


ISSUE 2, NOVEMBER 2000

UKHEC

High-end computing



Output of the Ray Tracer benchmark from EPCC's
Java Grande Benchmark Suite.

Contents

- | | | | |
|---|---|----|--|
| 3 | <i>Additions to the UKHEC website
UKHEC annual seminar</i> | 9 | <i>Object Oriented Programming workshop</i> |
| 4 | <i>Java for high-end computing seminar
Improve your practical software
development skills</i> | 10 | <i>Technical report: Building a Beowulf</i> |
| 5 | <i>EWOMP 2000</i> | 11 | <i>Technical report: Grid-based high
performance computing</i> |
| 6 | <i>Grid workshop and seminar
WOMPAT 2000</i> | 12 | <i>Benchmarking Java against F90 and C++</i> |
| 7 | <i>11th Daresbury Machine Evaluation
workshop
HPC users' meeting</i> | 13 | <i>Projects to exploit SuperJanet4 and
Internet2 links</i> |
| 8 | <i>OpenMP standard evolves
Technology Watch report: Mixed Mode
MPI/OpenMP Programming</i> | 14 | <i>MRCCS summer school on Linux for
HPC</i> |
| 9 | <i>Technology Watch report: Programming
Models for Parallel Java Applications</i> | 15 | <i>EuroGrid project
Seminars and courses</i> |
| | | 16 | <i>Feedback form
Contact information</i> |

Editorial

Robert Allan, Daresbury Laboratory

This is the second issue of the UKHEC newsletter providing information about the research and technology evaluations of the UK High-End Computing collaboration of Daresbury Laboratory, EPCC (University of Edinburgh) and MRCCS (University of Manchester).

There has been a lot of interest recently in the Computational Grid for HPC (see pages 6, 11). We have therefore taken an active role in a number of projects on the Grid and related topics such as meta-computing. A technical report has been produced explaining the concepts and surveying some of the more widely used software and techniques. This is accompanied by a more general survey. Both are available from the publications section of the UKHEC website.

In addition to the Grid, other topics of interest highlighted in the newsletter include Java, Linux, Beowulf

clusters, OpenMP and mixed mode programming (OpenMP plus MPI for distributed systems with SMP nodes).

Several seminars and workshops have been organised to pass on this information to others. Major events included organisation and participation in the MRCCS Summer School on 'Linux for HPC', the Cray/SGI MPP workshop, the EWOMP2000 conference and the courses and workshops on code development and object-oriented programming.

For further information about our work please review the website and do not hesitate to contact the authors if you have further questions.


UKHEC main website: <http://www.ukhec.ac.uk>

Robert Allan



CENTRAL LABORATORY OF THE
RESEARCH COUNCILS

epcc



Additions to the UKHEC website: www.ukhec.ac.uk

Robert Allan, Daresbury Laboratory

We are continuing to populate the UKHEC website as our primary vehicle for dissemination of the results of our investigations and evaluations.

The Technical Report on Grid technology (see page 11) is available together with a collection of other Technical Reports and Technology Watch (TW) Reports which we have published to date:

- TW Report. *Mixed-Mode MPI/OpenMP Programming* (L. Smith)
- TW Report. *The Tera MTA* (K. Taylor). Available soon.
- TW Report. *Mixed Mode MPI/OpenMP Programming* (L. Smith and M. Bull, EPCC)
- TR. *Grid-base HPC* (R.J. Allan, J.M. Brooke, F. Costen and M. Westhead)
- TR: *How to build a Beowulf* (R.J. Allan, S.J. Andrews, M.F. Guest, P.M. Oliver, D.S. Henty, L. Smith, S. Telford and S. Booth). Available soon.

- TR: *Data Management Software* (C. Greenough, J.V. Ashby, K. Kleese, R.J. Allan, J.M. Brooke and D.S. Henty). Available soon.
- TR: *Code Development Tools* (M. Antonioletti, E. Breitmoser and R.J. Allan). Available soon.
- Special Report: *A Review of UK HEC Grid Infrastructure: state of the art and next steps*

The content of some of these reports is described in more detail in the following articles. There are also copies of conference papers and links to the web with pages on detailed research carried out by the collaborators. Some software, eg for Java benchmarking, is available for download, and we encourage you to try it.

Comments on the website and our material are welcome, as is feedback on particular areas of computing software or hardware technology which you believe we should investigate.

UKHEC Annual Seminar and Visualisation workshop

John Brooke (MRCCS)

UKHEC will be holding its Annual Seminar at the University of Manchester on Tuesday 5th December. There will be presentations from all the UKHEC sites, two each from EPCC and Daresbury and one from MRCCS.

The Annual Seminar is an opportunity for UKHEC to present to the user community a summary of the most important work carried out by UKHEC over the year. There will also be an invited Guest Speaker of international standing to present a topic of current interest in high end computing.

Provisional details of this year's seminar are now available on the UKHEC website. The event will be held in the Old Dining Hall of Hulme Hall, one of Manchester's oldest halls of residence. This is situated within walking distance of the University and is easy to reach by road or public transport. Travel details and booking details will be available via the web.

Visualisation Seminar: 4th December


The UKHEC Annual Seminar is arranged as a two-day event with a specialist seminar being linked to the Annual Seminar itself. Since MRCCS is hosting the event this year, the specialist seminar is on visualisation as this is one of the areas of particular interest for MRCCS within UKHEC.

We are very pleased to be able to offer UK researchers a rare opportunity to meet a team who are pioneering

techniques of eScience in bioengineering, linking the rapidly growing area of biological applications with techniques from the other physical sciences. The team (D.P. Bullivant, P.M.F. Nielsen, A.J. Pullan, P.J. Hunter from the Bioengineering Group at the University of Auckland) are giving a presentation on Computing the Virtual Human. The aim of the virtual human project (also known as the physiome project) is to construct a computational model that reproduces the physiological function and anatomical structure of the human body. This is a vast and ambitious project with a wide range of scales, from whole body to organ to cell to protein, and a wide range of processes, mechanical, electrical and biochemical. The Group's progress and plans will be described, with emphasis on the heart and on the tools for creating models and visualising the computational results.

