epcc news

The University of Edinburgh Issue 45, MARCH 2002

CONTENTS

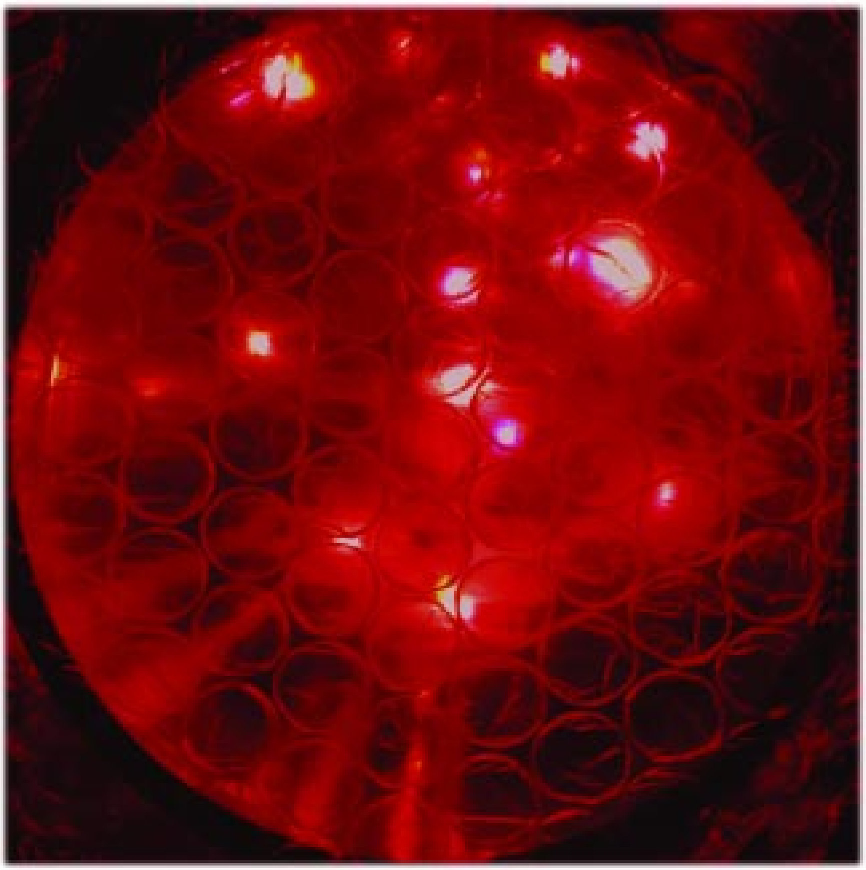

Pushing back the boundaries – research at EPCC

Editorial

JUDY HARDY

Welcome to this research-themed issue of *EPCC News*. EPCC has always been a focus for research, both through the support of work carried out on our HPC facilities and through the development of new technologies for high performance and Grid computing. In this issue you will find articles on many aspects of the research work that is carried out here.

We believe that Grid computing is a vital new technology for science and commerce. We have articles on a range of our Grid-related activities, including the Grid service requirements survey undertaken as part of the ENACTS programme, a discussion of the simulation of internetwork traffic and the recent Global Grid Forum (GGF3). EPCC has a leading role in the benchmarking group of the Java Grande Forum and our work in this area was showcased at Supercomputing 2001, where we ran a very successful tutorial on Java for HPC. The activities of the recently formed National eScience Centre, a centre of education and research in Grid computing, are also presented here. You can even find out what the Access Grid really is – and what it's used for!

Of course, our traditional commitment to providing support to UK and European academia continues. We have articles on a number of Grand Challenge problems – research to simulate flow in colloids, to model the world's oceans and even to model the fundamental forces of nature. At the European level, our TRACS programme has now welcomed over 350 researchers from many European centres and we present an example of their research here.

We hope that you enjoy this issue and that we have given you a taste of the range of research activities undertaken by EPCC, by our collaborators and by the research communities that use our HPC systems.

The Access Grid at EPCC

JOSEPHINE BEECH-BRANDT

Imagine being able to enter a virtual room to join a seminar being given in a different country by a world-leading authority. Or to meet with collaborators scattered around the world who you don't often see. Or to find your colleagues analysing results from a collaborative experiment. Or even just to sit around to have informal chats with people in various countries. Imagine doing all this without having to leave your building. Well, you no longer have to imagine because you can do it now using the Access Grid (AG).

An AG node consists of multimedia and computational equipment including video cameras, data projectors, microphones, an echo canceller, four computers and a very large screen.

There are currently over 80 fully operational AG nodes all around the world and numbers are increasing steadily. Although the majority of nodes are in the US, the UK now has 10 nodes and more are planned. EPCC has an AG node which is fully operational and we have already participated in events using the AG.

The access group community is growing and finding new ways to make use of the technology every day. Seminars are held in virtual rooms, and secure virtual rooms are available for confidential meetings. Events vary from computer science seminars to dance events. This was demonstrated at SuperComputing 2001 where geographically separated dancers and musicians were united across the AG. This is not like the radio – interaction is involved! On an even frothier note participants can look forward to an AG symposium entitled 'Beer'. The proposers are hoping informal events such as these will put people more at ease with using the AG technology.

If you would like to use the EPCC AG node please send an email to: epcc-support@epcc.ed.ac.uk.
Further information can be found at: http://www.accessgrid.org

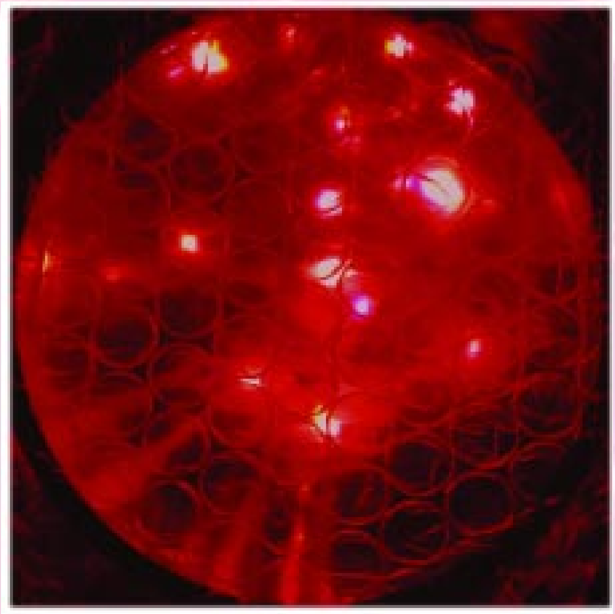

National e-Science Centre opens in Edinburgh

PROF. MALCOLM ATKINSON, NESC DIRECTOR

The National e-Science Centre was created in July 2001 to lead the development of e-Science and Grid computing for the whole of the UK.

To accomplish this ambitious plan, the e-Science Institute (eSI) was formed as a centre of education and research. Based in Edinburgh, eSI has undertaken a vigorous programme of activities: hosting meetings for UK Grid activities such as AstroGRID, Grid for Particle Physics and Biological Grid Users; tutorials on the Grid by US and UK grid experts; and serving as a forum for UK e-science directors.

We have a lively programme planned for 2002, starting with a Grid tutorial, a forum for Scottish e-science, and a Grid Portals workshop. We have joined up with IBM to host the BlueGene conference in March, and with the Global Grid Forum organisers to host GGF5/HPDC11 in July.

To host all of these e-science activities we are bringing the Victorian church that houses eSI into the 21st Century. The building has been shut down for the month of February to allow the installation of data ports, power points, extra phone lines, a central 'internet café' to accommodate visitors to eSI, and an Access Grid node (see article opposite), to allow virtual meetings with colleagues around the world. These changes will allow us to make material from our programme available for self-paced study courses, will enable web-casting of our conferences and tutorials and virtual participation in meetings via the Access Grid.

