

ISSUE 5, SPRING 2003

UK HEC

The newsletter of the UK high-end computing initiative



*HPCx:
a new resource for
UK scientific research*

Contents

- | | | | |
|---|--|----|--|
| 3 | Programming shared memory systems | | UKHEC Annual Seminar |
| 4 | HPCX: a new resource for UK computational science | 9 | UKHEC Seminar:
GUIs for scientific computing |
| 6 | Portable Parallel Programming Profiling Tools: Vampir | 10 | Deploying the UK e-Science Grid
Web services, software engineering and the Grid |
| 7 | Potential of commodity graphics hardware for scientific computation | 11 | UKHEC reports |
| | EPCC and OpenMP | 12 | Feedback form
Contact information |
| 8 | Motivation and requirements for a user interface to a computational grid in the UK | | |

Editorial

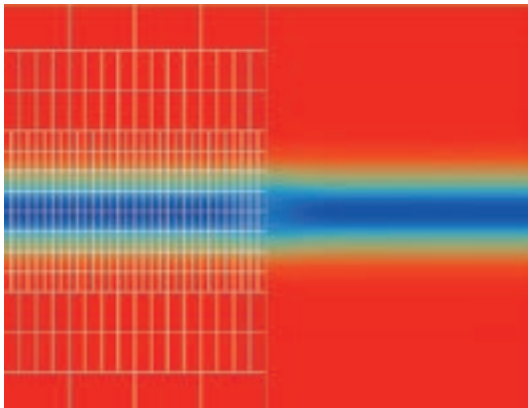

Mike Ashworth, Daresbury

This is the 5th and final UKHEC newsletter. The UKHEC programme, which began in January 2000, officially ended on 31st December 2002. The goal of the programme was to track and disseminate information on international activities in computer architectures, software and programming tools and to promote good programming practice for the HEC community via workshops, seminars, reports and mentoring. In the three years we have produced 33 Technical Reports and 13 Technology Watch Reports, provided 17 Training Courses (a total of 36 Course Days), carried out 15 Research Topics and Visualization Case Studies, held 3 two-day Annual Seminars and 10 other Seminars and Workshops and published 5 editions of the Newsletter, along with a number of other activities.

As one programme completes another is born. UK High End Computing has received a significant boost in resources with the award of the HPCx contract and the start of service on the UK's first single-system terascale facility. This augments and complements the existing facilities at CSAR, giving UK researchers greater computing power and greater choice.

This newsletter gives an introduction to the HPCx service as well as highlighting some of the UKHEC activities from the final year.

EPSRC will soon be calling for proposals for follow-on activities in high end computing, so we hope to be back with you in the future in some shape or form. UKHEC is dead! Long live UK HEC!



UKHEC SEMINAR

Programming Shared Memory Systems

Gavin J Pringle, EPCC

A UKHEC seminar on Programming Shared Memory Systems was held at EPCC on the 8th of November 2002. The event was also broadcast over the Access Grid allowing for anyone in the UK to attend the seminar simply by going to their nearest Access Grid node rather than travelling all the way to Edinburgh. Indeed, one of the speakers, namely Kevin Roy from CSAR, gave his talk from his home institute over the Access Grid.

But what are shared memory systems? Well, a significant number of parallel computers utilise shared memory architectures. This can take the form of a single box containing many processing elements, all accessing the

same memory, or a cluster of such boxes, linked with some interconnect, allowing the boxes to be used as a single resource.

The half-day seminar consisted of three talks concerned with the programming of these architectures. The three talks were as follows.

- Mark Bull, EPCC, University of Edinburgh, 'OpenMP: Current Implementations and Future Directions'.
- Kevin Roy, CSAR, University of Manchester, 'Is programming OpenMP really as easy as everyone says?'.
- Jurgen Dreher, Applied Mathematics, University of St Andrews, 'Parallel Mesh-Adaptive Fluid Simulations Using a POSIX-Threads/MPI Hybrid'.