Other talks by UK speakers will present the wide variety of uses of visualisation in physical science. There is a strong emphasis on creating problem-solving environments using visually-oriented tools. This is one of the important areas identified in the report of the UK Technology Watch Panel. The use of visual and virtual reality tools will increase rapidly within the field of high end computing; one only needs to look at the situation in the US and Japan to see this. This seminar provides an opportunity for UK users to explore the relevance of visualisation in their own work.

Java for high-end computing



Monday November 20th, 2000

The Cramond Room
The University of Edinburgh Conference and Training Centre
11 South College Street

Speakers:

ROLDAN POZO

Java performance analysis for scientific computing

MARK BULL

Java for high end computing: OpenMP and benchmarking

OMER RANA

Problem solving environment and legacy code integration with Java

MARK BAKER

MPJ: a reference message-passing system for Java

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

Improve your practical software development skills

One of the key roles of the UKHEC collaboration is to provide opportunities for continuing professional development for academic researchers in computational science and engineering. EPCC has prepared a course, Practical Software Development for Computational Scientists and Engineers, which is aimed at researchers who wish to improve their skills in practical software development. The next run of the course will be 29 February – 1 March 2001.

The first day of the course will cover Advanced Programming Tools and Techniques. Subsequent days will discuss Software Design and Software Project Development. The focus throughout will be on the

application of practical techniques allowing researchers effectively to develop and maintain high quality, portable code.

The previous two runs of this course have been very popular, with more than 30 attendees, so please contact us early if you are interested in attending.

Further information on this course is available at:
<http://www.ukhec.ac.uk/training/software>

To register, please email: EPCC-Support@ed.ac.uk



EWOMP 2000, the Second European Workshop on OpenMP was held at Murrayfield Conference Centre, Edinburgh on 14th and 15th September 2000.

EWOMP 2000 – the Second European Workshop on OpenMP

Mark Bull, EPCC

EWOMP was hosted by the EPCC in collaboration with Dr. Barbara Chapman of the University of Houston, Texas. The event attracted nearly 70 delegates from 10 European countries, Israel, the United States and Japan.

The technical programme, consisting of three invited presentations and 17 contributed talks, was of a very high standard, and initiated much discussion and debate, culminating in a lively panel session on the future of OpenMP. The topics covered included applications experiences, tools, language developments and implementations, mixed MPI/OpenMP programming, performance analysis and benchmarking.

The invited speakers were: Larry Meadows of Sun Microsystems, Jonathan Harris of Compaq Computer Corp. and Sanjiv Shah of Kuck and Associates Inc., representing the OpenMP Architecture Review Board.

EPCC gratefully acknowledges Sun Microsystems for sponsoring this event.



Above (left to right): Barbara Chapman (University of Houston), Jonathan Harris (Compaq Computers Corp), Mark Bull, (EPCC), Larry Meadows (Sun Microsystems), Sanjiv Shah (Kuck and Associates), Jim Cownie (Etnus LLC).

Electronic versions of the papers and presentations can be found at the workshop website:
<http://www.epcc.ed.ac.uk/ewomp2000>

Extended versions of selected papers from EWOMP 2000 and the sister workshop WOMPAT 2000 (see page 6) will be published in a special issue of Scientific Programming.

Grid workshop and seminar

Robert Allan, Daresbury Laboratory

UKHEC organised a seminar and PDRA workshop on computational grid technology on 21-22 June. This was open to all users of top-end computing and distributed resources who are interested to know more about how internet technology might be used in future scientific research (eScience).

The seminar, on 21 June, consisted of a series of presentations from those active in developing and evaluating computational grid systems.

The workshop, on 22 June, included an overview of current activities in the UK and a discussion about middleware and application requirements. It also showed potential end-users the steps involved in using computational grid software, such as PACX-MPI for meta-computing applications.

Grid-based computation is an exciting development which is intended to make computational resources more easily and flexibly accessible. The seminar and workshop featured leading international experts in Grid computing from the US, Europe and Japan.

The workshop was held concurrently with a Globus workshop and tutorial at the Rutherford Appleton Laboratory. Globus is the Grid middleware currently adopted for the majority of projects. This was primarily aimed at the needs of the particle physics community. We were able to benefit from video and NetMeeting broadcasts of the event, which will be made available via the workshop web pages.

Workshop on OpenMP Applications and Tools (WOMPAT 2000)

Lorna Smith, EPCC

Under funding from UKHEC, EPCC recently presented a talk on the development of mixed mode MPI/OpenMP applications at the 1st Workshop on OpenMP Applications and Tools (WOMPAT 2000), which took place on 6-7 July, in San Diego, CA.

This was a fairly small meeting, with around 40 people. Talks included application experiences, tools and compiler development, and mixed-mode programming.

An initial version of the OpenMP Fortran 2.0 Specification has recently been released for public comment (see <http://www.openmp.org>) and a number of representatives from the OpenMP Architecture Review Board were at the meeting to discuss this. One of the panel sessions focussed on this and considered a number of the new features of the standard.

<http://www.cs.uh.edu/wompat2000>

The presentations over the two days were as follows:

- Rob Allan and John Brooke. Introduction
- Steve Tuecke (Argonne National Laboratory). Overview of the Globus Project and Toolkit (via video link to Rutherford-Appleton Laboratory)
- Dr. Martin Tuori (Senior Product Manager, Platform Computing, Toronto). Grid Computing with the LSF MultiCluster
- Steve Tuecke. Globus User Tutorial (via video link to Rutherford-Appleton Laboratory)
- Hans-Christian Hoppe (Pallas GmbH). The UNICORE HPC Portal
- Yoshio Tanaka (ETL, Tokyo). Grid Software Development in Japan:
 - 1) Introduction to the Ninf project and its current status
 - 2) Globus enhancement in compliance with firewalls
 - 3) Sharing concepts and ideas of the ApGrid
- Thomas Boehme (Technical University Ilmenau). The Web Operating System
- Steve Tuecke. Globus Developers' Tutorial (via video link to Rutherford-Appleton Laboratory)
- Rob Allan and Steve Andrews (Daresbury Laboratory). Introduction to UKHEC and DisCo Grid trials
- Discussion of Technical Issues (see below)
- Edgar Gabriel (HLRS, Stuttgart). Using PACX-MPI for Metacomputing Applications
- Stephen Pickles (CSAR, Manchester). The Supercomputing'99 Metacomputing Demonstration
- Steve Tuecke - Globus Developers' Tutorial (via video link to Rutherford-Appleton Laboratory)

We would like to thank our sponsors for the UKHEC event: Quadrics Ltd., Hewlett-Packard and Pallas GmbH.