We have invited researchers for extended visits with us after our re-opening in March. We will also be entering into several commercial partnerships to stimulate the use of e-Science and the Grid.

We are excited about the programme we have planned for the coming year, but we are always open to suggestions for the eSI programme and we welcome applications from researchers wishing to visit eSI. We also welcome joint arrangements with IRCs, standards groups, networks, project groups and other users or creators of the Grid.

Contact: director@nesc.ac.uk
eSCI website: http://www.nesc.ac.uk/esi/index.html

Java and OpenMP benchmarking

MARK BULL

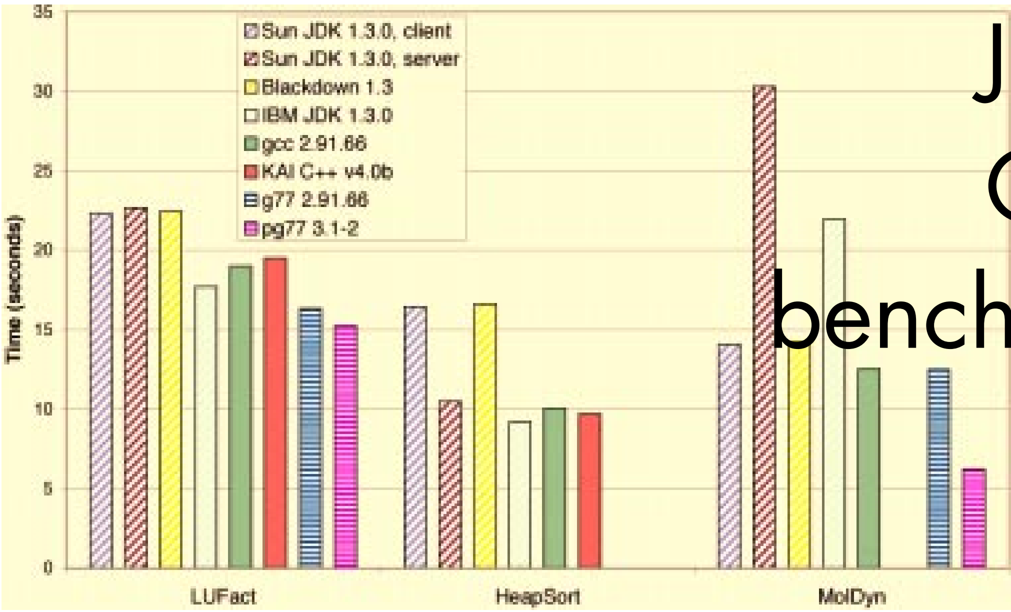

One of the principal remits of the UKHEC project is the tracking and study of new software technologies for high-end computing. Two such technologies which have been investigated in detail at EPCC are OpenMP and Java.

OpenMP research

OpenMP research at EPCC has focused on benchmarking. A set of microbenchmarks has been designed, with the goal of measuring the overheads associated with OpenMP directives on different platforms.

Overheads due to synchronisation, loop scheduling and data privatisation are an important factor in determining the performance of shared memory parallel programs. In OpenMP, the cost of these operations is dependent on their implementation in the OpenMP run-time library.

The basic technique used in the OpenMP microbenchmark suite is to compare the time taken for a section of code executed sequentially to the time taken for the same code executed in parallel enclosed in a given directive. Particular attention is paid to deriving statistically stable and reproducible results. The current suite has proved very popular both with users and OpenMP implementors. An update to the suite, reflecting changes introduced in OpenMP version 2.0, has been produced, and will be released early in 2002.

EPCC is an associate member of the OpenMP Architecture Review Board, the body responsible for the language specifications, and is taking an active role in considering possible extension to OpenMP for future versions. EPCC also has a post on the board of cOMPunity (www.compunity.org), the recently created OpenMP users and developers forum.

As part of this activity, we have been studying the performance of OpenMP codes on distributed shared memory platforms, in order to inform the debate on whether data distribution directives are a necessary addition to OpenMP.

Java research

Over the past four years, EPCC has taken a serious interest in the use of Java for high end computing. As an active participant in the Java Grande Forum (www.javagrande.org), we have led a benchmarking effort, resulting in a reasonably extensive and stable set of Java benchmarks.

Recent efforts have been focussed in two directions. First, we have provided comparison codes in C and Fortran for a significant subset of the Java benchmarks, which allow a direct and, we hope, meaningful comparison between the languages. This has produced some interesting results. For example, on Intel hardware, where significant effort has been invested in Java execution environments compared to that invested in C or Fortran compilers, Java is a strong competitor to C in performance on typical scientific and engineering code. The parallel and language comparison benchmark suites, along with associated publications can be found at: www.epcc.ed.ac.uk/javagrande.

The other direction we have been investigating is parallelism. Versions of the benchmarks have been developed for Java threads, MPJ (an MPI-like interface) and JOMP (an OpenMP-like interface). Results from these versions have shown that there are still some teething problems with parallel Java implementations, but that under the right circumstances, good scalability can be achieved.

Recent investigations have ported some larger applications (in Lattice-Boltzmann fluid simulation, and discrete event scheduling) from C or C++ to Java. These investigations have shown that the porting process itself is quite straightforward, due to the syntactic similarities between the languages, and that

GGF3: the third Global Grid Forum

STEPHEN BOOTH

The third Global Grid Forum (GGF3) was held near Rome in Frascati, Italy, from 7–10 October 2001. Despite the international situation causing a number of last minute cancellations, GGF3 had over 200 participants from over 20 countries.

The Global Grid Forum is a community-initiated forum of individual researchers and practitioners working on distributed computing, or 'Grid' technologies. GGF is the result of a merger of the Grid Forum, the eGrid European Grid Forum and the Grid community in Asia-Pacific.

The main impression that you take away from a GGF meeting is the sheer level of activity going on and the wide diversity of the Grid community.

Distributed computing is a very broad subject and can be approached from a variety of different directions. Currently the academic supercomputer centres and their users are still very strongly represented at GGF, though the forum is keen for all relevant communities to contribute.

European projects had a high profile at GGF3. The European Data-grid project was particularly prominent, although this was only one of the many distributed data handling projects that gave presentations at GGF3. Other hot topics included the accounting working group, which was very well attended, and the discussion of Web services applied to the Grid.

GGF4 will be held in Toronto, Canada, from 17–20 February 2002. GGF5 will take place in Edinburgh from 21–24 July 2002.

Further information: http://www.gridforum.org/

Java and OpenMP benchmarking *continued*

performance comparisons are both encouraging and in line with the results obtained from the benchmark suite.

OpenMP for Java: JOMP
With an interest in both these areas, a natural question is whether an OpenMP-like interface for Java would be readily implementable and useful. JOMP is a prototype system which implements such an interface. It consists of an API definition, and a proof-of-concept compiler and run-time library. The implementation is itself written in Java, making the whole system extremely portable.

Recent activities have involved integrating the compiler with the performance visualisation system Paraver (www.cepba.upc.es/paraver/), developed at UPC in Barcelona, to allow both profiling and trace visualisation of parallel Java codes.

JOMP was also successfully used to implement a parallel version of the Lattice-Boltzmann application, obtaining scalability comparable to that with C and OpenMP. Future work will enhance the interface and investigate performance portability issues.