Above. Break-up of an electric current sheet in a magnetized plasma by the process of 'magnetic reconnection'. A slight perturbation (a) grows into a violent instability in which the plasma is heated and accelerated at shock-like structures (b). A hierarchy of recursively refined numerical meshes covers the critical regions efficiently with the necessary local resolution. Magnetic reconnection plays an important role in many phenomena like Solar eruptions and flares, geomagnetic activity and the (unwanted) disruption of fusion plasmas.

Below. A simulated Rayleigh-Taylor instability in which an unstable density stratification of a fluid develops into a turbulent flow under the influence of gravitation. The mesh coverage of the critical regions in the flow is displayed in the right half of the panel, where the colours indicate the mesh distribution among eight computing nodes. The developed code implements a hybrid of multi-threading and MPI distribution with explicit dynamic load balancing.

The seminar was attended by around 40 people: about 30 locally at EPCC, including attendees from Aberdeen and Liverpool, plus nine and four in the two remote Access Grid nodes, namely the E-Science North West Regional Centre at Manchester and the Oxford E-Science Centre, respectively.

For more information on the Access Grid, visit: www.accessgrid.org.

PDF versions of the slides from all three talks, including two movies shown as part of Dr Dreher's talk, are available from the UKHEC website, at: <http://www.ukhec.ac.uk/events/smp/>





HPCx: a new resource for UK computational science

Mike Ashworth, Daresbury

HPCx is the name of the UK's new National High Performance Computing Service. It is a 1280-processor IBM p690 cluster whose configuration is specifically designed for high-availability capability computing. The Engineering and Physical Sciences Research Council (EPSRC) is overseeing the project, on behalf of the UK scientific community. HPCx is a joint venture between the Daresbury Laboratory of the Central Laboratory for the Research Councils (CLRC) and Edinburgh Parallel Computing Centre (EPCC) at the University of Edinburgh. IBM (UK) Ltd has been chosen as the hardware supplier for the six-year duration of the project.

1. The HPCx project

The British Government, through EPSRC, has funded the project to approximately £53M (~\$85M). The scope of the project is to operate and support the principal academic and research high performance computing service for the UK. The principal objective is to provide a capability computing service to run key scientific applications that can only be run on the very highest performing computing platforms.

The project is a collaboration between three partners: CLRC Daresbury Laboratory, Edinburgh Parallel Computing Centre (EPCC) and IBM. HPCx, formally known as UoE HPCx Limited, is a wholly-owned subsidiary of the University of Edinburgh. It is under contract to EPSRC to provide HPC services and it sub-contracts appropriate packages to the consortium partners. The partnership combines two of Europe's leading academic HPC, e-Science and technology transfer centres, with significant experience in operating and supporting national grand challenge and capability services on novel architectures.

CLRC Daresbury Laboratory has been an HPC service provider to the UK academic community for over 25 years – the first Cray 1 in Europe was delivered there in

1979. CLRC's Computational Science and Engineering Department is the UK's premier research, development and support centre for leading-edge academic science and engineering simulation codes. CLRC also provides distributed computing support for COTS processor and network technologies, evaluating scalability and performance. CLRC hosts the UK Grid Support Centre.

EPCC, located at the University of Edinburgh, was established in 1991 as the University's interdisciplinary focus for high-performance computing and its commercial exploitation. EPCC has hosted specialised HPC services for the UK's QCD community since 1989. The 5Tflop/s QCDOC system is due to be installed in 2003, in a project with Columbia, IBM and Brookhaven National Laboratory. From 1994 until 2002, EPCC operated and supported UK national services on Cray T3D and T3E systems.

Within HPCx, CLRC and EPCC both provide 'Added Value' services (primarily scientific support, code optimisation, and code development) and operations and system support services. The physical accommodation is provided at CLRC's Daresbury Laboratory.