For further technical and background information see the workshop web pages at:
<http://www.dl.ac.uk/TCSC/UKHEC/GridWorkshop>.

These contain a complete set of the presentation material, including the Globus tutorials. We have also recorded a set of videos of the Globus presentations, which will be prepared for the web soon. Please contact Rob Allan for further information:
r.j.allan@dl.ac.uk.

Following this seminar and workshop we have set up a Grid working group for HPC applications. So far this is in its early stages, but information will be posted on the web site as it continues. See:
<http://www.dl.ac.uk/TCSC/UKHEC/GridWG>.

Related activities at CLRC, via a virtual 'eScience Centre' are being coordinated to help the science community realise the potential of the Grid and related techniques.

11th Daresbury Machine Evaluation workshop

Martyn Guest, Daresbury Laboratory

The 11th Daresbury Machine Evaluation Workshop will be held from 22–24 November at the CLRC Daresbury Laboratory as part of the EPSRC's Distributed Computing Support Programme. This Workshop is now established as a leading national event dedicated to distributed high-performance scientific computing. The principle objective is to encourage close contact between the research communities from the Mathematics, Chemistry, Physics, Engineering and Materials Programmes of EPSRC and the major vendors of mid-range computing systems, workstations, servers, software and peripherals.

You are invited to attend this year's workshop, which will have a format similar to that of previous years. About a dozen major vendors will make 30-minute presentations on topics such as hardware, compilers, graphics, storage, networking etc. In past years the audience has been very appreciative of the technical content of these talks and we have encouraged the speakers nominated by the vendors to focus their presentations accordingly. This year we have chosen to extend the workshop to three days, so that we can fully include cluster computing (Beowulf) systems.

An important component of the workshop is the availability of systems for benchmarking evaluation

purposes, plus the exhibition itself. Vendors are providing Internet access to these systems prior to the event, and the availability of products on the day will help delegates evaluate the machines. We hope to make loaned systems provided by the vendors available on the Internet from 15th November onwards. Note also that we will be producing formal Workshop proceedings, a copy of which will be distributed to each delegate.

The programme will start on the morning of 22 November, and finish around 5pm on the 24th. We would be grateful if you would confirm your attendance before 31st October 2000. Additional information may be obtained from the workshop secretary, Mrs Shirley Miller, or Damian Jones via: Machine_Evaluation_Workshop@dl.ac.uk

The website has an on-line registration form and will also contain up-to-date information on exhibits and benchmarking systems as it arrives.

We look forward to hearing from you, and hopefully meeting at the event. Please pass this information to anyone else who you think might be interested.

Programme including up-to-date speaker information:
<http://www.cse.clrc.ac.uk/Activity/DisCo+1177>

HPC users' meeting

Robert Allan, Daresbury Laboratory

The Research Councils held their annual joint HPC Users' Meeting, organised chiefly by EPSRC, at the Queen Elizabeth II Conference Centre, Westminster, on 19 September. Some 130 people attended the event and the three centres, Daresbury Laboratory, University of Edinburgh (EPCC) and University of Manchester (MRCCS) presented material both separately and collectively as the UKHEC Collaboration.

There were talks on infrastructure and scientific requirements from: John Taylor (DGRC, IT, Informatics HPC and the Grid), Spencer Sherwin (IC, Direct Simulation of Geometrically Complex Flows), Mike Gillan (UCL, Ab-initio Calculations on Very Large Systems), Paul Jeffreys (CLRC, The LHC Computing Challenge), Rod Hubbard (York, Understanding and Exploiting Biology), Sid Karin (SDSC, Future of HPC in the USA), Ron Perrott (Belfast, The European Grid), Martyn Guest (CLRC, UKHEC and the Grid).

The meeting was an opportunity for the attendees to consider future policy for academic computing provision,

including the Grid and electronic access to other services – so-called eScience. A policy document, *High-end Computing: Strategic Issues for the UK*, was distributed for consideration. This is available from EPSRC for anyone who did not attend the meeting.

The meeting ended with the announcement of HPC2000, which is to begin its tendering exercise in December and run for 12 to 15 months, resulting in the procurement of a 5 Tflop/s system by April 2002 at the latest. The final discussion focused on user requirements around this new HPC system, and related Grid technology.

Material from the meeting will be made available via EPSRC's website soon.

Naturally UKHEC has spent a large amount of its effort so far investigating eScience, the Grid and underlying technology. A report on its work entitled *Proposal for a UK HEC Grid Infrastructure: state of the art and next steps* was presented to the Research Councils via EPSRC on 1 June.

OpenMP standard evolves

Robert Allan, Daresbury Laboratory

The OpenMP Architecture Review Board has carried out a period of public comment on the next release of the Fortran Application Program Interface (API), v.2.0. This comment period closed in September and the final API is expected to be available by Supercomputing 2000 in November.

Major new features expected to remain in the 2.0 specification include:

- Array reductions
- Parallelisation of F90 array syntax
- COPYPRIVATE for broadcast of sequential reads

- Privatisation of module data
- Portable timing routines
- Nested locks
- Control of the number of threads for multi-level parallelism

In addition to the above, several minor features have been added and many interpretations addressed and incorporated into the text.

