ISSUE 1, SPRING 2000

UKHEC

High-end computing

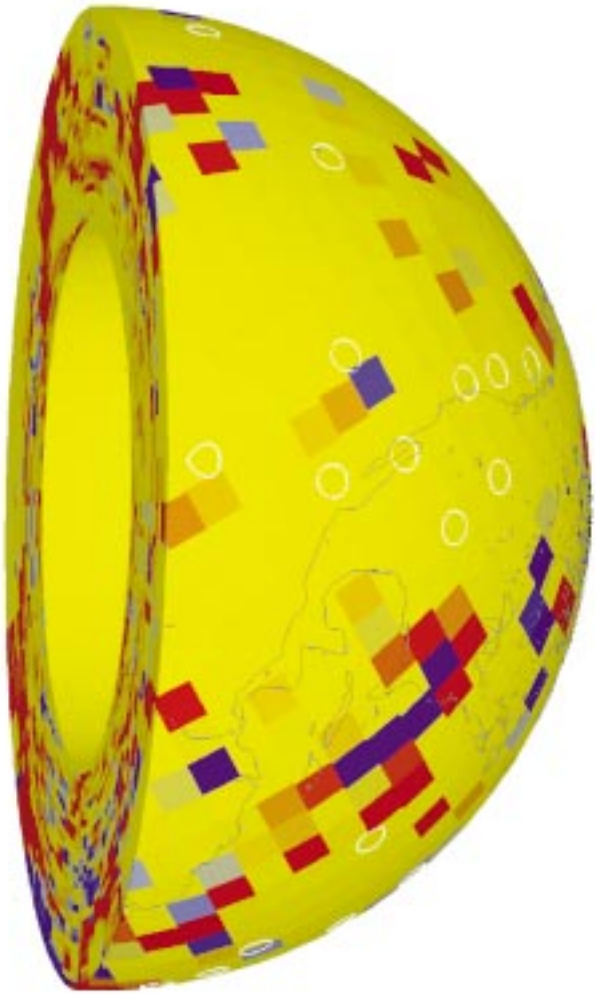

The Earth's mantle as visualised by the Terra project.
See page 10 for more information.

Contents

Editorial

Alan Simpson, EPCC

Welcome to the first issue of *UKHEC*, the new national newsletter for all researchers using high-end computing. This newsletter is produced by the new UKHEC collaboration of EPCC, Daresbury Laboratory and MRCCS, which has been recently funded by EPSRC.

The primary aim of UKHEC is to investigate emerging areas of computing and to inform and provide advice to the user community in hardware and software developments, in new tools, in best practice code development and in data management. We want to help you develop world-class application codes and so make effective use of current and future national facilities. We are keen to ensure good communication with HEC users and can be emailed at **ukhec@ukhec.ac.uk**. *UKHEC* will be a key part of the dissemination of our results and it will be published spring and autumn for at least the next three years.

This issue necessarily focuses on ongoing collaborative activities at the centres, which will be extended under UKHEC. We cover important new software developments in Java and OpenMP, as well as discussing the state of the art in visualisation and metacomputing. You will also find information about upcoming events in data management and software development, and about major national and international conferences on various aspects of HEC. This information will also be available on our new WWW pages at **http://www.ukhec.ac.uk**. I hope you find this first issue interesting.

If you have any comments or feedback, please email **ukhec@ukhec.ac.uk.**

CLRC CENTRAL LABORATORY OF THE RESEARCH COUNCILS
epcc
MRCCS

CLRC Daresbury Laboratory and the universities of Edinburgh and Manchester are collaborating to provide support for scientists using high-performance computers in the UK.

UKHEC: a joint initiative to support high-end computing

R.J. Allan, Daresbury Laboratory

The UKHEC Collaboration answers the need for a core activity to gather and spread information on international developments in computer software and programming tools, and to promote good programming practice. Funding is provided by the Engineering and Physical Sciences Research Council (EPSRC) which called for proposals from appropriate groups, and we are pleased to announce that our collaboration, which involves staff from CLRC's Daresbury Laboratory, the Edinburgh Parallel Computing Centre (EPCC) and the Manchester Research Centre for Computational Science (MRCCS), has been chosen to carry out this work. Our task is to help UK academic researchers in their use of advanced supercomputing facilities and to ensure the rapid exploitation of new computer technology.

UKHEC is particularly noteworthy in that this is the first time that there has been a strategic collaboration between the main centres offering nationwide academic computing support. It underlines the fact that computational modelling and the simulation of natural processes are now mainstream research activities, enabling lines of work which would otherwise be impractical and producing insightful results. A wide range of subjects is now benefiting from these techniques, from the quantum simulation of materials and chemical reactions, bio-sciences, fluid dynamics and aircraft design up to modelling the Earth's climate changes.

The UKHEC Collaboration will investigate critical issues in high-performance computing (HPC) hardware and software technology, and will aim to help researchers enhance their effective use of new and future computing resources in the UK. This includes programming languages (e.g. C++, Java, Fortran95), software engineering and software Quality Assurance (QA), data management, visualisation and Virtual Reality (VR), and metacomputing. Technical reports will be produced and short courses and workshops held in the key areas. Information will be announced through this newsletter and the UKHEC website (see below).

The HPCI centres

In the six years since the inception of the High Performance Computing Initiative (HPCI) in 1994, the Centres at Daresbury Laboratory and the Universities of Edinburgh and Southampton, working with the consortia of researchers they support, first ported and optimised their computer programs, and then contributed to the longer-term implementation of new scientific methodologies and numerical algorithms, building on experiences from previous massively-parallel implementations.

This continues the work of the Collaborative Computational Projects (CCPs), the earliest of which started in the late 1970s with support from Daresbury Laboratory. The work of the HPCI Centres and Consortia was effectively demonstrated at a conference held in Manchester, January 1998 [*1*].

We particularly welcome the renewed funding of the HPCI Centres at Daresbury and Edinburgh, which was announced in November 1999, to continue to provide direct support to program developers. The UKHEC Collaboration, which is funded separately, will dovetail with the Centres, by providing a central technology watch mechanism, assessing how new techniques may benefit computational scientists using the national HPC facilities.

The development of new HPC techniques is proceeding internationally at great speed. For example, in the USA there are several large collaborative activities funded by the federal government. These include the Accelerated Strategic Computing Initiative (ASCI) at the three largest US laboratories, which aims to simulate nuclear weapons and undertake stockpile stewardship in conjunction with the world nuclear test ban of 1994. This is stimulating the production of larger, faster computers with novel software, capabilities that are to be deployed in academia. It is essential for the scientific competitiveness of the UK that we make use of this cost-effective technology for our own future research needs.