IBM is the technology partner, providing the hardware, system software and appropriate maintenance and support



services. IBM is the world's largest commercial supplier of capability computers and a world-ranking research organisation. They offer world-leading and competitively priced HPC technology, an aggressive road map over the next 6 years and significant resources for science support

HPCx is a six-year project and the technology will be provided in three phases, with defined performance levels at year 0, year 2 and year 4. The performance targets for the three phases are defined in terms of Linpack Rmax performance. The targets for the three phases are approximately 3 Tflop/s, 6 Tflop/s and 12 Tflop/s, approximately following Moore's Law. Online and offline storage will increase proportionately, with 50, 100 and 200 TB in the three phases. There is a flexible approach to the technology refreshes in phases 2 and 3, and future upgrades are likely to feature the 'Federation' switch and POWER5 architecture.

The focus is one maximising the delivery of capability computing; the service is not intended to be used as a task farm or for multiple modest-sized throughput jobs. There is a comprehensive support service providing porting, optimisation, training and new applications outreach. The service has a 7 x 24 operating regime with high RAS requirements.

2. The HPCx Phase 1 System

Delivery of the Phase 1 system commenced on 4th October 2002. The full user service opened officially on 9th December 2002, although many users had benefited from up to a month's early access prior to this date. The system comprises 40 Regatta-H SMP compute nodes connected by the 'Colony' SP Switch2.

Each shared memory node has 32 1.3GHz POWER4 processors and 32 GB memory giving the system an aggregate memory of 1.28 TB. The POWER4 processor has dual floating point units each of which, through the fused multiply-add instruction, is capable of delivering two results per clock cycle. This gives each processor a peak performance of 5.2 Gflop/s and the whole system a peak of some 6.6 Tflop/s.

In order to increase connectivity to the switch and improve communications throughput, each compute node is configured as four 8-way logical partitions (LPARs). The 'Colony' SP Switch2 allows each LPAR to have two

connections (PCI adapters) into the switch fabric (dual plane), providing approximately 20 usec latency and 350 MBytes/sec bandwidth.

Two additional 16 processor nodes are provided as I/O systems. The system runs AIX version 5, with GPFS for file system support (18TB EXP500) and HSM for backup and archive to tape storage (35TB LTO tape library).

3. HPCx and CSAR

HPCx joins the established CSAR service, based at the University of Manchester, to form the UK's two major national HPC services. CSAR provides a range of supercomputers which complement the HPCx system.

The systems at CSAR include a new 256 processor SGI Altix 3700 Itanium2 system, a 512 processor SGI Origin 3800 with 512GB of shared memory, a 816 processor Cray T3E MPP system, a 128 processor Origin 2000, and others. These are supported by a storage area network, providing a common filestore across all the SGI computers, with a flexible range of disk types from Ultra High Performance for maximum performance for I/O bound tasks, to High Volume for large volume of data to be online stored cheaply. A 125GB tape store is provided for archiving.

CSAR will continue to run alongside HPCx for the next 3 years, until June 2006.

The two services will continue to provide UK researchers with world-class facilities for the next few years, and the community can look forward to this trend continuing with the proposed HPCy service in 2005 - an order of magnitude faster again. An HPCz also appears on the UK HPC roadmap. The strength of the existing facilities and of the roadmap, demonstrates the UK Research Councils commitment to the role of HPC, and the importance of computational science as a means of conducting scientific research.