JOMP, and associated publications can be downloaded from www.epcc.ed.ac.uk/research/jomp
UKHEC: www.ukhec.ac.uk
The current benchmarking suite is available from:
www.epcc.ed.ac.uk/research/openmpbench
www.epcc.ed.ac.uk/javagrande

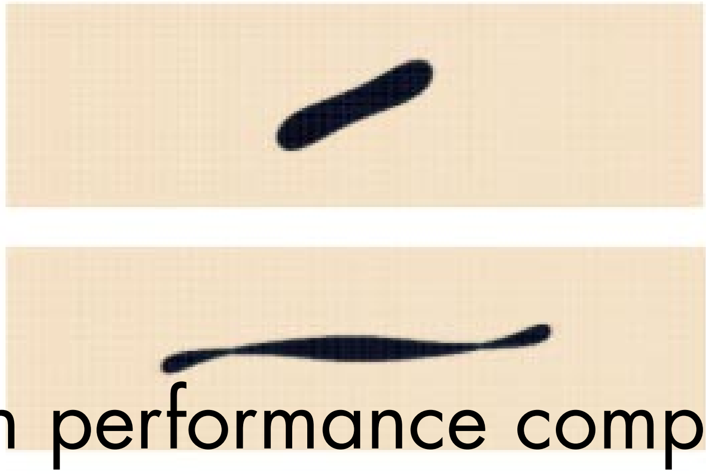

High performance computing for complex fluids

PROF. MIKE CATES, PHYSICS DEPT, UNIVERSITY OF EDINBURGH

The collaboration between EPCC and the Soft Condensed Matter Group at the University of Edinburgh is an excellent example of EPCC's long-term commitment to provide support to UK academia.

It was in 1996, as part of the EPSRC-funded High-Performance Computing Initiative (HPCI), that EPCC first undertook the development of several parallel codes for the E7 consortium 'Colloids Hydrodynamics'. This Edinburgh-led consortium was a national activity aimed at developing accurate and fast simulation codes for the flow and diffusion of colloidal suspensions. Among these codes was one based on the Dissipative Particle Dynamics (DPD) simulation technique [1], and a second one based on Lattice Boltzmann (LB) [2].

The LB code, named LUDWIG (after Ludwig Boltzmann) has since been used to carry out a series of studies of the spontaneous demixing (known as spinodal decomposition) of binary mixtures of immiscible fluids (eg see [3,4]). The performance of the LB simulation technique and scalability of our code on the Cray T3D and Hitachi SR-2201 at EPCC allowed us to investigate an unprecedented range of length and time scales for this problem. Our studies were the first in the world to attain unambiguously by simulation the 'inertial demixing' regime, in which the interface between the two fluids moves in an underdamped fashion.

When the Maxwell Institute was set up (see *HPC News 8*, July 1998), our group led one of its first pilot research programmes 'Fluid Flow in Soft and Porous Matter' which gave us the opportunity to continue our collaboration with EPCC. The functionalities of both the DPD and LB simulation codes were extended further to study the hydrodynamic and rheological properties of multicomponent fluids. For instance a novel predictor-corrector scheme for the thermodynamically consistent simulation of the wetting phenomenon was developed by P. Bladon and parallelised [2]. The values obtained for the contact angle for some simple tests in 2D agree with the predictions of the model used [2]. The work is currently being extended from wetting of flat walls to that of curved solid surfaces. This functionality is particularly important as it will enable thermodynamically accurate simulation of the flow of multicomponent fluids in solid pores, and also of colloidal particles suspended in binary fluids.

Since 1998 the collaboration has continued as part of the EPSRC-funded project 'HPC for Complex Fluids' (grant GR/M56234, with collaborators from Queen Mary University of London, Unilever and Oxford University). This phase of the work has seen the development of Lees-Edwards (sliding periodic) boundary conditions to allow ongoing investigations of breakup of fluid droplets under shear, and of the effect of shear on spontaneous demixing (cf.[4]). Meanwhile our definitive analysis of the inertial demixing regime was completed [5]. The implementation of sliding periodic boundary conditions is less straightforward for LB than it is for off-lattice algorithms such as DPD; group member A. Wagner and EPCC TRACS visitor I. Pagonabarraga devised a new scheme described in [6], which has now been fully implemented and parallelised in LUDWIG.

EPCC's expertise in the field of scientific visualisation has also proved a cornerstone of this collaboration. Visualisation is now a key tool for obtaining direct scientific insight into the dynamics of complex fluid flows, transcending its earlier, more illustrative role. In the case of binary fluid demixing, it has allowed us to gain direct evidence of several complex physical processes (such as recoil after fluid-fluid pinch and transient capillary waves) which could not be extracted from the simulation data by any other means. This recent animation study has resulted in a milestone publication in the web-based refereed journal, the *New Journal of Physics* (www.njp.org) [7]. Apart from its scientific

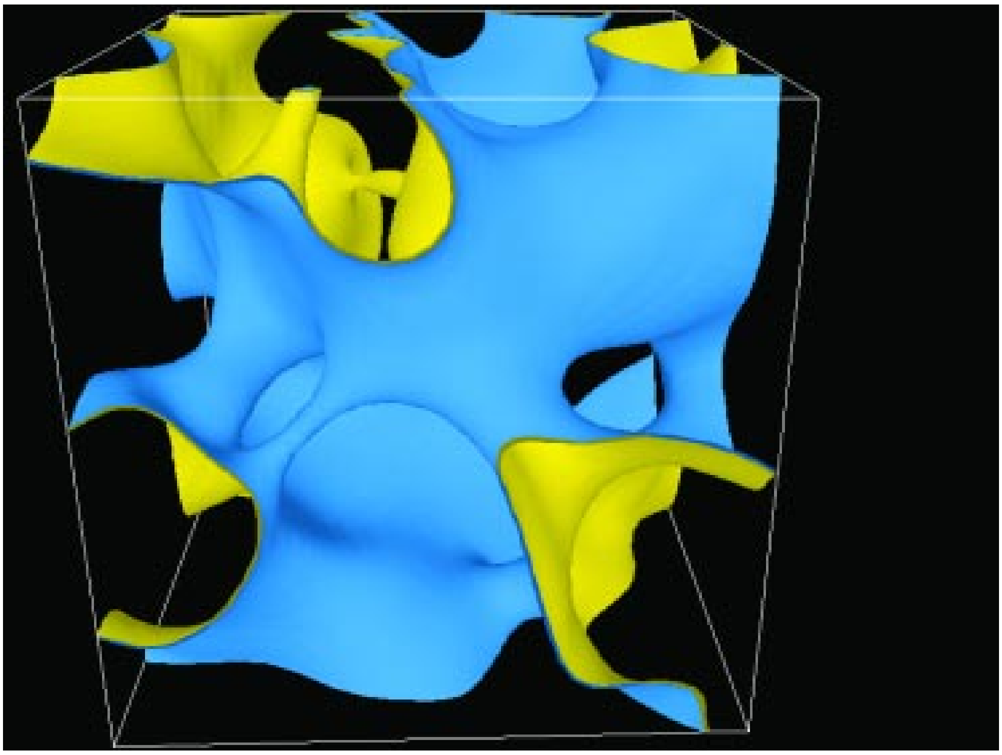

Right. The interface between two fluids undergoing spinodal decomposition in the inertial demixing regime.

Opposite. Droplet breakup under shear in two dimensions studied using lattice Boltzmann. Depending on Reynolds and capillary numbers (controlled by shear rate, droplet size, fluid viscosity and surface tension) the droplet shape just before breakup can be either of those shown. The transition between these different breakup modes is currently being investigated, as is breakup in three dimensions.

role, the importance of on-line visualisation for code validation, debugging and parameter steering is well documented. LUDWIG has built-in support for graphical output, both for its serial and parallel implementations. The functionalities offered by EPCC's DIVE library add stream-based visualisation capabilities within LUDWIG, thus allowing the visualisation on local workstations in Edinburgh of parallel runs carried out on the national facilities at CSAR. Full computational steering capabilities will soon be developed and tested within a Grid context as part of the EPSRC-funded Grid testbed project, RealityGrid [8]. This project, led by Prof. P. Coveney (Queen Mary University London), gathers researchers from six major UK universities and aims to Grid-enable the realistic modelling of complex condensed matter structures at the meso and nano-scale levels. Our research group, which will play a major part in the mesoscale modelling and simulations in RealityGrid, expects continuing major benefits from the presence of EPCC within this consortium. The presence in Edinburgh of NeSC (the National e-Science Centre, and successor to the Maxwell Institute) will strengthen these links still further.