OpenMP website at: <http://www.openmp.org>

Technology Watch report

Mixed Mode MPI/OpenMP Programming

Lorna Smith and Mark Bull, EPCC

Shared-memory architectures are gradually becoming more prominent in the HPC market, as advances in technology have allowed larger numbers of CPUs to have access to a single memory space. In addition, manufacturers are increasingly clustering these SMP systems together to go beyond the limits of a single system. As clustered SMPs become more prominent, it becomes more important for applications to be portable and efficient on these systems.

Message-passing codes written in MPI are obviously portable and should transfer easily to clustered SMP systems. However, it is not immediately clear that message passing is the most efficient parallelisation technique within an SMP box, where in theory a shared memory model such as OpenMP should be preferable. 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.

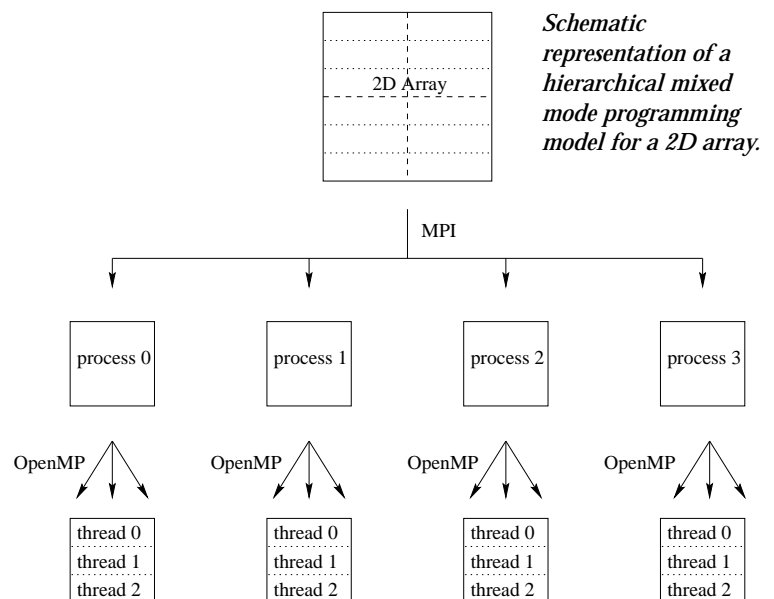
The majority of mixed-mode applications involve a hierarchical model, MPI parallelisation occurring at the top level, and OpenMP parallelisation occurring below. For example, Figure 1 shows a 2D grid which has been divided geometrically between four MPI processes. These sub-arrays have then been further divided between three OpenMP threads. This model closely maps onto the architecture of an SMP cluster, the MPI parallelisation occurring between the SMP boxes and the OpenMP parallelisation within the boxes.

As part of UKHEC's commitment to monitor and disseminate current developments in hardware and software, a Technology Watch Report has been written on Mixed Mode MPI / OpenMP Programming. This report discusses the benefits of developing mixed-mode MPI/OpenMP applications on both single and clustered SMPs, examining:

- codes which scale poorly with MPI (e.g. due to load

- imbalance or fine grain parallelism problems);
- replicated data codes which suffer from memory limitations;
- applications which run on a restricted number of MPI processes;
- systems with unoptimised and poorly-scaling MPI implementations.

Performance figures are presented for a simple test case code and for a large application code, with MPI, OpenMP and mixed MPI/OpenMP results compared and contrasted.



Mixed-Mode Programming Technology Watch Report by Lorna Smith, EPCC, may be found at: <http://www.ukhec.ac.uk/publications>

Programming Models for Parallel Java Applications

Mark Bull and Scott Telford, EPCC

Java has yet to make a significant impact in the field of traditional scientific computing. However, there are a number of reasons why it may do so in the near future. Perhaps the most obvious benefits are those of portability and ease of software engineering. The former will be particularly important when grid computing comes of age, as a user may not know when they submit their job, what architecture it will run on. Automatic garbage collection, thorough type checking and the absence of pointers make Java development significantly less error prone than more traditional languages such as C, C++ and Fortran.

Of course, using Java is not without its problems. Perhaps the prime concern for scientific users is performance, though just-in-time compilers are making rapid advances in this field. Other issues include lack of support for complex numbers and multi-dimensional arrays. A lack of suitable standards for parallel programming is also a concern: this technology watch report discusses the current status of parallel programming models for Java.

One fundamental characteristic of any Java-based parallel system is the distinction between single virtual machine (single-VM) environments and multiple virtual machine (multi-VM) environments. A single-VM environment, as the name suggests, provides a single Java virtual machine, running a single multi-threaded Java

application, with the Java threads distributed across the processors in the system. A multi-VM environment provides multiple Java VMs, each running a discrete Java application, with some form of inter-VM communication facility allowing interaction between the applications running on different VMs.

The report reviews programming models for both these types of environments: Java threads, directives and auto-parallelisation for single-VMs, and Remote Method Invocation, Message Passing and the Virtual Interface Architecture for multi-VMs. Data parallel Java dialects are also discussed.

Code for a simple CFD kernel is included, using three of the models described above: Java threads, mpiJava (a Java binding to MPICH), and JOMP, a shared-memory directive-based system in the spirit of OpenMP, which is currently under development at EPCC. Performance of the parallel Java versions is compared to a Fortran90 + OpenMP version of the code.

JOMP details and documentation, and to download a trial version of the software:

<http://www.epcc.ed.ac.uk/research/jomp>

Programming Models for Parallel Java Applications may be found here: <http://www.ukhec.ac.uk/publications>

Object Oriented Programming workshop: Advanced Software Design with Objects

Alan Simpson, EPCC

As part of UKHEC's ongoing workshop series, EPCC ran an event called 'Advanced Software Design: Object Oriented Programming' on Friday 6 October, immediately after the run of 'Practical Software Development.' These workshops are intended to be interactive and informal and aim to provide a mechanism for keeping UK researchers up-to-date with the latest techniques.

