Capability Computing

The newsletter of the HPCx community

[ISSUE 12, AUTUMN 2008]

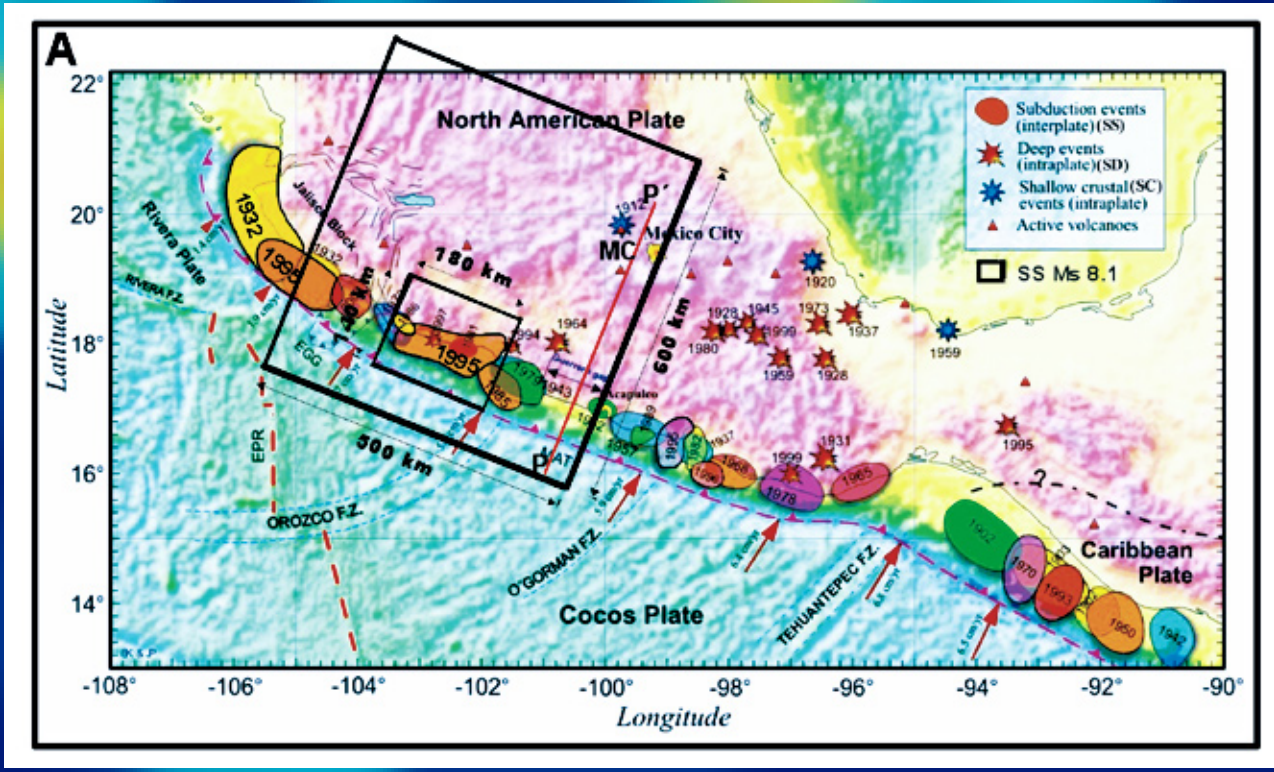

Earthquake simulations: ground-breaking science on HPCx

CONTENTS

2 Editorial
Cray X2 Vector goes online

3 Two-phase jet simulations on HPCx

4 HPCx: reaching out to life scientists

7 Forthcoming events

8 Helicopter fuselage flow simulations

11 ParaView – parallel visualisation on HPCx

12 3D parallel modelling of large subduction earthquakes

14 PRACE surveys the current European HPC ecosystem

15 UK scientists to get access to DEISA resources

16 HPCx Annual Seminar
Petaflop computing unveiled at ISC2008

Editorial

Liz Sim, EPCC, University of Edinburgh

Welcome to the latest edition of *Capability Computing*, the newsletter of the HPCx national computing service. This edition comes at the turning point between the end of the original HPCx service contract and the start of the extension until January 2010. It's an exciting time for all. By the end of Q4 2008 the HPCx hardware upgrade to the IBM System p5 should be complete. We also wait with baited breath for the outcome of EPSRC's High Performance Computing Complementary Capability Challenge call, which brings the prospect of yet more diverse science and novel computing on HPCx.

One of the key aims of *Capability Computing* is to share the science that the HPCx service enables, and this edition is no exception to that rule. We are delighted to share more success stories with you from other users, including the use of HPCx in gas-liquid two-phase jet simulations.

We also take a look at four different computational strategies used in the simulation of helicopter fuselages in cruise conditions. Our feature article looks at the use of dual-core platforms in complex 3D modelling of large subduction earthquakes.