For more information, please see:
<http://www.hpcx.ac.uk/>

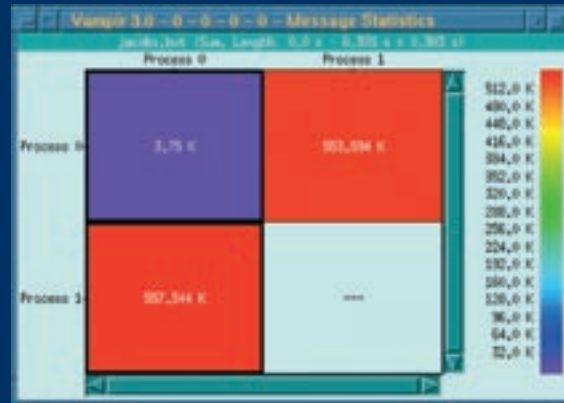
<http://www.csar.cfs.ac.uk/>

<http://www.epsrc.ac.uk/hpc>

Portable Parallel Programming Profiling Tools: Vampir

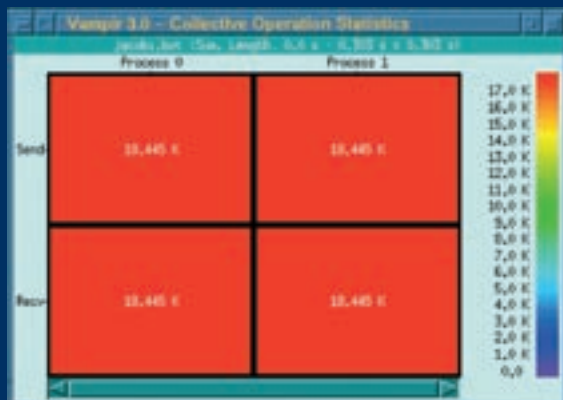
Chris Johnson

Vampir (Visualisation and Analysis of MPI Programs) is a portable tool for profiling MPI codes written in Fortran (77 or 90) and C. It is used extensively on MPI systems at HPC centres all around the world, and is available on most of the UK's HPC resources. For example, Vampir is available on HPCx, the UK's new 1280-node IBM system. The aim of this report is to provide an overview of



VAMPIR and to describing the more interesting features of the system. A simple user guide is provided and the ease-of-use of the vampir system described from the point of view of a new-user.

The report is available at http://www.ukhec.ac.uk/publications/tw/ppppp_tools.pdf



DIVE: Visualisation and Steering of Remote Applications

Jake Duthie, Kevin Stratford, JC Desplat, Lorna Smith, EPCC

Computational steering is an area of increasing interest to the UK High End Community, especially in terms of the computational grid. Providing demonstrators is therefore of interest. The aim of this project was to provide a simple visual movie of one particular computational steering project, the Distributed Interface to Visualisation Environment (DIVE) project.

DIVE is a library which allows computational steering and visualisation of distributed applications across a network. In the movie, we demonstrate steering and visualisation using a model simulating the mixing of two different fluids. The visualisation and steering component

runs on a workstation while the fluid model runs on a separate, parallel, machine. DIVE reassembles complete data fields from the parallel model and sends them to the visualisation environment. In this case AVS/Express is used. DIVE also allows model parameters to be changed from the visualisation environment as the model is running ie computational steering.

You can view the movie at:
<http://www.ukhec.ac.uk/publications/DIVE.avi>



Potential of commodity graphics hardware for scientific computation

Andrew Murdoch, EPCC

The aim of this report is to explain the potential value of commodity graphics hardware in the context of desktop scientific computation.

Much effort has been put into improving performance at the high end of scientific computation. Many tasks, especially those involving visualisation or real-time processing of data, can be performed using desktop machines, such as high end workstations which include multi-processor systems.

Graphics hardware has traditionally been used for visualisation, entertainment and interactive design. However, it is possible to use such equipment for

computational purposes. This situation can take advantage of the advances and rapid improvement in performance of graphics technology. In this report we discuss the potential of these devices and their possible applications.

The report is available at http://www.ukhec.ac.uk/publications/reports/gfx_report.pdf

Above. An example screenshot from the graphically intensive computer game 'Quake 3 Arena'. Games like this have driven demand for high-performance consumer 3D hardware.



EPCC and OpenMP *Mark Bull, EPCC*

EPCC has a strong interest in the development of OpenMP, the industry standard for shared memory parallel programming, as part of our longstanding commitment to standards in High Performance Computing.