A final area of interest, albeit more on the computational side, is the port of LUDWIG to Java undertaken as part of EPCC's Summer Scholarship project. A cut-down version of LUDWIG was successfully ported to Java and parallelised using JOMP (OpenMP for Java, developed by Mark Bull at EPCC [9]). Surprisingly good results were obtained with the Java/OpenMP matching the scalability of the C/OpenMP implementation and achieving between 50% and 80% of the performance of the native C implementation depending on the system size. These results have been presented at 'Últimos avances en Informática – 2001' and will be published shortly in the open literature. (See the forthcoming issue of the UKHEC newsletter.)

Throughout the years of our collaboration with EPCC, the input of EPCC staff has not been limited to code development, but has also involved genuine and sustained engagement with the scientific goals and objectives of the group. This has allowed a far more creative dialogue to take place between software development and science than would have been possible within a short-term, contract-by-contract interaction. Such dialogue has been of immeasurable benefit to our research work. Long may it continue.

Contact: M.E.Cates@ed.ac.uk.
See also: www.ph.ed.ac.uk/cmatter/soft.html

References

[1] *S. I. Jury, P. Bladon, S. Krishna and M. E. Cates, Phys. Rev. E. 59 (1999) R2535.*
[2] *J.-C. Desplat, I. Pagonabarraga, and P. Bladon, Comp. Phys. Comm. 134 (2001) 273.*
[3] *V.M. Kendon, J.-C. Desplat, P. Bladon, and M.E. Cates, Phys. Rev. Lett. 83 (1999) 576.*
[4] *M. E. Cates, V. M. Kendon, P. Bladon and J.-C. Desplat, Faraday Disc. R. Chem. Soc. 112 (1999) 1.*
[5] *V.M. Kendon, M.E. Cates, I. Pagonabarraga, J.-C. Desplat, and P. Bladon, J. Fluid Mech. 440 (2001) 147.*
[5] *A. Wagner, and I. Pagonabarraga, to appear in J. Stat. Phys. (cond-mat/0103218).*
[6] *I. Pagonabarraga, J.-C. Desplat, A. Wagner, and M.E. Cates, New J. Phys. 3 (2001) 9.*
[7] *EPCC's DIVE library: www.epcc.ed.ac.uk/industry/DIVE*
[8] *EPSRC Grid testbed project, the "Reality Grid". www.realitygrid.org*
[9] *OpenMP for Java (JOMP). www.epcc.ed.ac.uk/research/jomp*

ENACTS: the HPC technology roadmap

JAN FAGERSTRÖM, TORGNY FAXÉN AND JC DESPLAT

ENACTS is a Co-operation Network which brings together many of the key High Performance Computing (HPC) centres and key user groups from around Europe. It provides a co-operative structure for reviewing the impact of Grid computing technologies, enabling the network to formulate a strategy for increasing the quantity and quality of access. (See *EPCC News* issue 42 for more information).

The topics covered in 2001 include:
- Grid service requirements
- HPC technology roadmap
- Grid enabling technologies.

The HPC Technology Roadmap has been undertaken by the National Supercomputer Center (NSC) in Linköping, Sweden in collaboration with the Center for High Performance Computing in Molecular Sciences (CSCISM) in Perugia, Italy. The objective of the study is to determine the likely technological and economic trends which will prescribe the hardware architectures of HPC systems over the next 5 to 10 years, and to evaluate the effects that this will have on applications software.

The study has been divided into two parts. First, NSC has provided a survey of the technology roadmap for processors, memory, networking, data storage, custom-built solutions, and software paradigms and standards. Second, CSCISM has provided a case study focusing on the usefulness and implications of the technologies discussed in the technology roadmap, for the key molecular science community. The complete study will soon be presented in an ENACTS report.

The technology roadmap survey was conducted via interviews with several major HPC vendors. Vendors are often portrayed as having quite different views on supercomputing, especially in areas such as processors, vectorisation, memory architecture and programming models. Interestingly enough, the interviews showed that there is general agreement on many of the major trends that will be important for future computer architecture in general and HPC computing in particular. Most future computer systems can very generally be described as scalable parallel architectures based on the notion of clustered SMPs. When looking 'under the bonnet' however, there is actually a wide array of different solutions presented. New advances in all relevant computer technologies, made possible by an increasing consumer market as well as the open source movement, seems to make the number of available solutions even larger than before! Actually, if anyone believes that future HPC systems will all move towards a homogeneous computer architecture, it seems likely that they will be disappointed!

The second part of the study provides an assessment of the implications of future computer hardware and architecture for the key molecular science community. The assessment was made by several experts on molecular applications at CSCISM. Using the results from the vendor interviews in part one, two typical future HPC solutions were configured. These were used when investigating the impact of future computer HPC development on three main approaches to the definition of interatomic potential:

- Model potential based on Force-Field (FF) parameterizations.
- DFT (Density Functional Theory) derived potentials.
- Ab-initio potentials.

The final report from the HPC Technology Roadmap study is expected to be available in March 2002, and will be downloadable from the website.

Up-to-date information about the HPC Technology Roadmap study within the ENACTS project is published at the ENACTS HPC Technology Roadmap web site: http://www.nsc.liu.se/rd/enacts.

The Grid Service Requirements report is already available for download from the ENACTS web site: http://www.enacts.org

ENACTS is funded as part of the EC's `Improving Human Potential, Access to Research Infrastructures' programme.

Contacts: Alison Kennedy (EPCC), a.kennedy@epcc.ed.ac.uk and Jan Fagerström (NSC), fager@nsc.liu.se

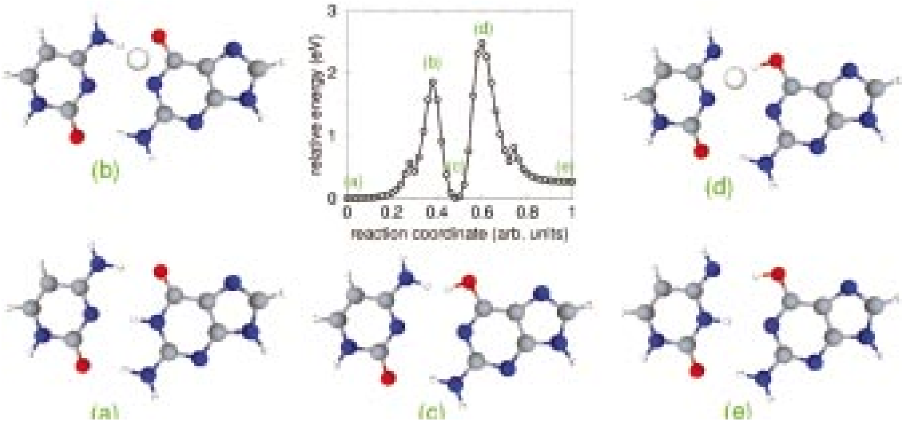

TRACS: furthering European research collaboration

CATHERINE INGLIS

Since 1993, TRACS (Training and Research on Advanced Computing Systems) has been helping to stimulate collaboration between researchers from all over Europe and those based here in Edinburgh's universities, and recently also in universities in Glasgow, St Andrews and Dundee. Through TRACS, researchers from any discipline to which HPC is applicable can come to Scotland for an EC-funded visit of up to three months, working with the department of their choice while also having access to EPCC's facilities and consultancy.