The workshop was divided into two sessions – an Object-Oriented (OO) tutorial in the morning and an applications focus in the afternoon. The tutorial was produced by Geoff Cawood (EPCC) who has significant OO experience both in industry and in academia. The tutorial included OO concepts and programming tips and examples, as well as a video of the history of OO. This session provoked useful and interesting discussions on a variety of topics.

The speakers for the afternoon session were John Brooke (MRCCS), Stephen Booth, Elson Mourao, Lorna Smith and Kira Smyllie (all EPCC). The talks covered

applications of OO techniques in C++, Java and F90, as well as a discussion on the use of the Unified Modelling Language.

The workshop was well attended with over 25 researchers from as far afield as Southampton, Newcastle and even Strathclyde, as well as representatives of each of the UKHEC sites. The many interesting contributions from the audience made this a valuable day and we were pleased that the workshop had such a positive atmosphere.

For information on the next UKHEC Workshop on Visualisation, see page 3. Also, if you are interested in OO programming, you may wish to attend the UKHEC Seminar on Java which EPCC will be running on 20 November (see announcement on page 4).

Further information can be found in the article on page 4 or via <http://www.ukhec.ac.uk/events/javahec>.

Building a Beowulf

Robert Allan, Daresbury Laboratory

Beowulf-class systems or 'piles of PCs' are rapidly becoming the most cost-effective way to build mid-range computing resources. These systems are sometimes referred to as 'clusters of workstations' (COWs), 'networks of workstations' (NOWs), 'commodity off the shelf' (COTS). The principle is that they should use commodity components and cheap widely-available software, e.g. based on GNU and Linux, enabling mid-range parallel computers to be constructed for a very modest cost.

One of the first projects, The Beowulf, was started at the NASA Goddard Space Flight Center in 1994 for image processing. Commodity cluster systems are therefore often known as Beowulf-class computers. It is not clear why the name was originally chosen, so we use the name Beowulf for everything, whether or not it has a Pentium processor.

As well as the hardware components, the software (including MPI, Linux, GNU, compilers and batch schedulers) is evolving rapidly, as is the total cost effectiveness. New processors and network components are also constantly emerging for evaluation. Understanding systems hardware and the software necessary for building, administering and using a Beowulf-style system is part of the remit of the EPSRC-funded DisCo programme.

A cluster computer could be built using the PCs or workstations on people's desks, connected by a standard network. Software could make use of idle CPU cycles by distributing tasks across such a cluster. This is very successful in commercial environments, and LSF from Platform Computing is currently the software of choice to manage such a system.

A Beowulf is built from a collection of commodity CPU, memory and disk components, probably purchased as a PC system unit or workstation minus keyboard, mouse and screen, and connected together via some kind of private network. There is usually a 'master' or 'head' node, which may not be identical to the 'worker' nodes, and which has an additional network interface card (NIC) by which it is connected to a public network for remote login. This enables one to get into the machine, either for system administration or for preparing data, compiling codes and submitting jobs. The master will have a keyboard, mouse and screen, which could also be plugged into a worker node for testing or connected via a switch. A stand-alone system does not necessarily need public access. In addition the master node will have a CD-ROM drive to install software, e.g. Linux, and all the nodes will have a 3.5" floppy drive for software installation. It is worth thinking about the sizes of hard disk required on the master and worker nodes and how the master node can be backed up to preserve data, operating system and

codes. Finally the entire system will need careful planning in terms of network configuration, number of switches, wiring etc. and should be mounted in some kind of robust rack. Careful layout and numbering of the components, perhaps colour coding the cables, will also be useful for any future maintenance.

With a high rate of change in the field there is a bewildering choice of processors, motherboards, interconnects, resource management tools etc., and it is vital for potential owners to be well informed on the pros and cons of individual components and their combinations, both in general and for the envisaged application.

By far the most popular CPU choices are the Pentium and Alpha, although the AMD Athlon is often considered and IA-64 based architectures are in the pipeline. The logic behind using commodity processors is not hard to see, given their ubiquitous presence in a variety of vendor platforms and their creditable performance against many proprietary RISC processors in floating point applications (see for example <http://www.cse.clrc.ac.uk/Activity/DISCO+921>). When considering an operating system the overwhelming preference is for some flavour of UNIX, either proprietary or, increasingly, Linux. Again there is a critical mass forming here, leading to an installed base of tools and devices, which will create a driving force towards an ad-hoc standard.

Interconnect subsystems (networks) can be classified as cheap but slow, or fast but expensive. Fast (sic) Ethernet is almost universally used for the former but is currently limited in both latency and bandwidth. This is also a fast-changing market.

In order to achieve a 'balanced' system it is usual to try to match the deliverable processor performance to the network bandwidth. For example, a 400-500 MHz Pentium delivering 100-200 Mflop/s on typical applications will work with a 100 Mb/s Fast Ethernet, whereas a 667 MHz Alpha typically delivering over 500 Mflop/s will require something better.

In this report we consider a number of issues which must be addressed, both in hardware and software, to build a dedicated system which could be used to run a parallel program. See the UKHEC website under 'Publications'.

Links to many Beowulf clusters: <http://www.beowulf.org>

The DisCo programme at Daresbury:
<http://www.cse.clrc.ac.uk/Activity/DisCo>

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

Grid-based high performance computing

R.J. Allan (Daresbury), J.M. Brooke and F. Costen (MRCCS) and M. Westhead (EPCC)

Grid computing describes the linking together of distributed computational resources to provide flexible access and a common interface for the user. This is sometimes referred to as an eServices environment in commercial applications. Meta-computing extends this concept to enable distributed systems to compete with supercomputers in order to try to overcome the limitations of a single computing system. To achieve these goals software systems must be provided which use Internet technology, now common in eCommerce, for the benefit of the computational science community. This has recently been referred to as eScience.

Distributed computing systems offer more than just a large CPU resource. A software environment of unprecedented quality and functionality is emerging along with the use of the Internet for commercial and leisure purposes. This is driven by a combination of the computer industry and the loose collection of worldwide 'freeware' programmers. Geoffrey Fox has referred to this as the 'Distributed Commodity Computing and Information System'.