The OpenMP Application Program Interface supports multi-platform shared-memory parallel programming in C/C++ and Fortran on all architectures, including Unix platforms and Windows NT platforms. OpenMP is a portable, scalable model that gives shared-memory parallel programmers a simple and flexible interface for developing parallel applications for platforms ranging from the desktop to the supercomputer.

The caretaker body for the OpenMP language is the OpenMP Architecture Review Board (ARB), which consists of vendor companies and academic institutions, including EPCC.

As well as the ARB committee itself, which meets to take care of administrative matters, the work of the ARB is carried out by a number of committees which meet regularly by email and phone conference:

- Language committee – following the publication of the Fortran 2.0 specification in November 2000 and the C/C++ 2.0 specification in March 2002, the language committee has recently started work on a combined version of the Fortran and C/C++ specifications, in order to remove inconsistencies, and to make further developments to the language easier.

- Futures committee – considering new features for inclusion in OpenMP 3.0. Topics under discussion include task queues, semaphores, control of idle threads and thread stack-size, automatic data scoping and new loop scheduling strategies.
- Performance/debug committee – working towards specifications for interfaces to OpenMP for performance analysis tools and debuggers. The aim here is to allow a single tool to operate with any OpenMP implementation.
- Runtime committee – looking at issues of job scheduling, including the interaction between OpenMP and MPI. Many batch systems do not adequately handle the mixture of threads and tasks present in mixed OpenMP/MPI jobs. This work aims to produce standard ways of requesting resources for such jobs.

EPCC is currently chairing both the Language and Futures committees, giving us a strong voice in the future development of the language.

EPCC has also been involved in the establishment of cOMPunity, the OpenMP users group. cOMPunity is a forum for the dissemination and exchange of information and research about OpenMP. As well as providing web-based resources, it organises three workshops each year (EWOMP in Europe, WOMPAT in the US and WOMPEI in Asia/Pacific) devoted to research and development in OpenMP.

Links

For more information on the OpenMP ARB: <http://www.openmp.org/>
For more information on cOMPunity: <http://www.compunity.org/>

Motivation and requirements for a user interface to a computational grid in the UK

Ali Anjomshoaa and Lorna Smith, EPCC

In an area such as High Performance Computing (HPC), where the lifetime of application codes often exceeds that of most HPC machines, the ability to access these machines through a simple generic end-user interface is of real interest. By developing an end-user interface to a wide area computational Grid that includes HPC machines, we may be able to hide from the users some of the complexity associated with the diverse and changing nature of the UK's HPC systems and their administrative domains, from the users.

This report examines the need for an end-user

interface to a wide area computational Grid, which will be comprised of a heterogeneous environment of machines in diverse administrative domains, within various institutions across the UK. The motivation is described, followed by a discussion of the requirements for an end-user interface. This discussion also addresses possible implementations of the end-user interface.

The report is available at:

http://www.ukhec.ac.uk/publications/reports/ui_require.pdf

Getting the most from your processor: The Third UKHEC Annual Seminar

Rob Allan, Daresbury

The Third UKHEC Annual Seminar was held from 9th to 10th December 2002 at Daresbury Laboratory, a few weeks after SC'02 and immediately before the annual Daresbury Machine Evaluation Workshop on 11th and 12th December. The focus of this year's seminar was a traditional one, 'Getting the most from your processor'. This theme ensured some interesting presentations on both hardware and software, plus a lively debate. Around 60 people attended the event.


Speakers, who addressed issues such as EPIC vs. RISC vs. vector architecture, hierarchical cache and memory optimisation, compiler design, performance monitoring tools, numerical libraries, multi-threading and languages, included:

- Mike Ashworth (CLRC), The UKHEC Collaboration
- Jan Boerhout (NEC ESS), High Performance through Vector-Parallel Processing.
- Bob Carruthers (Cray), Optimisation Strategy for exploiting the Cray X1 Architecture.
- Judit Gimenez (CEPBA, Barcelona), Obtaining useful Information from raw Performance Data.
- Martyn Guest (CLRC), Benchmark Performance of current Processors.
- John Hague (IBM), Optimisation for IBM's p690 Power4 system.
- Gernot Hoyer (Intel), Getting the most out of the Intel Itanium Architecture.
- Werner Krotz-Vogel (Pallas), Software development Tools for getting the most from your Processor.
- Rick Kuftrin (NCSA, USA), Experiences with first-generation Itanium at NCSA.
- Deborah Salmond (ECMWF), Implementation of a global weather forecasting system on an IBM highly parallel scalar system with 960 Power4 processors.
- Jennifer Scott (CLRC), Numerical Libraries, HSL and large sparse Systems.

Thanks to Shirley, Damian, Emma and Laura for making this event run smoothly.

Copies of presentations are available on the UKHEC website:

<http://www.ukhec.ac.uk/events/annual2002>.



UKHEC SEMINAR

GUIs for scientific computing

Mike Ashworth, Daresbury

A one-day UKHEC seminar on Graphical User Interfaces for Scientific Computing was held at Daresbury on 4th July 2002. The seminar addressed the issues of designing and developing GUIs for scientific applications, with the primary goals of improving ease-of-use, enabling computational steering and allowing the application to be marketed as a product.

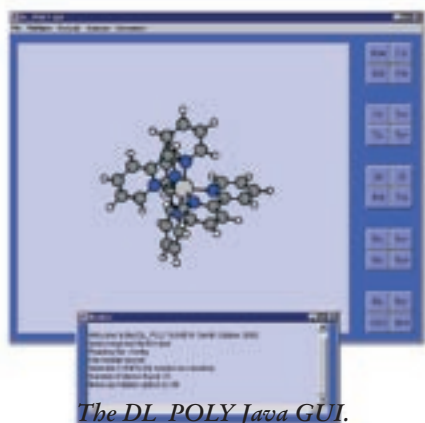
The major issues in building a GUI are:

- ease of download, build, use and learning
- ease/speed for manipulation of GUI objects and 3D graphics
- ability to link to native code (Fortran, C)
- quality of support for multi-threaded applications
- choice of toolkits
- choice of open-source solutions
- whether toolkits can be distributed as part of a package
- portability and cross-platform aspects

The programme included speakers who have been working on GUIs for a range of applications and using a range of different software toolkits, including Java, Python, VTK and AVS. There was no attempt at a consensus on the best available software environment, but those who attended benefited from the exchange of views and the comparison of experiences given in the presentations.


- Mike Ashworth (CLRC) 'Welcome and Introduction'
- Paul Sherwood (CLRC) 'Interfacing Quantum Chemistry Codes using Python'
- Barry Searle (CLRC) 'DLVisualize'
- Liz Potterton (University of York) 'CCP4i: The GUI for the CCP4 Macromolecular Crystallography Suite'
- Daniel Hanlon (CLRC e-Science Centre) 'Java techniques for GUIs'
- Bill Smith (CLRC) 'The DL_POLY Java GUI'
- Mike Ashworth (CLRC) 'The POLCOMS Java GUI'

One of the problems is that whereas everyone agrees that having a GUI for a scientific code is a great idea, there is very rarely funding for this work within regular research council grants, so it tends to be done very much as a *spare time* activity.



The DL_POLY Java GUI.

A technical report entitled 'Graphical User Environments for Scientific Computing' is also available from the UKHEC Web Site at:
<http://www.ukhec.ac.uk/publications/reports/guienv.pdf>



Deploying the UK e-Science Grid

Rob Allan, Daresbury

The UK e-Science Grid is being deployed with resources from the e-Science Centres and Centres of Excellence funded by the e-Science Core Programme receiving funding from the Research Councils and DTI. CLRC and the UKHEC sites are leading this activity through the

Grid Engineering Task Force which is part of the UK Grid Support Centre. Staff of all Centres are participating.