[*1*] R.J. Allan, M.F. Guest, D.S. Henty, D. Nicole and A.D. Simpson (eds.) *High Performance Computing* Proc. HPCI '98 Conference 1998, (Plenum/Kluwer Publishing, 1999) ISBN 0-306-46034-3

UKHEC
http://www.ukhec.ac.uk

MRCCS
http://www.man.ac.uk/mrccs

CLRC HPCI Centre, Daresbury Laboratory
http://www.cse.clrc.ac.uk/Activity/HPCI

EPCC
http://www.epcc.ed.ac.uk

You can also contact the people listed on the back page.

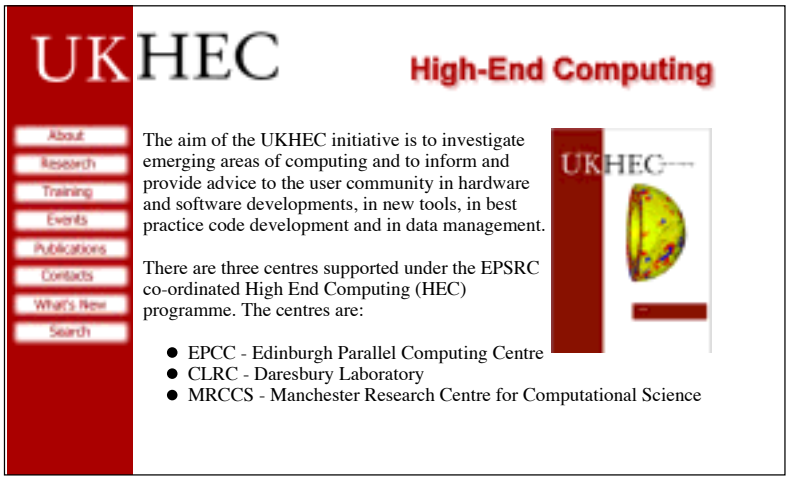

Announcing the UKHEC website: www.ukhec.ac.uk

David Henty, EPCC

Every project has to have a website, and as UKHEC is a collaborative project we didn't simply want to attach it to one of the centres' existing sites. We have therefore recently registered UKHEC as a separate UK academic domain, and not surprisingly the WWW pages are available via **www.ukhec.ac.uk**.

The details of the site and its structure will evolve over time, but we will be using it for the following:

- information on upcoming events including workshops, seminars and courses
- details of recent technical documents
- surveys of HPC hardware, software and literature
- reports from UK and international meetings
- a repository for all UKHEC presentations and documents, including these newsletters.

Please visit our WWW site regularly, as the content will be growing rapidly throughout our first year.

All comments on the web pages should be sent to: ukhec-www@ukhec.ac.uk.

Professional training for HEC researchers

Alan Simpson, EPCC

One of the key aims of UKHEC is to provide ongoing professional training for researchers involved in high performance computing. As well as running various seminars and meetings which will be advertised in this newsletter, we will be providing a series of courses over the next few years. The UKHEC courses will be complementary to those on HEC techniques already provided at the centres and will provide ongoing professional training for researchers.

The first pair of courses will be on:

- Advanced Programming Tools and Techniques
- Practical Software Development for Computational Scientists and Engineers.

The first course will include revision control, effective use of makefiles, compilers and modern tools for debugging, profiling and browsing. These tools and techniques are important in enhancing the efficiency of the software engineering process, permitting rapid code debugging, profiling and subsequent software re-engineering.

The second course will be a two-day course covering design, object-oriented techniques and software project development. The focus will be on the application of practical techniques allowing researchers effectively to develop and maintain high quality, portable code. Together these courses will provide researchers with valuable software skills which are vital for research in computational science and engineering.

The first runs of these courses should be in June and there will be ongoing runs every few months.

The UKHEC collaboration is committed to providing further courses over the next three years. Likely topics include *Data Management*, *Visualisation and Virtual Reality for HPC* and *Metacomputing*.

The dates of the courses will be advertised on:
http://www.ukhec.ac.uk/training
http://www.epcc.ed.ac.uk/epcc-tec/courses/timetable.html

New techniques are needed to manage the very large datasets which today's HPC applications produce.

Workshop on Advanced Data Storage and Management Techniques for HPC

Kerstin Kleese and R.J. Allan (Daresbury Laboratory), Michael Lautenschlager (DKRZ, Germany) and Michael King (University of Essex)

This international conference was held at Daresbury Laboratory on 23-25 February 2000, just as we go to press. We will provide a fuller report on the event via the UKHEC and conference web pages (see below).
The conference was organised jointly by the Computational Science and Engineering Department (CSE) at Daresbury Laboratory and the DIRECT Concerted Action of the European Commission.

DIRECT, which is facilitated by EPCC, unites fourteen large-scale computing facilities throughout the EU. It aims to address the role of emerging computer technologies in defining the future direction of scientific computing, and to represent to the Commission the interests and needs of all classes of users of HPC. Daresbury Laboratory's partners in organising the conference were two of DIRECT's Enterprise Groups: Data Storage and Data Management, and Data Inter-Operability.

The aim of the conference was to provide a forum for vendors, data and computing centres, developers and users to discuss available solutions, new technological developments and user requirements.

Topics on the programme included:

- High Performance Storage Systems
- Data Management Tools
- High Performance I/O on Supercomputers
- Data Mining
- Data Processing and Visualisation
- Data formats and Data Interoperability

Over 100 people attended the conference, which was generously sponsored by EPSRC, DIRECT, Sun and FileTek.

There is a growing awareness that it is not enough to provide leading-edge computational facilities on their own. Many of today's modelling codes produce large amounts of data, which have to be handled appropriately. Keeping records, archiving, preserving and exploring results, granting access to and disseminating new findings, as well as providing back-up capacities are major tasks and without adequate data archival and exploration, vital research data may well be lost.

The rapidly increasing amount of stored data provides new challenges for the existing computing and data centres, as new techniques have to be developed to capture, store, manage and explore the data efficiently.

The conference provided a forum to discuss future data storage and management requirements of various scientific disciplines. We invited contributions from the areas of: environmental science, quantum chromodynamics, materials chemistry, computational chemistry, biology, engineering, astrophysics, particle physics, medical and social sciences. If readers have comments on data management issues, or are aware of other important areas please contact the authors.

A full set of proceedings will be available at a cost of around £20.