In the USA and Japan there are several alliances of computing centres separated by large distances (e.g. NPACI). In Europe, Germany has taken a lead because of the regional computing centres. In the UK the JREI-funded centres may be (but are not yet) a source of similar resources.

The whole concept is often referred to as a 'computational grid'. Computers on a grid can solve very large problems requiring, for instance, more main memory than is available on a single machine. However, resources are typically heterogeneous in nature with fast internal message passing or shared memory, but a slow wide-area network connecting the separate computers. The implementation of a message-passing standard (e.g. MPI or PVM) on a meta-computer must therefore be able to take into account different architectures and hierarchical networks with differing bandwidth and latency, in a way transparent to the user. Issues of security, resource management and load balancing are also fundamental to successful exploitation.

A key concept is that of 'ownership'. A grid is a 'federation' of resources which can be accessed in a transparent way by authenticated grid users. This raises the fundamental question of 'accounting' for resource usage, whether it be CPU time, disk, memory, licensed

software or preserved data. Whilst this is perhaps the most important issue to be considered in implementing a national grid environment we do not consider it further in this report. Instead we focus on how the scientific user might benefit from such an ideal environment.

Grid-based computing is likely to become an important key technology for future UK high performance computing. In this technical report we give a very brief description of some current international developments. Further information is available from larger technical surveys carried out by the collaborators, and from the book by I. Foster and C. Kesselman. We do not attempt to provide an introduction to all the

underlying distributed-computing techniques, which are both complex and diverse. There are numerous discussions in the computer-science literature which should be consulted for background information.

A more detailed description is given of some specific computational grid environments and tools:

- LSF from Platform Computing
- Globus from Argonne National Laboratory and the University of Southern California
- UNICORE from Pallas and Jülich Rechenzentrum
- Seamless Thinking Aid from ETL, Japan
- Jini Legion
- StaMPI from ETL
- PACX-MPI from Stuttgart
- MPICH-G from Argonne National Laboratory

In addition to surveying ongoing world-wide research activities the UKHEC Collaboration is involved in practical experiments to demonstrate the utility of a grid computing environment. CSAR in Manchester has already been actively involved in a meta-computing demonstration which won a prize at SuperComputing'99 last November. This was summarised in CSAR Focus, edition 4 and is briefly described in the report. It is hoped to make this meta-computer available in a persistent manner for other UK researchers to evaluate via a funded Internet-2 project (number 10): A Global Meta-computer for UK Applications.



Benchmarking Java against F90 and C++

Y.F. Hu, R.J. Allan (Daresbury) and K.C.F. Maguire (European Southern Observatory)

Java is now widely recognised as a good object-oriented language for writing portable programs quickly. However, its penetration to computationally-intensive numerical calculation is still low. One of the main reasons is its poor performance, or the perception of it, for such numerically intensive computing. On the other hand, Java is very attractive for scientific and engineering calculations that involve many researchers using different platforms, or over the Internet as a Computational Grid.

That Java can suffer from performance problems is perhaps not at all surprising, due to its emphasis on reproducibility of results on different platforms, safety of execution, and ease of programming and testing. It is also a relatively new language and its compiler technology is still evolving. Java also differs from C and Fortran in that Java codes are first compiled into bytecodes and then interpreted on any platform, using a Java Virtual Machine (JVM). However, there are already major advances in the Java compiler technology, including JIT (Just-In-Time) compilation, HotSpot technology, and effort to reduce the number of array-bound checks by code transformations.

We have compared the performance of Java with C and F90 for sparse matrix based calculations, in terms of the computing speed as well as I/O speed. In addition we have also compared Java with C using the SciMark2 benchmark. A summary of our findings is given below, and a more detailed report can be found on the JASPA webpage (see below).


The experiments were carried out on the following systems:

- dual-processor Pentium II with both the IBM and the Sun 1.3.0 compiler
- Sun Ultra 80 with the Sun 1.3.0 Java compiler
- IBM Power3 with Java 1.1.8 compiler.

The benchmark we used is a simple yet important operation: that of multiplying two sparse matrices. This operation appears when forming the normal equations of interior point methods for large scale numerical optimisation, as well as in large scale unstructured calculations where a multilevel/multigrid scheme is used. The benchmark is known as JASPA (JAva SPArse benchmark), and is available via the JASPA webpage.

Comparison of Performance

The figure summarises the key findings when comparing Java with C and F90 using the JASPA benchmark. Similar conclusions can be drawn from the SciMark2 benchmark when comparing the performance of Java with C.



Conclusions

The I/O and computing performance of Java has been found to be between 30 per cent to about 3 times slower than F90, depending on the platforms and the compilers (and the JVMs). The best relative performance is achieved on a Pentium II, where the IBM Java yields code that is about 30-40% slower in computing and I/O performance. The performance of Java seen in this work is certainly a big improvement over the situation only 1-2 years ago when it was generally perceived that Java could only achieve, say, 20% the speed of Fortran and C.

Currently a number of factors still make it difficult for Java to be used in large scale scientific computing. These include its lack of efficient multidimensional arrays, complex number support or operator overloading, its inability to take advantage of fused multiply-add and the associativity of operations in compiler optimisations, its lack of a Math Library that produces the same results on all Java platforms, and the difficulties in interfacing Java with other languages. However, with many activities to address these issues, and with improving compiler technology, we may be closer to a time when new scientific and engineering applications could equally well be written in Java to benefit from its advantages, without the need to worry about the loss of performance.

SciMark2 benchmark:
<http://math.nist.gov/scimark2>

JASPA: <http://www.dl.ac.uk/TCSC/UKHEC/JASPA>

Projects to exploit SuperJanet4 and Internet2 links

John Brooke (MRCCS)

Background

The UK backbone academic network SuperJanet is being upgraded from SuperJanet3 to SuperJanet4 starting in December 2000 and continuing through 2001. This will have a major impact on the way that scientific computation develops in the UK, in that it provides the networking speeds necessary for the collaborative scientific working described as eScience by John Taylor in the keynote speech to the HPC Users Meeting on September 19 (see the report by Rob Allan in this issue).