HPC is applicable to many different fields, including astrophysics, chemistry, meteorology and geophysics. Here is a report by a TRACS visitor from outside the sphere of pure informatics.

Parallelisation of a General Nudged Elastic Band Method

PAUL MARAGAKIS, IESL-FORTH/DEPARTMENT OF PHYSICS, UNIVERSITY OF CRETE
HOST: DR GRAEME ACKLAND, DEPARTMENT OF PHYSICS AND ASTRONOMY, UNIVERSITY OF EDINBURGH

An important problem in condensed matter physics and in theoretical chemistry is the identification of a lowest energy path for a rearrangement of a group of atoms from one stable configuration to another. This allows the calculation of rates of chemical reactions or of diffusion events. One of the most promising methods for obtaining the 'minimum energy path' is the Nudged Elastic Band method [1], which guarantees to find a homogeneously sampled trajectory connecting the stable configurations.

We have implemented a general version of this method[2] so that it can use any existing electronic structure code for the calculation of the total energy and forces. During the visit to EPCC the code has been parallelised and two parallelisation schemes have been evaluated. In addition the project has been ported to the EPCC Linux Beowulf cluster, and two electronic structure codes (one ab-initio and one semi-empirical) have been interfaced to it. Due to a combination of file-system caching and parallelisation we obtained a speedup of almost two orders of magnitude when using the semi-empirical electronic structure code on the 16 processor cluster compared to the serial version on the original SGI development machine.

The figure above shows the tautomerisation of a cytosine-guanine pair [3], two nucleic acids that are used for information coding in DNA. The initial state (a) shows the 'amine' cytosine form binding with the 'keto' guanine form which is a stable pair appearing in normal DNA. The final state (b) shows 'imine' cytosine binding with 'enol' guanine. This tautomerisation process is extremely significant because it changes the bonding properties of the molecule and may lead to mispairing during duplication and thus mutations. In the figure we show the intra-molecular proton transfer a reaction path-way that reveals a very interesting double peak structure with two saddle points and an intermediate stable configuration.

REFERENCES

[1] *H. Jonsson, G. Mills, and K. W. Jacobsen, in Classical and Quantum Dynamics in Condensed Phase Simulations, edited by B. J. Berne, G. Ciccoti, and D. F. Coker (World Scientific, Singapore, 1998).*
[2] *P. Maragakis and E. Kaxiras, General nudged elastic band method (2000), in preparation.*
[3] *P. Maragakis, D. Margetis, E. Kaxiras, M. Elstner, and T. Frauenheim, Transition paths between tautomeric structures of nucleic acid bases and pairs (2000), in preparation.*

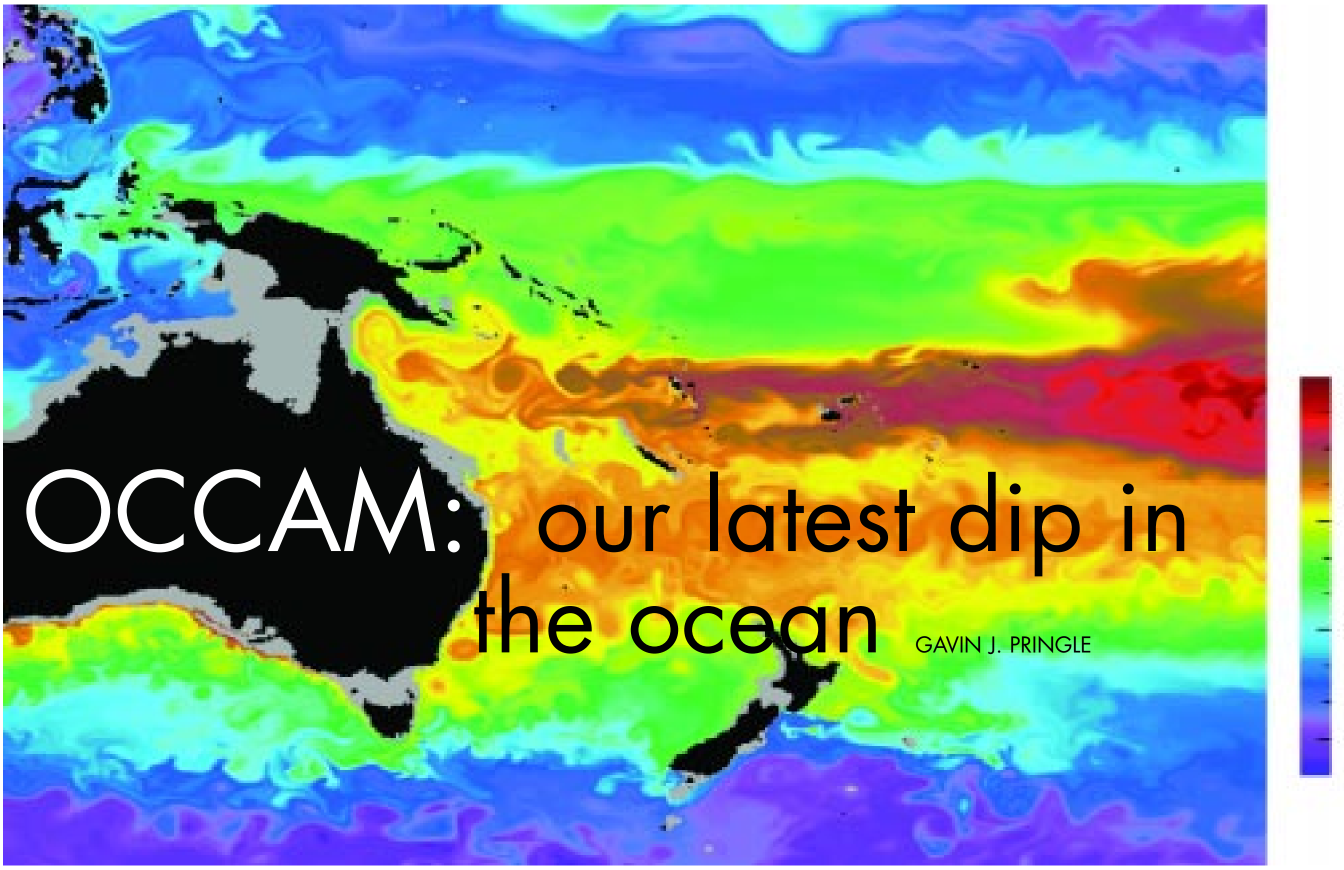

OCCAM: our latest dip in the ocean

GAVIN J. PRINGLE

EPCC and the OCCAM consortium worked closely, once again, to produce a seamless method to assimilate real world data into a simulation of the world's oceans.

The simulation of Earth's climate is highly complex and demands vast amounts of computing power. One very important component of this endeavour is the accurate simulation of the world's oceans.

A very wide range of timescales are required as the deep ocean takes many decades to respond to changes at the surface, yet fast surface gravity waves will restrict numerical timesteps to a few hours. Length scales need to be large enough to capture the whole globe, but small enough to capture important eddy fields: the Gulf Stream is, in places, only 50km wide.

The Ocean Circulation and Climate Advanced Modelling (OCCAM) project was started in 1992 with the aim of utilising the Cray T3D at EPCC to develop and run a parallel, high-resolution global ocean model. The model now runs in production on the Cray T3E at CSAR (Computer Services for Academic Research) and is currently being ported to the Origin 3000. EPCC has some previous experience of the OCCAM code from performance optimisation work done under the HPCI initiative.