Conference:
http://www.dl.ac.uk/TCSC/datamanagement/conf2.html

DIRECT:
http://www.epcc.ed.ac.uk/DIRECT

HPC facilities in the UK *R.J. Allan, Daresbury Laboratory*

The current national facilities for academic research include:

- 616-processor Cray T3E-1200E system at University of Manchester
- 8-processor Fujitsu VPP300 at University of Manchester
- 40-processor and 16-processor Silicon Graphics Origin2000 machines at University of Manchester
- 256-processor Cray T3E-900 system at University of Edinburgh
- 24-processor Compaq system at CLRC Rutherford Appleton Laboratory
- 60-processor IBM system at CLRC Daresbury Laboratory.

Other facilities are available to individual universities through the Joint Research Equipment Initiative which is run by EPSRC and the Higher Education Funding Council for England (HEFCE). These include the 256-processor Hitachi SR2201 at University of Cambridge.

A number of commercial companies and research institutions also have very large computers, including the 876-processor Cray T3E-900 and 636-processor Cray T3E-1200E systems at the UK Meteorological Office, a 116-processor Fujitsu VPP700 at ECMWF, the European Centre for Mid-range Weather Forecasting in Reading, a 252-processor IBM at AWE, Aldermaston and a 302-processor IBM at British Airways.

OpenMP: an emerging standard

Lorna Smith and Mark Bull, EPCC

OpenMP is a relatively new industry standard technique for shared-memory parallel programming. It is a set of directives and library routines for both Fortran and C/C++, and for the first time enables shared-memory parallel programs to be made portable. An important theme in the work of UKHEC will be to study the performance and application of OpenMP and to assess its relevance to high end computing users, and EPCC has been working on several key areas for this initiative.

OpenMP applications

EPCC has been involved in developing a number of real scientific parallel applications using OpenMP. One such application, a code developed to calculate the Earth's magnetic field, highlighted a number of issues. For example, implementing this application in OpenMP was significantly easier than implementing the MPI version, but to make the application perform satisfactorily required more time and effort.

OpenMP benchmark

The overheads of synchronisation and loop scheduling strongly influence the performance of shared-memory parallel programs. The cost of these operations depends on how they are implemented in the OpenMP run-time library. EPCC has designed a set of microbenchmarks to measure these classes of overhead for language constructs in OpenMP. These measurements can serve a number of purposes:

- to compare the efficiency of different implementations of OpenMP
- to compare the performance of semantically-equivalent directives (e.g. CRITICAL vs ATOMIC vs lock routines)
- to allow developers to estimate these overheads in their code by counting the number of directives executed and multiplying by the overhead time for each directive.

Our basic technique is to execute a section of code first sequentially, and then when enclosed in a given directive. Timings from the two runs are then compared. We take great care to derive statistically stable and reproducible results.

The synchronisation overheads which we measure are those associated with the directives: PARALLEL (with and without a reduction clause), DO/FOR, PARALLEL DO/FOR, BARRIER, SINGLE, CRITICAL, ATOMIC, ORDERED and lock/unlock. We also measure the loop scheduling overhead of the DO/FOR directive when using the STATIC, DYNAMIC and GUIDED options and different chunk sizes.

Both Fortran and C versions of the benchmarks are available (see illustrations).

So far, the benchmarks have been run on a Sun HPC 3500, a Silicon Graphics Origin 2000 and a Compaq Alpha server. They have shown up important differences between these versions of OpenMP and highlighted various ways in which improvements might be made. We would like to encourage OpenMP users to download the benchmarks and run them on their local machine. Comments, suggestions and results are all very welcome.

Mixing OpenMP and MPI

Interest in mixed OpenMP/MPI codes has recently increased with the advent of clustered SMP systems. EPCC has developed a mixed OpenMP/MPI version of a large QMC application code for the highly successful Electronic Structure of Solids consortium.

The original parallel version of the code was written in MPI and was highly successful, resulting in numerous publications for the consortium. This work focussed on re-writing the code to allow for an arbitrary mix of OpenMP and MPI parallelism.

The majority of the execution time is spent within the loop over electron configurations. We therefore placed compiler directives around this loop allowing the work to be distributed between the threads. By carrying out the simulation at this level of course-grained parallelism the MPI and OpenMP versions are

Continued on next page

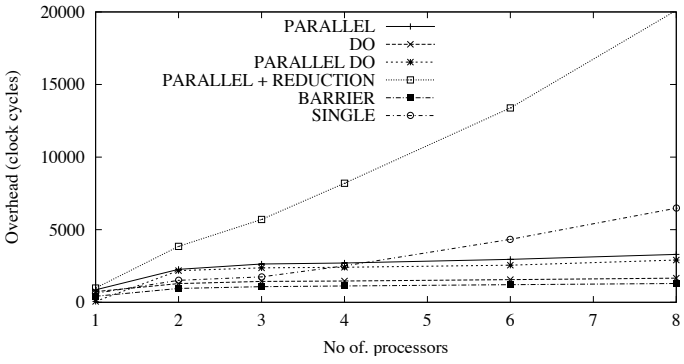

Overheads of OpenMP synchronisation constructs on an 195MHz SGI Origin 2000, using the MipsPro f90 compiler.

Optimisation of application codes is still the key to successful exploitation of the T3E.

Specialised optimisation for users of UKHEC

Currently the UK high-end computational facilities are heavily dominated by the very successful Cray T3E architecture, and major systems are available at CSAR and EPCC. The HPCI initiative was very influential in enabling the codes of UK researchers to exploit the parallel capabilities of the MPP architecture. However there is still much work to be done, and in particular the experience of the first year of the CSAR service has shown that the optimisation of codes to exploit the capabilities of the T3E processors can be very important.

One of the problems of MPP machines based on standard RISC processors, as compared to vector processors, is that typical performance on a single processor can be a low percentage of the theoretical peak. Typical scientific codes often run at less than 10% efficiency. A recent MRCCS initiative has been addressing this problem by offering to evaluate the codes of CSAR users and to consider whether optimisation of this kind might be useful.

The results have confirmed earlier experience that processor-level optimisation can result in dramatic increases in performance. UKHEC has access to a wealth of experience in such techniques at all of the participating sites and we intend to provide training and outreach work to enable UK science to benefit from this.

OpenMP *continued*

distributing the work in a similar manner allowing direct comparison between the two methods.

To ensure that the code is portable to systems without thread-safe MPI implementations, MPI calls are only made from within serial regions of the code. Hence the OpenMP loop parallelisation occurs beneath the MPI parallelisation. Initially, electron configurations are distributed between the MPI processes. The work is then further distributed by the OpenMP directives, so that each of the loops is executed in parallel between the OpenMP threads.

The code has been run on an SGI Origin 2000 and a Sun HPC 3500. For both systems the code scales well with increasing OpenMP threads, although not as well as with MPI processes.