The essential difference that the upgrade will bring is a move from a backbone at 155 Megabits/s using ATM technology to at 2.5 Gigabits/s backbone using SDH (Synchronous Digital Hierarchy). The technology switch means in particular that the speed of the backbone can be upgraded without having to replace the fibre connections. In fact, the current plans are to upgrade to 20 Gigabits/s by 2002.

This will be of vital importance to the fledgling Grid and eScience projects currently being investigated by UKHEC and others, e.g. the DataGrid based around the High Energy Physics community. Similar technology is already being deployed in the US and given the important advances being made in scientific exploitation of what is called Internet2, JISC (Joint Information Systems Committee) have decided to sponsor projects in the UK that exploit links between the UK and US to provide the benefits of technology transfer and collaborative working. Eight projects have been announced. For details see: http://www.ja.net/development/internet2/collab_projects.html

In this article we describe the project which has direct relevance to high-end computing and which provides the UK high-end computing community with an opportunity to try out eScience and collaborative ventures on a global scale.

Importance of a global metacomputer

The essence of this project is to make available to the JANET community a global supercomputer spanning Europe, United States and Japan. The two hubs of the metacomputer will be at University of Manchester (UoM) and Pittsburgh Supercomputer Centre

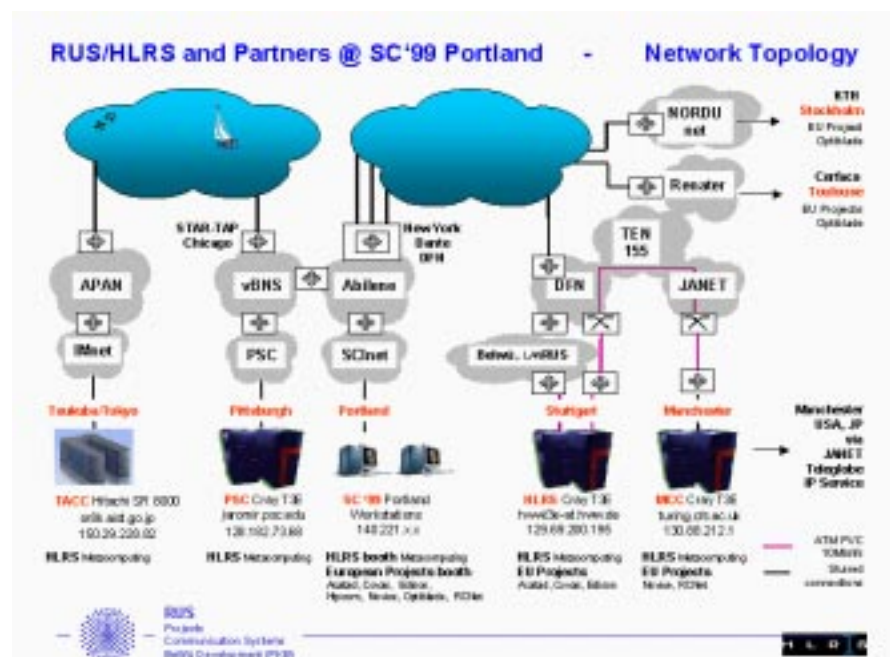
(PSC) and the topology is shown in the figure. Other sites can join the global metacomputer for specific experiments but the core system will be made persistently available. The project will work with users in the computational science community to enable them to transfer their applications to the global metacomputer, providing an opportunity for some ground-breaking computational experiments. We have a variety of applications already identified which will test the capabilities of the international links, particularly with regard to QoS issues. Connectivity with other international networks is one of the major thrusts of the Internet2 networking tracks and this project will drive Grid computing forward on both sides of the Atlantic and the Pacific.

A global metacomputer was demonstrated at Supercomputing 99 and will be demonstrated at Supercomputing 2000. Four applications were run, from the fields of molecular dynamics, computational fluid dynamics, processing of data from a radio telescope and collaborative working via visualisation. The project aims to make this global computing network persistently available from October 2000 to October 2001.

Other UK sites can also join the metacomputer via SuperJanet4 and it is intended to invite participation from other European, US and Japanese sites for the showpiece demonstrations. The aim is to test scientific applications involving collaboration over trans-continental networks, and applications which need exceptional processing power not available in a single machine.

Continued overleaf

The network topology of the proposed global metacomputer. The two 'ponds' of water shown in the diagram at the Atlantic and Pacific Oceans.





MRCCS summer school on Linux for HPC

Robert Allan (Daresbury Laboratory)



The Summer School on Linux for HPC, 4-15 September, was jointly organised by MRCCS and Daresbury Laboratory, using facilities at Manchester Computing and Daresbury. These included a motley collection of obsolete PCs borrowed from the Manchester University Engineering Department and Library. Students were expected to assemble these into a complete Beowulf system by the end of the course!

A number of speakers gave tutorial presentations over the two weeks of the School. These included Doug Johnson (San Diego Supercomputer Centre), Neils Walet (UMIST), Cos Ierotherou (University of Greenwich) and Rob Allan (Daresbury). Other individual technology presentations were given by:

Rob Allan (Daresbury), Niels Walet (UMIST), Werner Krotz-Vogel (Pallas, Bonn), Bill Sparks (SGI, USA), Cos Ierotherou (U. Greenwich), Doug Johnson (Ohio Supercomputer Center), Keith Taylor (MRCCS), Richard Watkins (Platform Computing), Jonathan Nash (Synchron, Oxford), Nick Davies (IBM), John Merlin (Portland Group), Robin Harker (Workstations UK), Jon Taylor (Quadrics), Martyn Guest (Daresbury) and Einer Rustad (Scali).

The practical work was led very competently by Chris Mueller whose enthusiasm made the event a great success. We also thank the other staff at Manchester who handled registration, catering etc. and in particular Gareth Poulton.