The present stage of deployment is known as 'Level 2'. This means that resources will have one or more application packages pre-installed and capable of offering a service with local accounting and tools for simple user management, discovery of applications and description of resources in addition to MDS. This Grid will be opened to users at the end of April 2003. CSAR computers are already on the Grid, and HPCx will follow soon.

At the previous Level 1, resources were registering with the UK Monitoring and Discovery Service (MDS) and capable of running the Grid Integration Test Script (developed at University of Southampton) to test core functionality of the Globus GT2 toolkit.

At Level 3, there will be an increasing community of resource providers and users on a production GT2 Grid,

widely accessible application base, distributed user and resource management, auditing and accounting. Resources signing up to Level 3 will be monitored to establish their availability and service level offered. Level 3 will also include a GT3 (Grid Services) testbed and the project will evaluate migration strategies for projects.

A lot more information about this large-scale distributed activity is available from the Grid Support Centre: <http://www.grid-support.ac.uk/etf>.

We have developed InfoPortal as a means of monitoring Grid resources and providing information on their status, see: <http://esc.dl.ac.uk/InfoPortal>.

HPCPortal, described in a previous issue of UKHEC Newsletter will be used as a client tool to invoke remote applications, see: <http://esc.dl.ac.uk/HPCPortal>.

A brief introduction to the portal tools is available at http://esc.dl.ac.uk/TechReports/ukhec_portal/ukhec_portal.pdf.

Web services, software engineering and the Grid

Rob Allan, Daresbury

Traditional software development, certainly in the area of scientific applications, has derived its modularity, maintainability, tunability and re-usability from subroutine or class libraries. Numerical libraries started to move into parallel and distributed computing in the mid 1980s, with examples such as MPI, ScaLAPACK and NetSolve. A new revolution is now sweeping software engineering circles, based on division of software into autonomous persistent services which can be maintained, distributed, replaced and invoked as required.

Emerging standards such as XML and SOAP are enabling the new generation of so-called 'Web Services' that allow systems to make remote procedure calls to other systems over the Internet. The remote procedures are self-describing (via the XML-based WSDL) and language independent, with toolkits available in Perl, Python, C++ among others – and of course Java.

Services are defined in terms of logical behaviour rather than physical system components. There could be more than one component for each service, so a service could be

an 'aggregation' of components. A service could likewise be an interface to a traditional numerical library on a parallel computer via wrappers.

The recent announcement of the Open Grid Services Architecture, OGSA, gives us insight into how the Grid API may also be converted to a set of persistent Web services and time-limited Grid service instances. With the main vendors such as IBM, Oracle and Microsoft backing this approach it is clearly important to understand how future application developers will use it.

UKHEC and the UK Grid Support Centre have written a Web Services Tutorial which is available at: http://esc.dl.ac.uk/TechReports/WebServices/webServices_doc.pdf.

An additional website providing more information and access to software is at <http://esc.dl.ac.uk/WebServices>.

UKHEC reports

The following reports have recently been published on the UKHEC website.

To obtain the reports, please visit:

<http://www.ukhec.ac.uk/publications/>

Mixed Mode MPI/OpenMP Programming by Lorna Smith and Mark Bull

Message passing codes written in MPI are obviously portable and should transfer easily to clustered SMP systems. Whilst message passing is required to communicate between boxes, it is not immediately clear that this is the most efficient parallelisation technique within an SMP box. In theory a shared memory model such as OpenMP should offer a more efficient parallelisation strategy within an SMP box. Hence a combination of shared memory and message passing parallelisation paradigms within the same application (mixed mode programming) may provide a more efficient parallelisation strategy than pure MPI. This technology watch report discusses the benefits of developing mixed mode MPI / OpenMP applications on both single and clustered SMPs. Section 2 briefly describes SMP and SMP cluster architectures whilst section 3 provides a comparison of the different characteristics of the OpenMP and MPI paradigms. Section 4 discusses the implementation of mixed mode applications and section 5 describes a number of situations where mixed mode programming is potentially beneficial. Section 6 contains a short case study, describing the implementation of a mixed mode application and comparing and contrasting the performance of the code with pure MPI and OpenMP versions. Section 7 describes a real mixed mode application code and section 8 describes related work on mixed mode programming.