Training

EPCC has developed a two-day training course which looks at the techniques involved in parallelising programs using OpenMP. Topics covered include parallel regions, worksharing directives, synchronisation, advanced topics and performance tuning. In addition, there is also a set of demonstrator codes in C and Fortran, and a Technology Watch report on OpenMP which provides a background and introduction to OpenMP and its implementation.

One-day courses are also available from Manchester and Daresbury, and UKHEC will work to combine and extend this material.

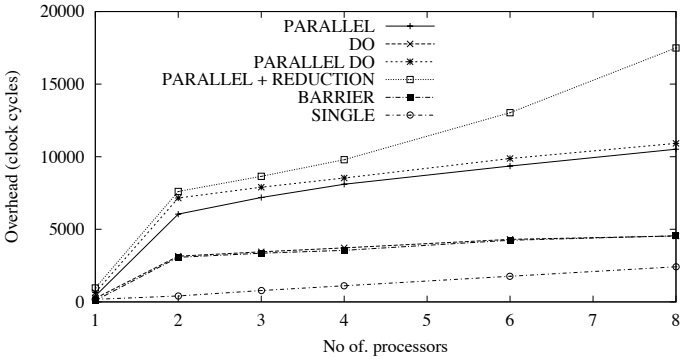

Overheads of OpenMP synchronisation constructs on an eight processor 400MHz Sun HPC 3500, using the KAI guidef90 compiler.

Calculating the Earth's magnetic field:
http://www.epcc.ed.ac.uk/ssp/1999/Reports/ss9903.pdf

EPCC OpenMP Microbenchmarks:
http://www.epcc.ed.ac.uk/research/openmpbench/

OpenMP training:
http://www.epcc.ed.ac.uk/epcc-tec/courses/OpenMP.html

http://www.epcc.ed.ac.uk/hpci/#openmp

http://www.epcc.ed.ac.uk/epcc-tec/documents/techwatch-openmp/

http://www.mcc.ac.uk/hpc/OpenMP

http://www.dl.ac.uk/TCSC/HPCI/reports.html

Java has the potential to offer a better programming environment for high end computing applications than other languages such as Fortran and C. UKHEC will see it as an important task to evaluate and monitor the development of these potential advantages.

Java and the future

Lorna Smith and Mark Bull, EPCC

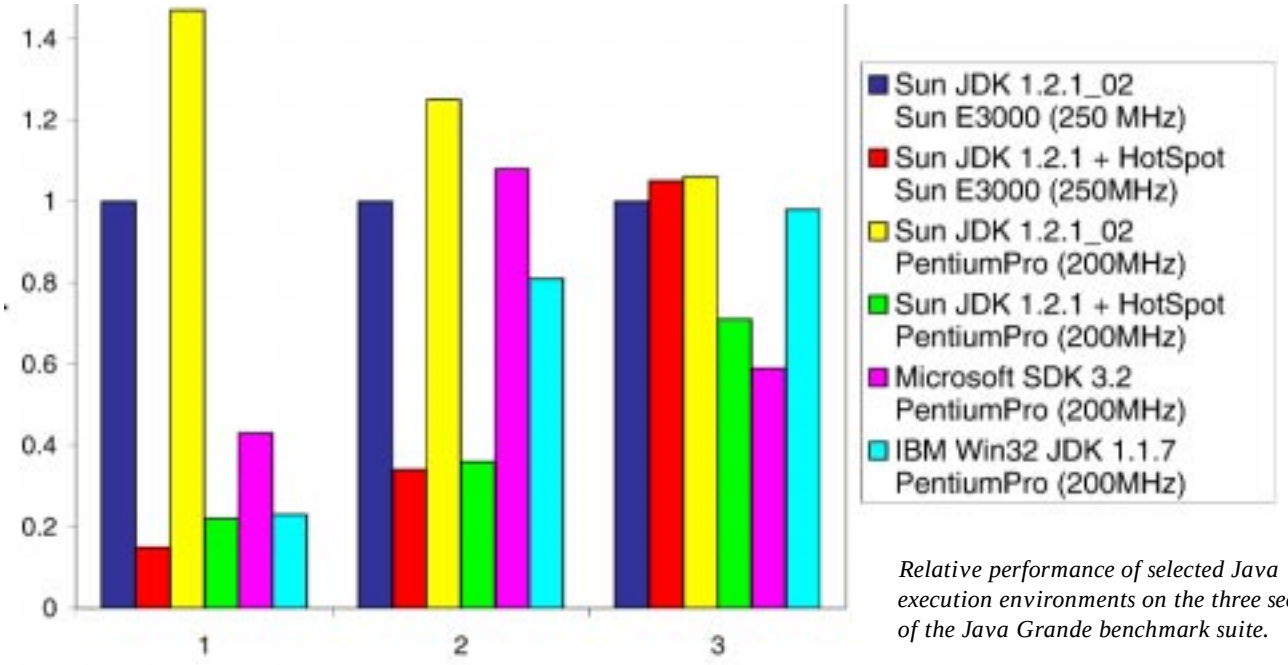

Relative performance of selected Java execution environments on the three sections of the Java Grande benchmark suite.

UKHEC's Java initiative will concentrate on a number of key areas.

The Java Grande Forum

Java Grande is a community initiative led by NPAC and Sun to promote the use of Java for "Grande" applications. A Grande application is one which has large requirements for any or all of:

- memory
- bandwidth
- processing power

Grande applications include high performance computing applications, as well as large-scale database applications and business and financial models. The Forum's aim is to propose modifications to the Java language to make it more suitable for the development of Grande codes, and to promote this use in the community. EPCC has already been actively participating in the Forum's activities, and is leading the initiative to develop a suite of benchmarks to measure the performance of Java execution environments.

The Benchmark Suite

The benchmarks have been gathered into a coherent suite which provides consistent output formats and a consistent definition of terms. A standard class library has been provided to execute and instrument the benchmarks, so that users can implement new benchmarks with the minimum of modification. The benchmarks are all open source, reflecting the desire of the Java Grande Applications and Concurrency Working Group (JGACWG) for the user to be able to understand exactly what the benchmark is testing.

The benchmark suite has been divided into three sections:

Section 1: Low level operations – measuring the performance of operations such as arithmetic and maths library operations, method calls and casting.

Section 2: Kernels – short codes that carry out specific operations which are frequently used in Grande applications, e.g. encryption, sorting and LU factorisation.

Section 3: Large scale applications – real Grande codes,

Continued on page opposite.

possibly less useful for comparative performance studies but worthwhile to demonstrate the potential of Java for tackling real problems. Examples include ray tracing and molecular dynamics simulation.