The OCCAM code is fully 3-dimensional, with ¼-degree latitude-longitude by 36 depth levels resolution. The vertical layers vary in thickness from 20m at the surface to 250m at depth. The model has four main variables, namely the meridional and zonal velocity components, the potential temperature and the salinity, plus several two-dimensional fields such as the depth-averaged velocity and the free-surface height.

The poles present a problem numerically as they create potential singularities, yet the North Pole is important in ocean modelling. OCCAM circumvents any numerical instabilities by employing two grids: 'One standard longitude-latitude grid covers the entire region from the Antarctic continent to the Bering Straits. Any sea points north of the equator in the North Atlantic are masked off in this grid and treated as land points. A second grid is then overlaid and used to cover the North Atlantic, Mediterranean and Arctic basins. This second grid is rotated through 90 degrees so that the grid poles lie on the geographical equator.'[1]

The code was designed to run on distributed memory parallel computers and is parallelised following an irregular domain decomposition paradigm. Each processor updates values within the volume of the mesh that contains ocean – no work is performed over land points. For the sake of load balancing, the ocean points are distributed so that each PE has almost the same number. During each time-step, each processor exchanges boundary information with its neighbouring processors to update shared 'halo' points. The communications use OCCAM's own communication harness. This, in turn, calls either Cray's own communication language, SHMEM, or either of the portable MPI and PVM libraries.

The focus of this project was data assimilation. At regular intervals, the OCCAM model is compared to real experimental data and appropriate changes are made to the simulation. This is done for two observables: the sea-surface height taken from satellite altimeter data, and the temperature at various depths

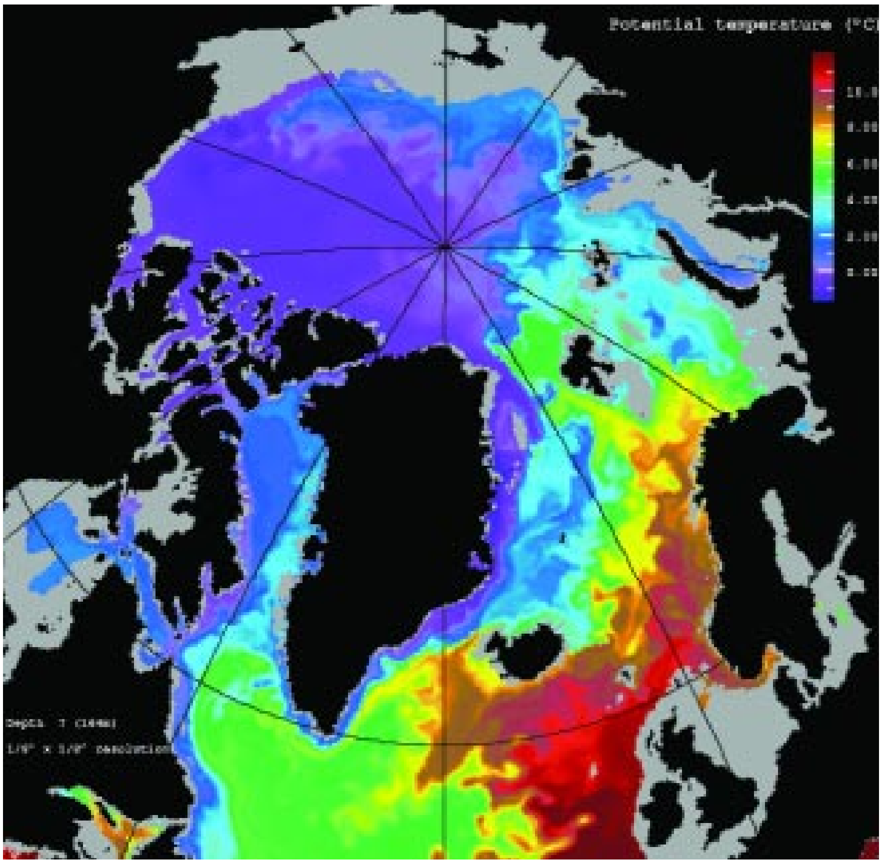

Opposite. The OCCAM model has shown the existence of a complex series of jets and boundary currents associated with island groups such as Fiji and New Caledonia in the South Pacific.

Right. The model has reproduced the Arctic current system in tremendous detail.

recorded by instruments dropped overboard from ships or arrays fixed to buoys. The 2D sea-surface data set is small enough that the assimilation can be performed by a single T3E processor, and this is done in-line as part of the simulation. For the 3D temperature data, however, the assimilation was performed off-line by a separate program which operates on checkpoint data files. This involved stopping and restarting the main simulation code and performing the assimilation could easily take a whole day.

There are around 10,000 discrete sample points, each containing temperatures interpolated to the 36 depth levels. The complete raw temperature data set can easily be stored on each slave. However, what is actually required are the differences, or 'anomalies', between this data and the simulation temperatures at the corresponding grid points.

Each processor computes the anomalies for those sample points that lie within its local 'core' grid points. However, the process of assimilation requires knowledge of the anomalies in an area that extends 3 degrees (many grid spacings) around each grid point. The anomalies of shared 'halo' grid points must therefore also be calculated. It was decided that, rather than perform the assimilation on the core points only and then update the halo points by performing a boundary update, we would update the halo points at the same time as the core points. This avoided having to alter the boundary update code. It is also likely that the employed approach is more efficient than invoking the communication routines. Once the anomalies have been generated, the temperature data on each slave can be updated without any further communications.

This project has now been successfully completed and the process of assimilation is now performed *within* the OCCAM code. Before, when the work was done off-line, it could take around a day to perform the assimilation and restart the code, although the assimilation itself took around 6 minutes. Within the OCCAM code, the assimilation now takes around 9 seconds, due to optimisation and parallelisation, and since the assimilations only occur about once every simulation month, the overhead of in-line assimilation is negligible. Hence, this work has increased the scientific productivity of the OCCAM model.

For far greater insight into the workings of the OCCAM code, please point your browser to: ww.soc.soton.ac.uk/JRD/OCCAM/welcome.html

References

[1] 'Ocean Modelling', Coward, A., HPC News, EPCC, 3, April, 1996. http://www.epcc.ed.ac.uk/t3d/documents/HPCNews/issue_3/pages5_6.html

[2] 'OCCAM on the T3D: A New View of the Global Ocean', Webb, D., HPC News, EPCC, 10, June, 1999. http://www.epcc.ed.ac.uk/t3d/documents/HPCNews/HPCNews10.pdf

Supercomputing 2001

10-16 November, Denver, Colorado
DAVID HENTY

Supercomputing is the foremost international conference on high-performance networking and computing. Last year EPCC was more involved than ever before. As well as having a research booth in the main exhibition hall and presenting research work as part of the Technical Program, we also ran a tutorial session for the first time.

It was clear that the surge of interest in the Grid has had a major effect. The majority of stands in the exhibition hall mentioned the Grid: researchers were keen to showcase their Grid projects, software vendors described how their products would operate in a Grid context and hardware vendors promoted their own technologies as the key Grid platforms.

The impact on the technical program was less dramatic – about a quarter of the technical sessions had 'Grid' in the title. However, the focus was still very much on the Computational Grid with little in the important area of Data Grids. It will be interesting to see if the focus changes in the coming year.

I was pleased to see that SC2001 still had space for old-fashioned parallel supercomputers! Sun, IBM, Compaq and Intel were all demonstrating new hardware, and Moore's Law shows little sign of tailing off.

Overall, SC2001 was an extremely enjoyable event and I very much hope that EPCC can maintain the same level of participation at the next meeting in Baltimore.

For more information see: www.sc2001.org

SC2001 tutorial session: 'Java for HPC: Performance and Parallelisation'

Java offers a number of benefits as a language for HPC, especially in the context of the Computational Grid. For example, it offers a high level of platform independence not observed with traditional HPC languages and its object-oriented nature facilitates code re-use and reduces development time.