Java Grande language comparison benchmarks on the IBM p690 system by Joachim Hein

In this study we compare the performance of Java, C and Fortran for a set of serial benchmarks on a p690 system. The performance of Java compares favourably with the equivalent Fortran and C versions. Some variation in run times was observed for all the benchmarks, due to different logical partitions of the p690 system not being fully independent. This variation is of the order of a few percent.

Visualizing Large Parallel Finite Element Simulations by J Leng, L Margetts

The Centre for Civil and Construction Engineering, UMIST and Manchester Computing have been collaborating on a project to develop a suite of parallel finite element programs for solving a range of large engineering problems. With huge data sets, visualizing and

interpreting the results of the simulations required some extensible software. AVS/Express and VTK were likely candidates. However, only AVS/Express supported the specialist data types being produced. Even so, the data could not be easily imported. A number of visualization modules and applications were developed to make the visual analysis of these FE engineering problems possible. This report also draws attention to the need for High End visualization capability.

Application Performance on High-End and Commodity-class computers by Martyn Guest

Beowulf clusters, on face value, offer the potential of a viable cost effective alternative for the provision of High Performance Computing. In this paper we compare the performance of Beowulf clusters built from commodity “off the shelf” components in the support of major research and production codes, with current high-end hardware such as the IBM SP, Compaq AlphaServer SC and SGI Origin 3800. The results concentrate on the application area of computational chemistry, materials science and computational engineering. Benchmark data on nine commodity-based systems (CS1-CS9) featuring Intel, AMD Athlon and Alpha CPU architectures coupled to traditional Beowulf interconnect, such as Myrinet and Ethernet, are presented. Furthermore, we provide performance data on systems utilising the SCALI SCI and Quadrics QSNNet interconnect technology.

Motivation and Requirements for a User Interface to a Computational Grid in the UK by Ali Anjomshoaa and Lorna Smith

In an area such as High Performance Computing (HPC), where the lifetime of application codes often exceeds that of most HPC machines, the ability to access these machines through a simple generic end-user interface is of real interest. By developing an end-user interface to a wide area computational Grid that includes HPC machines, we may be able to hide some of the complexity associated with the diverse and changing nature of the UK’s HPC systems and their administrative domains, from the users. This report examines the need for an end-user interface to a wide area computational Grid, which will be comprised of a heterogeneous environment of machines in diverse administrative domains, within various institutions across the UK. This need is first motivated, before a discussion of the requirements for an end-user interface is presented in the following sections. This discussion also addresses possible implementations of the end-user interface.

Bring on the new... sign up for the HPCx or CSAR newsletter

If you would like to be added to the mailing list for the new HPCx newsletter Capability Computing or CSAR's CSAR Focus, please supply your details as requested in the form below.

CSAR can be contacted at csar-advice@cfs.ac.uk

You can contact HPCx in various ways:

Email support@hpcx.ac.uk

Phone: 0131-650 5029

Fax: 0131-650 6710

Or fill in the form at

<http://www.hpcx.ac.uk/support/helpdesk/query.html>

Newsletter subscription

Your name _____ Job title _____

Company/institution _____

Address _____

Email _____

Telephone _____ Fax _____

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>

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>

MIKE ASHWORTH

Computational Science and Engineering
Department
Daresbury Laboratory
Daresbury
Warrington WA4 4AD
m.ashworth@dl.ac.uk
01925 603663
<http://www.cse.clrc.ac.uk>

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>

JON HUNT

Research Facilities Team
EPSRC
Polaris House,
North Star Avenue,
Swindon SN2 1ET
j.hunt@epsrc.ac.uk
01793 444263
<http://www.epsrc.ac.uk/hpc>

Main UKHEC webpage: <http://www.ukhec.ac.uk>