The benchmark suite has been tested on a range of hardware platforms and execution environments. This has resulted in a range of interesting results. For example, there is often no clear correlation between the performance of the various systems on the different sections of the benchmark suite. This suggests that the relative performance of the different environments depends strongly on the detailed nature of the computation.

Language comparisons

We are also interested in comparing the performance of Java and other programming environments. For this reason a number of the benchmarks are also available in C and Fortran. For example, the LU factorisation benchmark, which is based on the Linpack benchmark, is available in all three languages. Measurements carried out on a 1000 x 1000 problem size, using Java, C and Fortran on a Sun Ultra Enterprise 3000, demonstrate that execution times for Java are now less than two to three times those of the traditional native-code languages.

Parallel benchmarks

Developing parallel Java Grande benchmarks will allow us to assess the relative advantages and disadvantages of using Java for high end computing applications over other, more traditional parallel programming languages. Hence current work is focussed on developing a parallel benchmark suite. Parallel versions of a number of the benchmarks have already been developed using Java threads. For example, results obtained on a Sun HPC 3500 for the Series Benchmark, which computes the first N fourier coefficients of the function $f(x) = (x+1)^x$, demonstrated almost ideal speed-up on up to 8 processors.

An OpenMP-like interface for Java

The relatively new industry standard for shared-memory parallel programming, OpenMP, is enjoying increasing levels of support from both users and vendors in the high performance computing field. The standard defines a set of directives and library routines for both Fortran and C/C++ and has the advantage of providing a higher level of abstraction to the programmer than, for example, programming with POSIX threads.

It is, of course, possible to write shared-memory parallel programs using Java's native threads (see above). However, a directive system has a number of advantages over this approach. For example, the resulting code is much closer to a sequential version of the same program. Indeed, with a little care, it is possible to write an OpenMP program which compiles and runs correctly when the directives are ignored. This makes subsequent development and maintenance of the code significantly easier. Hence, EPCC has developed a prototype implementation of an OpenMP-like set of directives and library routines for shared-memory parallel programming in Java.

The Java Grande Forum:
http://www.javagrande.org/

Java Grande benchmarks:
http://www.epcc.ed.ac.uk/javagrande/

The first QBIT Software Testing Week was held on 4–8 October 1999 in London.

QBIT Software Testing Week

R.F. Fowler, C. Greenough (Rutherford Appleton Laboratory) and R.J. Allan (Daresbury Laboratory)

QBIT, a company specialising in software testing, organised the event for the first time this year. In previous years it was known as the "QCC/SQE Software Testing Week", and has been recognised as a high quality training event.

QBIT was formed after QCC Training Ltd was sold in 1997. The firm specialises in training in software testing and related topics including: test management; acceptance testing; testing techniques for testing specialists and developers; and data modelling.

The courses at the Software Testing Week ranged in cost from £375 to £1000 per day + VAT for the three-day ISEB Foundation Certificate Course and included topics: Using Automated Test Tools; Object Oriented Testing; Client Server and Internet Testing; Testing in the Real World; User Acceptance Testing; Managing Software Quality and Business Risk; Test Management; Implementing Technical Reviews; Usability Testing; Test Process Improvement; and Presenting the Case for Automated Testing.

The event also included The Tools Fair, which was a one-day free exhibition and lecture event held at the same location. This was attended by Ron Fowler and Chris Greenough representing the CLRC HPCI Centre.

Many of the major vendors of software testing packages were represented at the exhibition. Some of the products from these companies, and others of relevance to scientific application development on UNIX platforms, are described in a Technical Report *Software Quality Assurance for Fortran 90 – a Survey of Available Tools* RAL-TR-99-046 by Fowler and Greenough. This is available on-line at the URL below or from the authors. A fuller version of this report will appear on the UKHEC website.

Technical report:
http://www.dl.ac.uk/TCSC/HPCI/reports.html

Visualisation techniques can enable users of HEC to analyse large datasets interactively and intuitively.

Case studies in visualisation

Joanna Leng, Manchester Visualization Centre/MRCCS

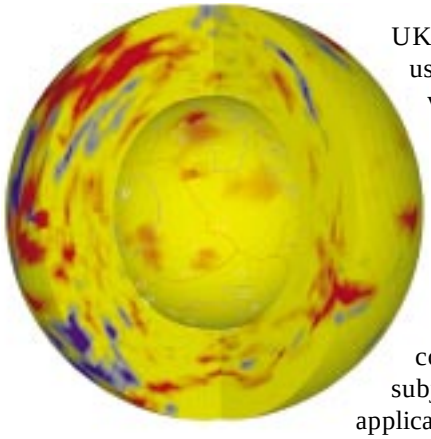

UKHEC will be investigating the use and promotion of visualisation techniques to aid HEC-based research. The intention is to examine three case studies in depth and to provide technical reports and consultancy based on this experience. The studies are designed to cover a range of different subject areas and types of application. In addition a post-doctoral research assistant workshop on visualisation will be provided based at MRCCS and using the facilities of Manchester Visualisation Centre. The MRCCS summer school last autumn had a strong component on visualisation and the analysis of the questionnaires provided to participants gives a starting point for identifying to concerns of UK HEC users. Users are welcome to contact the UKHEC service (details are given at the end of this article) for advice and assistance with visualisation.

Visualisation can be useful to users of HEC because it allows interactive visual analysis of large amounts of data. This visual analysis can be used for:

- Debugging
- Confirming a computational simulation is within reasonable bounds
- Giving 3D views of data with an inherent 3D physical geometry
- Visually comparing simulated data with images of real data
- Dissemination and education.

Some of the main topics new users of visualisation tend to need help with are reading the data into a visualisation system, constructing the right visualisation network to show the important features of their data, getting images or data out of the visualisation system and using the visualisation system efficiently.

The first case study which we shall examine is already in progress. The Terra project, lead from its UK base in the University of Liverpool by Prof. Huw Davies, involves simulations of the Earth's mantle. Its simulations run on 512 processors of the T3E and not surprisingly produce large data sets. The data is distributed on a rare type of spherical grid which has not been visualised in 3D before. Normally the researchers look at their data as a series of 2D slices but this simplifies the data by loosing its depth and distorts it by projecting the spherical shell into 2D. In addition, academic reviewers and journals are now demanding 3D reconstructions. With the size of this data and the amount of relevant reference data (the coastal boundaries, tectonic plate boundaries and geothermal hot spots) the visualisation techniques need to be carefully selected so that they reveal all the complexity of his data but remain efficient by extracting only the relevant components.

In other cases, we find that projects can be held up by problems in the analysis stage, and that visualisation techniques can help them to better understand their results and recommence their progress.