The students, of whom there were some 18 from all over the world (including Australia), were able to set up the master PCs starting from scratch and install Linux Red Hat v6.2. This was followed by an exercise in networking and installing the basic Linux kernel onto slave PCs (one master and 3 slaves were provided for each group of 3 students). Other software, such as gnu f77, PGI f90, MPICH, PVM and PBS, was also installed so that example and real programs could be run in parallel. With more help from Chris the individual clusters were transformed into a 24-processor Beowulf system with a grand master before the end of the School.

This cluster has been moved and is currently available on-line. Contact Dan Kidger (d.kidger@mcc.ac.uk) for further information.

Penguin by Larry Ewing (lewing@isc.tamu.edu).

Further information is available from the website at:
http://www.man.ac.uk/mrccs/summer_school/2000

SuperJanet4 and Internet2 *continued*

Benefits to the HEC Community

Grid computing in the UK is in a fledgling state and there are very few applications that can benefit from the use of distributed computing power over the Internet. There are some very ambitious projects attempting to change this, notably from the High Energy Physics (HEP) community; however, the timescales of this project are such that significant results are unlikely within the next year or two. The HEP projects look towards the year 2005 when the Large Hadron Collider results will come on-stream from CERN.

The global metacomputer, on the other hand, has been demonstrated to work already. The project links with international partners who have acknowledged expertise in this field and this will benefit the UK projects involved. It is intended also to link the global metacomputer to the grid network being developed by collaboration between the UKHEC sites and RAL. Other UK sites can be

included as required and subject to the technical challenges being successfully worked on.

Pittsburgh Supercomputing Centre (PSC)

The SuperJanet4-Internet2 projects are encouraged to work with a US partner currently involved in Internet2 projects. The partner for this project is Pittsburgh Supercomputing Centre, which has extensive involvement with US projects using High End Computing and is a major node on the US research networks.

For more information on Internet2 projects in general and PSC in particular see:
<http://www.internet2.edu> and <http://www.psc.edu>

For more information on the global metacomputer and details of contacts if you wish to try working on it, see:
http://www.man.ac.uk/mrccs/global_supercomputing

EuroGrid – uniform access to European computational resources

John Brooke (MRCCS)

The IST programme of the EU Fifth Framework call has funded two major and complementary Grid projects. The DataGrid project will centre around the need of the High Energy Physics (HEP) community to deal with the enormous quantities of data in a structured and effective manner. This project will be centred round the Globus software and will examine issues of multi-site collaboration, analysis and storage of data. Many of the techniques developed will have relevance to other scientific and commercial users of large volumes of data.

The sister EU project, EuroGrid, will have a different focus. The aim here is to facilitate the use of high end computing facilities across Europe. It is based on the Unicore software which is described in the EuroGrid documentation.

The UNICORE project was originally proposed to the BMBF, the German Ministry for Research and Development, in 1997. The aim was to create a software layer focusing on the practical needs of users of high performance computer (HPC) centres, and not to repeat the development of earlier US projects like Globus or Legion. To solve large problems in computational science the users have to use remote systems, access large volumes of remote data and exploit different architectures. For example, some parts of a complex problem can be run effectively on a massively-parallel system, others work best on a vector machine. In practice, it is difficult to use other resources of multiple systems, especially at multiple sites, since it is very time-consuming for the users to acquaint themselves with the different

environments. As a consequence, results are not obtained in a timely fashion or even at all, and the available HPC resources are often sub-optimally utilised. The UNICORE research and development project has set out to create a seamless, secure and intuitive access to distributed HPC centres while leaving the established policies and practices at the centres intact.

Thus it can be seen that Europe will have two contrasting developments in Grid technology. There will clearly be synergy between these two major projects, and it is to be hoped that Globus and Unicore will influence each other to create Grid middleware with the best characteristics of both. UKHEC is well placed to offer the latest techniques developed by both these projects, since it has links with DataGrid (via CLRC involvement in the UK side of this project), and with EuroGrid since MRCCS is the UK node of EuroGrid and will be developing a Resource Broker that can enable UK users to access computational resources on a European scale. These two projects can also benefit from the developments in Java technology currently being explored at EPCC. Indeed all these developments are described in the UKHEC Report on Grid Computing available from the UKHEC website.

Information about developments in GRID technology will be monitored by UKHEC and made available by all the usual means of dissemination. The EuroGrid project will be launched in November 2000, and the DataGrid project early in 2001.

Seminars and courses

The three UKHEC centres also organise a regular series of seminars and courses. Anyone is welcome to attend, but please let organisers know in advance.

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

Forthcoming events

13 NOVEMBER (CSAR)
Shared Memory Parallelisation with OpenMP

20 November (EPCC)
Java for High End Computing

22–24 NOVEMBER (Daresbury Laboratory)
11th Machine Evaluation Workshop and Beowulf Cluster Workshop
<http://www.cse.clrc.ac.uk/Activity/DisCo>

28 NOVEMBER (CSAR)
SGI Origin2000 Optimisation

4 DECEMBER (MRCCS)
UKHEC workshop: Scientific Visualisation

5 DECEMBER (MRCCS)
UKHEC Annual Seminar
<http://www.ukhec.ac.uk>

- Daresbury: http://www.cse.clrc.ac.uk/Activity/Daresbury_CSE_Seminars
- EPCC: <http://www.epcc.ed.ac.uk/epcc-tec>
- MRCCS: <http://www.man.ac.uk/mrccs/seminars>
- UKHEC: <http://www.ukhec.ac.uk/training>
- UKHEC: <http://www.ukhec.ac.uk/events>

7 DECEMBER (CSAR)
Parallel Vector Processing on Fujitsu VPP 300

15 DECEMBER (CSAR)
Introduction to High Performance Computing

18 DECEMBER (CSAR)
Fortran 90

27 FEBRUARY (EPCC)
Practical Software Development for Computational Scientists and Engineers

18–2 APRIL 2001 (University of Durham)
Royal Society of Chemistry
Faraday Discussion 118, on Cluster Dynamics
<http://www.rsc.org/pdf/confs/faradisc/fara118.pdf>

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>

JON HUNT

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

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

This site is under construction and will be used to provide our main source of information.