if you have a good MPI code, do not convert it to OpenMP just for the sake of it or to get better performance.

Wednesday was a combined event with the 17th bi-annual RAPS Workshop. RAPS (Real Applications for Parallel Systems) was set up to develop and maintain a collection of benchmark codes, to reduce the effort involved in benchmarking, both for customers and for vendors. There were a number of talks about the status of these benchmark codes.

Geerd Hoffman from the German Weather Service (DWD) in Offenbach talked about Grid computing. He gave some historical details about why this is important, the requirements and some examples, and then discussed one particular application which will be developed under the Eurogrid project (of which MRCCS is a partner). This application is the generation of detailed local weather forecasts on demand - something which will require substantial computing power, but can be distributed across systems relatively easily.

Bob Numrich of Cray described an application using Co-Array Fortran in the UKMET Unified Model. Co-Array Fortran offers an elegant way of using distributed memory systems like the Cray T3E, and Bob showed impressive performance figures indicating that the use of Co-Array Fortran in conjunction with MPI was better than pure MPI. At present, Co-Array Fortran is not widely used, but all reports of it seem very encouraging.

Finally, given the emphasis in this newsletter on the CSAR Technology Refresh, it is worth commenting that the German Weather Service (DWD) is following a very similar path to us - currently with a Cray T3E, but with a replacement system based on the IA 64 architecture from SGI.

Overall, this was an excellent workshop for HPC generally, not just for meteorologists.

EUROGRAPHICS 2001

Challenges in Computer Graphics for the 21st Century

Call for Participation

Location: Manchester, United Kingdom
Dates: 3-7 September 2001
Web site: http://www.eg.org/eg2001/
Email: eg2001-info@eg.org

Conference Theme

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" was chosen as the theme for Eurographics 2001.

Conference Program

The Conference Technical Programme spreads over three days, and is complemented with EG Workshops, industrial sessions, and two days of tutorials. Refer to the conference web site for full details.

The Conference Technical Programme will include three **invited talks** covering the three sub-themes. The speakers lined up are:

- Markus Gross (CH) – **"Surface Representations for Computer Graphics - A Personal Perspective";**
- Holly Rushmeier (US) – **"Rendering: Input and Output";**
- Chris Hecker (US) – **"Why Games Will Be the Preeminent Art Form of the 21st Century".**

Abstracts are available on the conference web site. **Research and application papers** are at the heart of the technical programme of the conference. **Short papers** will provide an alternative to the full paper programme, to include late results, fresh ideas, etc.. **State of the Art Reports** will cover the current state in selected key topics. The **Industrial Programme** will provide a forum on new and future trends with talks given by leading industrials. The planned themes will cover Aerospace, Computer Games and Process Industry.

Conference Topics

The Eurographics conference accepts papers from all areas related to the discipline of computer graphics, from hardware, through novel applications and algorithms, to human aspects of interaction with graphics, as well as artistic viewpoints. In support of the overall conference theme, three sub-themes have been selected as being of particular interest:

- Visualization
- Virtual Reality
- Computer Games

Submission Information

Authors are invited to submit original papers reporting research contributions, practice and experience, or novel applications. Independent referees will review papers and the International Programme Committee will make the final selection in March 2001. Conference proceedings will be published as the third issue of Computer Graphics Forum, Volume 20, both on paper and electronically, and will be available to all conference participants. Authors are requested to include a declaration that their paper has not been submitted to any other conference or journal and that they are prepared to present their paper at the conference. Papers must be written and presented in English.

Authors are also invited to submit proposals for short presentation (ca. 10 minutes is allowed per selected submission) of late results, work-in-progress, and new ideas. Submissions should be 5 to 10 pages in length. Papers accepted as short presentations will be printed in a separate volume, available to all conference participants.

State of the Art Reports (STARs) covering topics in research, practice, and applications, will be included in the conference programme. STARs are typically of 90 minutes duration and will be printed in a separate volume given to all conference participants.

Tutorials on computer graphics topics, especially those pertaining to the three sub-themes are encouraged. Each tutorial is typically organised as a one-day session, though proposals for half day tutorials will also be considered. A small honorarium will be paid to tutorial lecturers.

Full details of the electronic submission procedure are available at the conference web site:
http://www.eg.org/eg2001/

In case of difficulty, please contact the conference secretariat.

HSL 2000

The RAL Numerical Analysis Group

The Numerical Analysis Group at the Rutherford Appleton Laboratory is delighted to announce the release of HSL 2000. HSL 2000 is available to the UK academic community free of charge.

HSL (formerly the Harwell Subroutine Library) is a collection of Fortran packages for large scale scientific computation written and developed by the Numerical Analysis Group at the Rutherford Appleton Laboratory and by other experts and collaborators. The Library was started in 1963 and over the years it has evolved and has been extensively used on a wide range of computers, from CRAY supercomputers to modern PCs and high performance workstations. HSL now offers users a high standard of reliability and has an international reputation as a source of robust and efficient numerical software.

The latest release of the Library, HSL 2000, contains codes for Automatic Differentiation, Differential Equations, Eigenvalues and Eigenvectors, Mathematical Functions, Sorting, Linear Programming, Linear Algebra, Nonlinear Equations, Polynomials, Optimization and Non-linear Data Fitting. New packages include state-of-the-art routines for sparse linear programming, the solution of sparse linear systems, and optimization.