Important bottlenecks of this sort are caused not by a lack of HEC resources themselves, but in the interpretation of the results they produce. As large scale computational facilities grow in size and complexity, the use of visualisation and virtual reality becomes more important. It is noteworthy, for example, that this is a substantial part of the ASCI programme in the US.

Images: Joanna Leng and Huw Davies, Terra project.

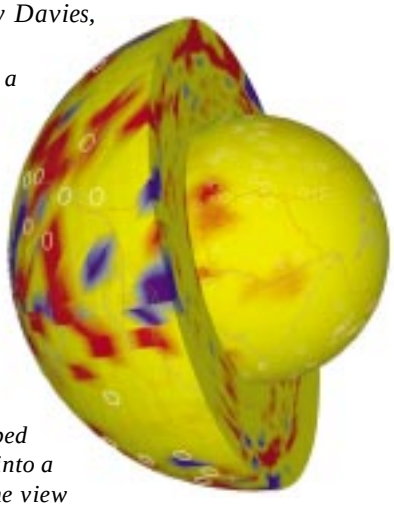

Top. Data has been mapped onto a globe which has been cut into a hemisphere along the equator. The view is perpendicular to the cutting plane. The innermost shell of the data has not been cut and on its surface are displayed the outline of the countries in black, the outline of the tectonic plates in red and "hotspots" are white rings.

Right. Data has again been mapped onto a globe which has been cut into a hemisphere along the equator. The view this time is towards America. The innermost shell of the data has not been cut. The outline of the countries are in black, the outline of the tectonic plates in red and "hotspots" are white rings.

Consultation or support:
Joanna Leng, Manchester Visualization Centre:
Tel: 0161-275 6252
Email: j.leng@mcc.ac.uk
http://www.man.ac.uk/MVC

MRCCS
http://www.man.ac.uk/mrccs/

In the future, metacomputing and grid-based techniques will enable groups of heterogeneous computers spread across the world to compete with supercomputers.

Metacomputing: spreading the load

J.M. Brooke and F. Costen, MRCCS, and R.J. Allan, Daresbury Laboratory

Metacomputing and grid-based computing are key technologies for the future of high end computing, and one important strand of work in UKHEC's initial year will be to compile and distribute a technical report on current developments in this field. A preliminary version, *Survey of Computational Grid, Meta-computing and Network Information Tools* by Rob Allan is available through the web (see below).

CSAR has already been actively involved in a metacomputing demonstration which won a prize at SuperComputing99 last November. This was summarised in CSAR Focus, edition 4 (see CSAR Web pages).

Recently Fumie Costen attended the IEEE Workshop in Cluster Computing in Melbourne, Australia to give a paper on experiences with load balancing and scheduling software on workstation clusters. Here she summarises her impressions of the current thinking on metacomputing as discussed at the workshop.

High-end metacomputing

Metacomputing is a term used to describe the linking together of distributed computational resources so as to compete with supercomputers. This is referred to as a "computational grid", and with it we can solve very large problems requiring, for instance, more main memory than may be available on a single system.

Such grids are typically heterogeneous in nature, with a hierarchy of communication links of differing bandwidths and latencies, from fast message passing or shared memory within each machine, to slow wide area networks connecting them. The implementation of a message-passing standard (e.g. MPI) on a metacomputer must take account of this in a way which is transparent to the user. Issues of security, resource management and load balancing are also crucial.

Message passing and resource management

Metacomputing is potentially applicable to all applications. Here we concentrate on tools and resources of interest to scientific programmers.

- Message-passing libraries. MRCCS has been working with HLRS Stuttgart, who have developed the PACX MPI library. It was this software that was demonstrated running metacomputer applications at SC99 via intercontinental links between Europe, the USA and Japan. PVM can also cope with heterogeneous clusters. Case studies involving coupled models on the different systems at CSAR are being tested and the results will be available to the community.
- Resource management software. RMS Systems allocate resources to jobs running on a distributed system, allowing users to find computers whose load is relatively low, or which satisfy a particular requirement (memory, vector processing, visualisation, etc). Examples are LSF (Load Sharing Facility) from Platform Computing, CODINE (COmputing in DIstributed NEtworked environment) from GridWare, as well as public domain products such as PBS (Portable Batch System) and DQS (Distributed Queuing System).

Metacomputing environments

Here are some examples of projects which provide the full functionality of a metacomputing infrastructure. An extensive survey will be part of the metacomputing report.

- GLOBUS is developing a basic software infrastructure for computations that integrate geographically distributed computational and information resources. This is a joint project of Argonne National Laboratory and the University of Southern California's Information Science Institute.
- Seamless Thinking Aid, from JAERI (the Japanese Atomic Energy Research Institute), is a Web-enabled Java-based environment which includes a number of tools for assisting parallel programming. Stampi has MPI-2 extensions in order to implement a heterogeneous computing resource. It can allocate child processes to computers dynamically, and provides a varying number of message routers to optimise communications.
- UNICORE is an initiative funded by the German Federal Government to provide infrastructure linking various supercomputing centres as a uniform resource. Participating sites do not have to alter their procedures for job submission, scheduling or resource control; instead there is a high level web-based interface, enabling very practical models for grid-based metacomputing to be developed.
- Legion has been deployed at a number of sites. The goal has been a highly-usable efficient and scalable system, designed to support a large degree of parallelism in application codes and to manage the complexities of the physical system for the user.

For further information on grid-based metacomputing see the book by I. Foster and C. Kesselman (eds.) *The Grid: Blueprint for a New Computing Infrastructure* (Morgan Kaufmann Publishers, 1998) ISBN 1-55860-475-8.

GLOBUS: http://www.globus.org

Stampi: http://ssp.koma.jaeri.go.jp/en/stampi.html

UNICORE: http://www.unicore.de/

Legion: http://legion.virginia.edu/overview.html

Preliminary survey:
http://www.dl.ac.uk/TCSC/HPCI/reports.html

SC99 demonstration:
http://www.csar.cfs.ac.uk/staff/costen/sc99

Optimisation and software tools on current and future IBM HPC systems were discussed at the ACTC meeting in Karlsruhe last September.

Working with IBM: ACTC workshop

R.J. Allan, CLRC Daresbury Laboratory

This three-day meeting at the University of Karlsruhe in Germany was part of the work of IBM's Advanced Computer Technology Center (ACTC).

ACTC, located at Yorktown Heights, NY, is part of IBM Research and is headed by John Levesque. The team of some 15 people is planned to soon expand to 20 or 30 with associated "competency centres", including Europe. ACTC was established to enable IBM to work more closely with application developers and provide tools and methods to exploit IBM SP systems which are not bound to the normal product lines. Information and free software will be made available via the web (see below).