The aim of our one-day interactive tutorial was to focus on Java performance issues relevant to HPC applications, based on experiences gained through EPCC's benchmarking activities for the Java Grande Forum (see 'Java and OpenMP benchmarking' on p4). Thanks to a loan of equipment by Sun Microsystems, we had access to some 20 Sun Ray terminals attached to a powerful server. This allowed us to follow the lectures with tutored practical sessions illustrating the key points.

The tutorial was a great success, with every terminal occupied during the busiest sessions. The SC2001 committee collected feedback from attendees of all 27 tutorials, and we received the second highest overall score.

Technical Program
Continuing the theme of Java for HPC, Mark Bull presented a paper entitled 'A Parallel Java Benchmark Suite' (available at http://www.sc2001.org/papers/pap.pap158.pdf). We also contributed to a workshop in the SC Global session, broadcast worldwide using Access Grid technology. Sitting in the US we were able to see Rob Baxter of EPCC give a talk live from the UK. A number of Access Grid nodes were linked together and Rob took questions from Manchester, Denver and two national labs elsewhere in the US.

Research Booth
An innovation this year was to group all the European booths into a single 'European Village' to provide a focus for all the HPC activities on this side of the Atlantic.

EPCC had a larger booth than last year (see photo above) which allowed us to have eight Sun Ray terminals showing a range of interactive demonstrations and presentations. These highlighted a number of projects in network computing and the Grid including Intersim-DiffServ, Java Grande, JOMP, JiniGrid and ePortal. We also showed a number of visualisations of large-scale computer simulations produced by the UK Virgo consortium and EPCC's TRACS visitors.

There was a lot of interest in the booth, particularly in the demonstrations, and it was definitely worth the effort of setting up the computing equipment on the stand. Our backdrop of the Forth Rail Bridge at night was also very popular, attracting interest from a lecturer who asked to have it for his own classroom and from a delegate who planned to take a vacation to visit all the world's great bridges!

For more information on the projects mentioned above please see: http://www.epcc.ed.ac.uk/grid/

Intersim: A hierarchical tool for parallel simulation of internetwork traffic

ROB BAXTER

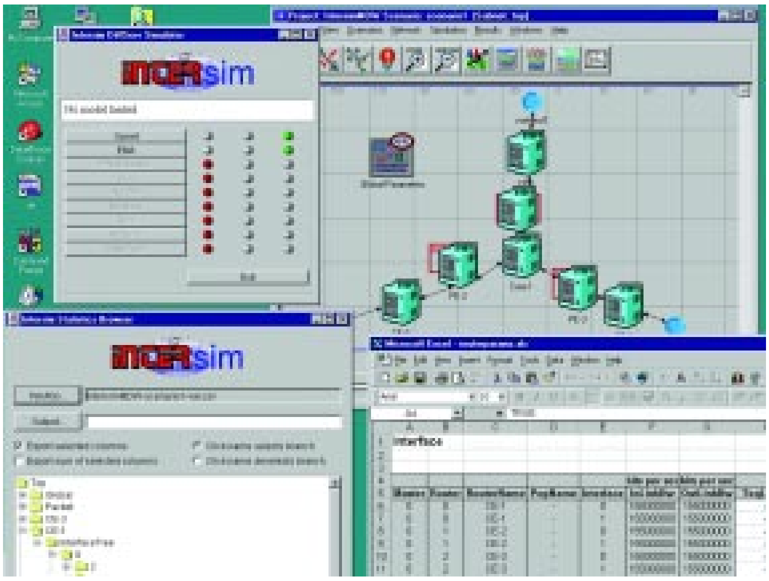

EPCC has completed the world's first packet-level simulation tool for differentiated service networks. Intersim is a tool that simulates Cisco equipment for IP quality of service on diffserv networks, to help designers build better networks.

Intersim has been built at EPCC over the last two years and includes some 45,000 lines of C++, Java and Visual Basic code. It is a discrete event simulator built around the Opnet network simulation tool and models the movement of individual IP packets through diffserv networks.

Quality of service

IP quality of service is an important new technology that will significantly improve the performance of the next generation of internet networks. In a computer network, just like a road network, traffic accumulates at junctions. If too much traffic arrives at a queue – a network 'junction' – the excess data packets must be discarded, or 'dropped'. In today's IP networks like the Internet, there is no way to discriminate between different types of packets. There is only one type of queue – all the packets go into that queue and have an equal chance of being delayed or dropped.

IP quality of service allows us to add more queues, like multiple lanes on a highway. Traffic that is sensitive to delay, such as voice or video, can be put into high priority queues (fast lanes), bulk data transfers, on the other hand, can be put into low priority queues (slow lanes). Different types of network service are thus differentiated, hence 'diffserv'. In the diffserv model, traffic is marked into different classes when it enters the network and is assigned into different queues at each junction, based on its class. A business customer leasing network capacity from a telecoms company may be willing to pay a premium to guarantee a certain bandwidth for some quality-critical types of traffic, while regular Web access over HTTP might be relegated to a best effort class.

A leading-edge tool for leading-edge network design

Quality of service for network traffic is a big area of research, both theoretically and in practical, engineering terms. Network providers and network equipment manufacturers need to understand in detail how to build diffserv networks, how best to use their existing network capacity to serve a broad customer base and meet a whole set of service level agreements.

Currently, designing these new networks is hard. The benefits of quality of service only come when service levels for different traffic types can be guaranteed. The complexity of the network dynamics makes meeting these guarantees a significant challenge.

To address these complexities, Intersim has been constructed in close collaboration with Cisco Systems. Intersim's network models are based on Cisco's source code and validated against the real hardware to ensure near packet-for-packet performance. Its aim is to support diffserv network design on two levels: for strategic decisions such as capacity planning and SLA targets; and for tactical decisions such as parameter tuning within router groups. Here are some examples:

- At present network providers have to grow traffic (ie customers) cautiously after deployment, because they don't know the network limits. Intersim can find these operating limits, allowing the possibility of rapidly growing the traffic to capacity without risking existing SLAs.

- Once deployed, it is difficult to tune the configurations of a live network. Intersim can test any configuration, allowing the possibility of safely experimenting with major changes.

- It is fundamentally difficult to determine the impact of new applications. Intersim can be used to test new application usage and give users a perspective on performance.

- Time to market is a crucial competitive edge for network providers. Intersim can reduce design time by allowing designers to validate designs more quickly and easily than testbeds and on a larger scale.

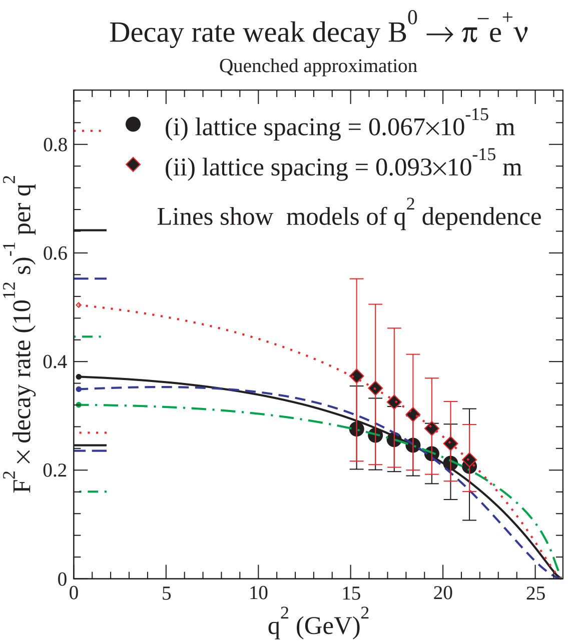

QCD: quarks and gluons on a lattice

CHRIS MAYNARD, UKQCD COLLABORATION

The UKQCD Collaboration aims to increase the predictive power of the Standard Model of elementary particle interactions through numerical simulation of Quantum Chromodynamics (QCD). EPCC provides the technical expertise to enable the UKQCD group to exploit the full range of High Performance Computing facilities available to them.