Free use of HSL 2000 by UK academics

While HSL 2000 is a commercial product, it is also available without charge to anyone working in an academic institution in the UK for teaching and academic research purposes. This innovation is a direct result of much of the core funding for the Numerical Analysis Group at the Rutherford Appleton Laboratory being provided by a grant from the Engineering and Physical Sciences Research Council. An academic licence for the use of any HSL 2000 package may be obtained without charge by UK academics through the HSL web page http://www.cse.clrc.ac.uk/Activity/HSL

Access is by means of short-lived individual password-controlled accounts. Potential users are asked for brief details of the use they intend to make of the package(s) they aim to download. Users are also asked to accept a set of terms and conditions (see http://hsl.rl.ac.uk/acuk/cou.html), and are not permitted to divulge their userid and password to anyone else, nor to distribute any HSL packages to a third party.
Information on the commercial use of HSL 2000 can be found by links from the main HSL web page.

About the Numerical Analysis Group at the Rutherford Appleton Laboratory

The Numerical Analysis Group is part of the Computational Science and Engineering Department (CSED) of the Central Laboratory of the Research Councils (CLRC). The Group is based at the Rutherford Appleton Laboratory and is one of the world's leading centres for the solution of sparse linear systems and large-scale nonlinear optimization problems. Group members hold visiting professorships at British universities, act on the editorial boards of many leading scientific journals, serve on international standards committees, and have written a number of recent books. The software developed by the Group has been described in a large number of published journal articles. The Group's activities, including the development of HSL 2000, are primarily supported by a strategic grant (GR/M78502) from the EPSRC.

CLRC comprises two major laboratories, the Rutherford Appleton Laboratory in Oxfordshire and the Daresbury Laboratory in Cheshire. CCLRC employs more than 1700 staff, who support the work of over 10000 scientists and engineers, mainly from the university research community.

For more information on the Numerical Analysis Group, visit their website at www.cse.clrc.ac.uk/Group/CSENAG. For details of CLRC, see www.clrc.ac.uk/.

Forthcoming Events

March 2001

4 - 9

First Global Grid Forum & European Datagrid Conference, Amsterdam - http://www.gridforum.org/

May 2001

16 -18

The 2001 International Workshop on Parallel I/O on Clusters (PIC 2001) http://www-unix.mcs.anl.gov/~thakur/pic2001/. Organized at the IEEE International Symposium on Cluster Computing and the Grid (CCGrid'2001), Brisbane, Australia - http://www.ccgrid.org/

21 -23

International Parallel CFD 2001 Conference, Parallel Computational Fluid Dynamics, Egmond aan Zee, The Netherlands - http://www.parcfd.org/2001conf/index.html

21-25

CUG Summit 2001, "2001 a Space Odyssey", The Jet Propulsion Laboratory (JPL), Indian Wells, California, USA - http://www.fpes.com/cugs01/

June 2001

6 -10

3rd International Conference on "Large-Scale Scientific Computations", Sozopol, Bulgaria - http://copern.bas.bg:80/Conferences/SciCom01.html

25 - 27

HPCN 2001, University of Amsterdam, NL - http://www.science.uva.nl/events/HPCN2001/index.html

August 2001

29 - 31

Europar 2001, University of Manchester, UK, http://europar.man.ac.uk

September 2001

3 - 7

Eurographics 2001, Challenges in Computer Graphics for the 21st Century, Manchester, UK http://www.eg.org/eg2001/

Europar 2001

CALL FOR PAPERS

Location: **The University of Manchester (MRCCS) and Manchester Conference Centre**

Dates: **28-31st August 2001**

Website: **http://europar.man.ac.uk**

The international Euro-Par conferences are held annually and are dedicated to the promotion and advancement of all aspects of parallel computing. The major themes can be divided into the broad categories of hardware, software, algorithms and applications for parallel computing.

Euro-Par 2001 Key Dates

29th January, 2001	Final date for submission
1st May, 2001	Acceptance notified
29th June, 2001	Early registration deadline
1st August, 2001	Late registration deadline
28-31st August, 2001	Conference

Euro-Par 2001 Provisional Topics

Core Topics:

1. Support Tools and Environments
2. Performance Evaluation, Analysis, and Optimization
3. Scheduling and Load Balancing
4. Compilers for High Performance
5. Parallel and Distributed Databases, Data Mining and Knowledge Discovery
6. Complexity Theory and Algorithms
7. Applications on High-Performance Computers
8. Parallel Computer Architecture
9. Distributed Systems and Algorithms
10. Parallel Programming: Models, Methods and Programming Languages
11. Numerical Algorithms
12. Routing and Communication in Interconnection Networks
13. Instruction-Level Parallelism and Architecture
14. Object Oriented Architectures, Tools and Applications
15. Architectures and Algorithms for Multimedia Applications
16. Cluster Computing
17. Metacomputing, Grid Computing and Web Computing
18. Parallel I/O and Storage Technology

New/Developing Topics:

19. Problem Solving Environments
20. Parallel and Distributed Embedded Systems

Workshops and Tutorials

We are also seeking topics for workshops and tutorials, please contact the Workshops and Tutorials Chair, John Keane (jak@co.umist.ac.uk) with ideas.

Further Details

Please see the website, or if you have any specific questions about the conference, **email europar@man.ac.uk** or contact Kaukab Jaffri on 0161 275 6824.