This was the first workshop held by ACTC in Europe; further ones are planned at roughly six-monthly intervals. Similar workshops are held in the USA. In addition to discussions on optimisation and tools supported by ACTC, the workshops allow early announcements of new products, in this case the Nighthawk I, and discussion of future hardware such as the Power 4 and Giga-processor. Early results from the Nighthawk installation at Karlsruhe were presented by Matthias Krauss.

A number of speakers at this workshop are also active in writing the so-called Red Books, which provide useful technical programming information for IBM systems.

Much of the workshop focussed on programming and optimising codes for the new Power3 systems with SMP nodes. It was pleasing to see that a combination of “standards”, including Fortran 90, MPI, OpenMP, Pthreads, TotalView and VAMPIR, were being used for this, together with portability tools such as a SHMEM library (based on LAPI) and a SciLib library. Some parts of MPI-2 are also implemented, including single-sided communications and MPI-IO; however, performance gains may be possible by using the native LAPI and gpfs.

Fortran 95 is included in the release of "xlf" due early in 2000.

A fuller report on the Workshop and copies of the presented material are available from the workshop website (see below).

UKHEC and ACTC are discussing possible future collaboration on general computing issues. We will report any new developments in this newsletter.

ACTC:
http://www.ibm.com/research/actc

Red Books
http://www.ibm.com/redbooks

Workshop:
http://www.uni-karlsruhe.de/~SP/sp-workshop

Karlsruhe Rechenzentrum:
http://www.uni-karlsruhe.de/RZ

The University of Karlsruhe and the German national supercomputer network

The University of Karlsruhe is the oldest technical university in Germany. The Karlsruhe Rechenzentrum (Computing Centre) offers a 256-processor IBM SP2-SC/120 with up-to-date software and campus-wide ATM network. The Forschungszentrum Karlsruhe (Research Centre) took delivery of a 64-processor 8-way SMP IBM SP3/220 MHz Nighthawk I system just two weeks before the workshop.

Karlsruhe RZ is part of the German national supercomputer network and operates alongside HLRS Stuttgart. HLRS houses a 512-processor Cray T3E-900 and two large NEC systems. Together with two industrial partners, debis Systemhaus GmbH (a Daimler-Benz company) and Porsche AG, and the University of Heidelberg and the State of Baden-Württemberg, these centres operate together as a company, hww GmbH. There is thus a very strong focus on computational engineering for the aerospace and automobile industries.

As well as a strong science focus there is a high quality user service with courses on systems, optimisation and tools and activities to develop a meta-computing environment, METODIS and UNICORE. The latter is funded by the German government to provide seamless access to all the regional German supercomputing centres. It uses software from Pallas GmbH, who presented their VAMPIR, DIMEMAS and TotalView tools at the workshop. Stuttgart demonstrated its metacomputing capabilities along with Manchester, Pittsburgh, San Diego and Tsukuba at SC99 (see next article).

UKHEC at SC99

Alan Simpson, EPCC

This year's SC conference was held in Portland, Oregon on 14-19 November. For the first time ever, all three UKHEC centres organised research exhibits and these included significant representation from key UKHEC staff. Overall, the SC99 exhibits had an increasing presence from non-US centres and it is clear that this remains the major HEC event worldwide. Most of the major European sites had organised exhibits, including CEPBA, INRIA, Vienna, HLRS and Juelich.

All the HEC vendors take the SC conferences very seriously and SC99 was the ideal place to find out about current developments in HEC hardware and software.

SC99 clearly showed a few key trends:

- clustered SMPs are the mainstream future of hardware and programming them effectively will be a key issue
- more and more people are looking at Beowulf clusters of PCs as a cost-effective alternative in the mid-range
- OpenMP is increasingly popular as it is easy to use although, in general, performance is not yet up to MPI
- the continuing exponential growth in the WWW makes it vital to understand metacomputing and the grid.

Overall, SC99 was invaluable in keeping up to date with key HEC issues and the UKHEC collaboration is committed to having an ongoing presence at the SC conferences.

Left: Alan Simpson and Terry Sloan staffing the EPCC stand. During SC99 we distributed more than 200 copies of our new interactive CD-Rom which gives case studies from our technology transfer projects, European activities and academic research.

Centre left: Mike Ashworth at the Daresbury stand. Daresbury provided information on its activities via a couple of very large, glossy posters.

Bottom left: the Technology Transfer Network stand. This showed the interest in HEC throughout Europe and demonstrated the impact that has been made for small- and medium-sized enterprises.

Below: Terry Hewitt and John Brooke chatting at the CSAR stand. Terry ran a successful one-day tutorial on visualisation before SC99. (See also page 11.)

14-15 September, University of Edinburgh

Second European Workshop on OpenMP: EWOMP 2000

Mark Bull, EPCC

See pages six and seven for more on OpenMP.

Following the success of the First European Workshop on OpenMP, held at Lund University in Sweden, in September 1999, a second workshop will be held at the University of Edinburgh, and will be hosted by EPCC.

This workshop will be a forum for discussion of the latest developments in OpenMP and its applications.

Topics of interest for the workshop include:

- OpenMP implementations
- proposals for, and evaluation of, language extensions
- applications development experiences
- comparisons with other approaches such as MPI and HPF
- benchmarking and performance studies
- compilers, debuggers and performance analysis tools for OpenMP.

Key dates:

- Submission of abstracts: 15 May 2000
- Notification of acceptance: 26 June 2000
- Final papers due: 31st July 2000

Details of submission deadlines and registration will be announced shortly.

For more information please see: http://www.epcc.ed.ac.uk/ewomp2000

Email ewomp2000@epcc.ed.ac.uk

4-15 September, University of Manchester

Summer School on Linux

Manchester Computing and MRCCS are pleased to announce a two-week summer school in high-performance computing using the Linux operating system and commodity hardware. The school aims to provide the skills needed to benefit from the next generation of MPP HPC solutions, to administer and tune a Linux cluster or to purchase or build a Linux-based HPC system oneself.

This event will bring together people from all areas of the Linux/HPC world, from vendors interested in creating integrated Linux Supercomputers on proprietary hardware, through to people with experience of trying to make such systems function on a massive scale. Issues surrounding performance problems on commodity hardware and the implications for scientific programmers will be examined, as will resource management on platforms not often used in this way. Participants will have many opportunities for "hands on" experience in all areas, from benchmarking – maybe even to soldering!

More details and contacts will be available from the UKHEC, MRCCS and CSAR web sites.

First announcement and call for papers.