Look around you, there are many things; desks, computers, other people, planets, stars and galaxies. They are all made up of protons and neutrons (matter). It is believed that while matter and anti-matter were created in equal amounts in the Big Bang at the start of the Universe, there is only matter around today. This is fortunate for us, as matter and anti-matter annihilate one another to leave energy, and so all the stable structures in the Universe, from galaxies to people, are only possible due to this matter/anti-matter asymmetry.

In the Standard Model of particle physics, matter/anti-matter asymmetry occurs in the weak nuclear interaction of quarks. Physicists at accelerator complexes are trying to measure and understand these weak interactions. However, quarks are are always confined in composite particles called hadrons, eg protons and neutrons, by the strong nuclear interaction. Understanding the theory of the strong interaction (quantum chromodynamics or QCD) is necessary to disentangle the weak interaction from it. But this may not be enough to explain the cosmic asymmetry and new physics may be required.

QCD describes particles that have one of three colour charges, analogous to electric charge in electromagnetism (EM). Two quarks interact by exchanging a gluon, which is analogous to the photon in EM. Gluons carry colour charge which allows them to couple to each other, unlike the photon which is electrically neutral. It is this difference that makes QCD hard to solve. The fundamental degrees of freedom are the quarks and gluons, and they are bound by the strength of the interaction into hadrons. It is the properties of the hadrons that we need to calculate.

By replacing spacetime with a finite grid or lattice, the calculation can be done numerically. The box size must be large enough so that the hadrons fit in the box, and the lattice spacing small enough to resolve the structure. The first stage is to simulate the vacuum, which is a sea of virtual quark/anti-quark pairs and gluons. A large number of random vacuum configurations is generated. Next, the quarks required to make a hadron are introduced onto the lattice and averaged over all the configurations.

Until recently, most calculations were done in the quenched approximation. This ignores the virtual quark pairs as it takes at least one thousand times more computation to generate a configuration with quark pairs than without. Most of the effort is used in computing a large matrix and its inverse on each configuration. This matrix is typically more than 10^6x10^6 in size. The lighter the mass of the virtual quarks in the sea, the more computation required. The problem is mostly local, which readily lends itself to parallel computation, each processor computing a small part of the matrix and communicating with its nearest neighbours. The UKQCD collaboration [1] has used the T3E at EPCC to generate configurations in full QCD, but with masses around that of the strange quark. With the new QCDOC machine we hope to start pushing the mass down to nearer zero quark mass, the up and down quark masses which contribute most to the vacuum.

Shown in the figure is a quenched lattice determination of the rate at which a B meson decaying by the weak interaction to a pion (meson) and an electron and its neutrino. q2 is the momentum transferred to the electron and neutrino. Once the experimental rate for this process has been measured, the ratio of experiment to theory is F2 which is related to the amount of matter/anti-matter asymmetry in the SM. The new QCDOC machine will enable us to do a similar calculation without the quenched approximation.

[1] *UKQCD collaboration: http://www.ph.ed.ac.uk/ukqcd/*

Machine vision – spreading the word

DIANA ENGESSER

EPCC, working with four other European technology transfer centres, has launched EXPANSIV – a new programme designed to publicise the benefits of integrated machine vision (IMV).

The benefits of integrated machine vision are recognised by the programme's participants but now these have to be communicated to a wider market. EXPANSIV will achieve this through a series of promotional actions over a two-year period. It will work to raise awareness among potential new end-users and exploiters of the benefits of IMV. It will also support the broad deployment of IMV and best practice throughout European industry.

EXPANSIV activities will include a widespread media campaign addressing the target market as well as attending relevant trade fairs and events to promote the new technology solutions. Local seminars and workshops aimed at business are also to be held.

EXPANSIV, which is funded by the European Commission, will be an accompanying measure to the activities of the existing EUTIST-IMV group of projects. The project is funded by the EC's Information Society Technologies (IST) and will run until the end of 2003.

For more information, contact: EXPANSIV@epcc.ed.ac.uk

An IMV installation at a sawmill.
Pic courtesy of inX Systems.

EPCC training and education course timetable

February–May 2002

Here are the forthcoming courses that will be held at EPCC. They are open to all and there is a limited number of free places for academics.

If you would like to register for a course, please do so as early as possible to ensure a place.

For further information and commercial rates, contact: EPCC-Support@ed.ac.uk

Updates to this timetable can be found at: www.epcc.ed.ac.uk/epcc-tec

February 2002	
26-28	Applied Numerical Algorithms
April 2002	
16-18	Fundamental Concepts of HPC
22	Performance Optimisation
23-25	Message Passing Programming
30-2 May	Scientific Visualisation
May 2002	
7-9	Practical Software Development
14-16	Shared Memory Programming

|epcc|

MSc in High Performance Computing

EPCC, a technology transfer centre within the University of Edinburgh, offers a one-year Masters course in High Performance Computing (HPC). A number of EPSRC studentships are available, which cover the full fees for EU residents. UK residents also qualify for a maintenance grant.

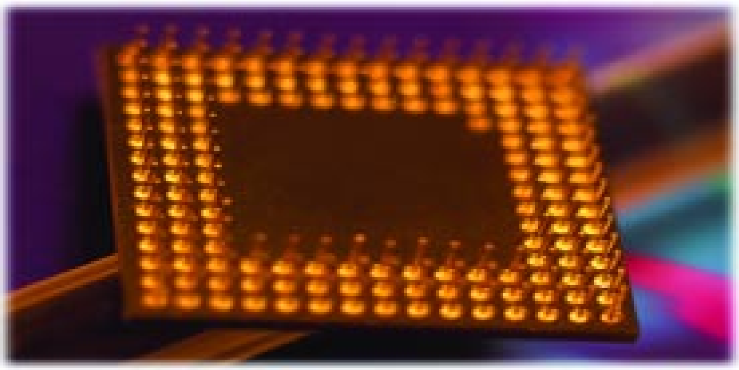

EPCC has an international reputation in the application of novel computing solutions to real-life problems. This postgraduate qualification, awarded by the University of Edinburgh, has a strong practical focus. It covers topics relevant to a wide spectrum of careers including computational science research and commercial software development.

MSc students will have access to an impressive range of leading-edge parallel platforms and HPC technologies. Graduates of this course will hold one of the few university-accredited postgraduate HPC qualifications in Europe.

The entrance requirement is a good honours degree or equivalent work experience. No prior HPC knowledge is assumed, but candidates must be competent in Java, C++, C or Fortran.

The MSc runs annually, starting each October, and can be studied full or part-time. The taught part comprises a series of intensive modules, each lasting around three days, with associated tutorials, course work and examinations. Students also submit a thesis based on a 16-week practical research project.

Taught modules:

- Fundamental Concepts of HPC
- Practical Software Development
- Message Passing Programming
- Shared Memory Programming
- Specialised Programming Models
- Object Oriented Programming for HPC
- Applied Computer Science
- Exploiting the Computational Grid
- Applied Numerical Algorithms
- Performance Optimisation
- Scientific Visualisation

Applications are encouraged from graduates of all areas of science, engineering, computer science and mathematics, and from those currently working in a relevant field.

For more information and application details see http://www.epcc.ed.ac.uk/msc/
Email: msc@epcc.ed.ac.uk

This MSc is supported by: EPSRC

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