The Sixth European CRAY-SGI MPP Workshop

This event will be held on 7-8 September 2000 at the University of Manchester. The workshop is aimed at computer scientists, computer engineers and end users, working on Cray and SGI MPP systems in Europe.

Previous workshops in this series (see below) have attracted contributions of a very high standard. As usual, there will be a strong emphasis on applications, and we expect the workshop to be of great interest to all users of high-end computing facilities in UK academia. It will overlap with the coming HEC workshop, to be held at the same time. Contributions from UK users of HPC, and their attendance at this workshop, will be very welcome.

The deadline for receipt of extended abstracts is 9 June 2000.

For more information, please see the web pages of UKHEC, MRCCS or CSAR.

Workshop series:

I.1995	EPFL, Lausanne, CH
II.1996	EPCC, Edinburgh, UK
III.1997	I.F.G. (CEA), Paris, F
IV.1998	RZG/IPP, Garching, D
V.1999	CINECA, Bologna, I
VI.2000	Manchester

Recent publications

R.J. Allan, *Survey of Computational Grid, Meta-computing and Network Information Tools,* Technical Report DL-TR-99-02 2nd edition (Daresbury Laboratory, January 2000).
http://www.dl.ac.uk/TCSC/HPCI/reports.html

M.F. Guest, R.J. Allan, K. Maguire and B.G. Searle eds, Proc. 10th Daresbury Machine Evaluation Workshop (Daresbury Laboratory, January 2000).
http://www.cse.clrc.ac.uk/Activity/DisCo

J. M. Bull, L. A. Smith, M. D. Westhead, D. S. Henty and R. A. Davey, *A Benchmark Suite for High Performance Java*, in Concurrency, Practice and Experience, to appear (2000).
http://www.epcc.ed.ac.uk/research/publications/journal/javabenchcpe.ps.gz

J. M. Bull, *Measuring Synchronisation and Scheduling Overheads in OpenMP*, in Proceedings of First European Workshop on OpenMP, Lund, Sweden, Sept. 1999, pp. 99-105.
http://www.epcc.ed.ac.uk/research/publications/conference/jmbewomp.ps.gz

L.A. Smith, *Development and Performance of a Mixed OpenMP/MPI Quantum Monte Carlo Code*, in Proceedings of First European Workshop on OpenMP, Lund, Sweden, Sept. 1999, pp. 6-9.
http://www.epcc.ed.ac.uk/research/publications/conference/lasewomp.ps.gz

L.A. Smith, *Comparison of Code Development Tools on Cluster*s, Technology Watch Report, (EPCC, 1999).
http://www.epcc.ed.ac.uk/epcc-tec/documents/techwatch-clustertools/tools-1.html

Seminars and courses

The three UKHEC centres organise a regular seminar series. Anyone is welcome to attend, but please let organisers know in advance.
- Daresbury: **http://www.cse.clrc.ac.uk/Activity/Daresbury_CSE_Seminars**
- EPCC: **http://www.epcc.ed.ac.uk/epcc-tec/seminars**
- MRCCS: **http://www.man.ac.uk/mrccs/seminars**

To book your place on the courses at EPCC or CSAR (Manchester) listed below please contact:
epcc-support@ed.ac.uk or **csar-advice@csar.cfs.ac.uk**

26–30 MARCH (Keble College, Oxford)
ICAMDATA: 2nd International Conference on Atomic and Molecular Data and their Applications
For further information contact:
icamdata@am.qub.ac.uk

3–4 APRIL (Keble College, Oxford)
International Conference on Exploiting Leading-Edge High Performance Computing. This is a joint event with AWE and EPSRC.
www.awe.co.uk/hpc

6 JUNE (EPCC)
Advanced Programming Tools and Techniques

7–8 JUNE (EPCC)
Practical Software Development for Computational Scientists and Engineers

9–21 JULY (Durham)
9th Annual EPSRC Summer School on Numerical Analysis
fourier.dur.ac.uk:8000/nass

4–15 SEPTEMBER (University of Manchester)
Second MRCCS Summer School on Supercomputing with Linux
Please check dates nearer the time via: **www.csar.cfs.ac.uk**

6–7 SEPTEMBER (University of Manchester, to be confirmed)
First UKHEC Annual Seminar

7–8 SEPTEMBER (University of Manchester, to be confirmed)
6th European SGI/Cray MPP Workshop

14–15 SEPTEMBER (EPCC)
EWOMP 2000
www.epcc.ed.ac.uk/ewomp2000

23–24 NOVEMBER (Daresbury Laboratory)
11th Daresbury Machine Evaluation Workshop
See **www.cse.clrc.ac.uk/Activity/DisCo** for information nearer the time.

Feedback form

We'd like to know what you think of this newsletter. You can pass on comments by faxing back this form to: 0131 650 6555. This form is also online at: http://www.ukhec.ac.uk/feedback.html

1. Did you find any articles particularly interesting or useful? If so, which?_____

_____

2. Why?_____

_____

3. What topics would you like to see covered in future issues? _____

_____

4. Please feel free to add any further comments _____

_____

_____

Your name_____ Job title _____

Company/institution _____

Address_____

_____

Email_____

Telephone _____ Fax _____

Was your copy of this newsletter properly addressed? Yes/No

Contact information

MARTYN GUEST
Computational Science and Engineering Department
Daresbury Laboratory
Daresbury
Warrington WA4 4AD
m.f.guest@dl.ac.uk
01925 603240
http://www.cse.clrc.ac.uk/Activity/HPCI

ALAN SIMPSON
EPCC
JCMB, Maxwell Buildings
University of Edinburgh
Mayfield Road
Edinburgh EH9 3JZ
a.simpson@epcc.ed.ac.uk
0131 650 5120
http://www.epcc.ed.ac.uk

TERRY HEWITT
Manchester Computing
University of Manchester
Oxford Road
Manchester M60 9PL
w.t.hewitt@mcc.ac.uk
0161 275 6095
http://www.man.ac.uk/mrccs
http://www.csar.cfs.ac.uk

ROBERT ALLAN
Computational Science and Engineering Department
Daresbury Laboratory
Daresbury
Warrington WA4 4AD
r.j.allan@dl.ac.uk
01925 603207
http:///www.cse.clrc.ac.uk/Activity/HPCI

DAVID HENTY
EPCC
JCMB, Maxwell Buildings
University of Edinburgh
Mayfield Road
Edinburgh EH9 3JZ
d.henty@epcc.ed.ac.uk
0131 650 5960
http://www.epcc.ed.ac.uk

ALISON WALL
Research Facilities Team
EPSRC
Polaris House,
North Star Avenue,
Swindon SN2 1ET
a.wall@epsrc.ac.uk
01793 444575
http://www.epsrc.ac.uk/hpc