The key focus for HPCx for the remainder of the service is Complementarity. In this edition our terascaling team introduce the use of ParaView as a means of producing high quality animations. Results of lattice Boltzmann simulations of fluid mixtures under shear are used as an example of what can be achieved. We also take a look at current outreach activities in the field of Life Sciences. In an article from our outreach team we examine some of the ongoing Life Science projects and the challenges they have set themselves on HPCx, such as developing parallel code equivalents to traditional serial Life Sciences codes.

Looking beyond our local shores, we have a status update from DEISA (Distributed European Infrastructure for Supercomputing). We also include a review of recent PRACE (Partnership for Advanced Computing in Europe) activity with regard to identifying HPC applications for benchmarking on Petaflop/s systems.

Both EPCC and STFC Daresbury Laboratory will be present at SC'08 and I look forward to seeing you there. I hope you enjoy this edition of *Capability Computing*!

HECToR X2 Vector goes online

Liz Sim, EPCC, University of Edinburgh and Steve Jordan, Cray Systems

The HECToR Vector X2 component successfully passed acceptance back in July. The integration of the X2 with the HECToR XT4 went extremely smoothly thanks to the efforts of all involved. The resulting 'XT5h' hybrid supercomputer is the first X2/XT4 integration on this scale.

The Cray vector system – known as Black Widow – consists of 28 vector compute nodes, each of which has 4 Cray vector processors, making 112 processors in all. Each processor is capable of 25.6 Gflops, giving a theoretical peak performance of 2.87 Tflops. Each 4-processor node shares 32Gb of memory.

The Black Widow interconnection network has a point-to-point bandwidth of 16 Gb/s and a bi-section bandwidth of 254 Gb/s. The average ping-pong MPI latency is ~ 4.6 µsec.

The X2 systems were shipped from Cray's factory in Chippewa Falls, Wisconsin on Friday 13th June and arrived on site for installation at the University of Edinburgh's Advanced Computing Facility (ACF) on Thursday 19th June. Specialist temperature controlled air-ride vehicles were used to transport the equipment to Edinburgh from Heathrow airport. The equipment was installed by the local Cray team at the ACF aided by X2 product specialists from Cray in the US.

The HECToR X2 at the ACF

The X2s connect directly into the High Speed Network of the XT4 systems and are adjacent to the XT4 cabinets. Unlike the XT4, which is air cooled, the X2 uses chilled water for cooling. One of the first tasks was to connect in the water hoses to the cabinets. Once power was applied, the cabinets were allowed to stabilise before being exhaustively tested with a suite of diagnostic software.

Integrating the X2s into the XT4s required a service shutdown and a number of configuration changes. Full accounting was enabled on the X2 in early August and early-access users from a number of projects are currently using the system. These early-access projects have been chosen to represent a broad range of sciences from across EPSRC 1, BBSRC2 and NERC3 funded research.

To find out more about the HECToR service, including how to apply and the cost of access for both the XT4 and X2 systems, see: www.hector.ac.uk.

1. www.epsrc.ac.uk
2. www.bbsrc.ac.uk
3. www.nerc.ac.uk

Two-phase jet simulations on HPCx

Xi Jiang, School of Engineering & Design, Brunel University (UK Turbulence Consortium e01).

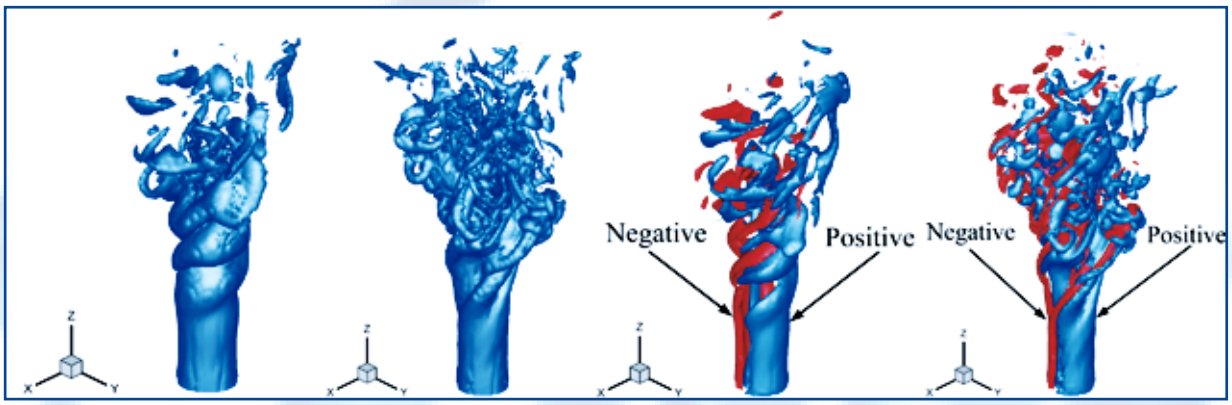

(a) Enstrophy $\rho_l / \rho_g = 2.5$ (b) Enstrophy $\rho_l / \rho_g = 5.0$ (c) x-vorticity $\rho_l / \rho_g = 2.5$ (d) x-vorticity $\rho_l / \rho_g = 5.0$

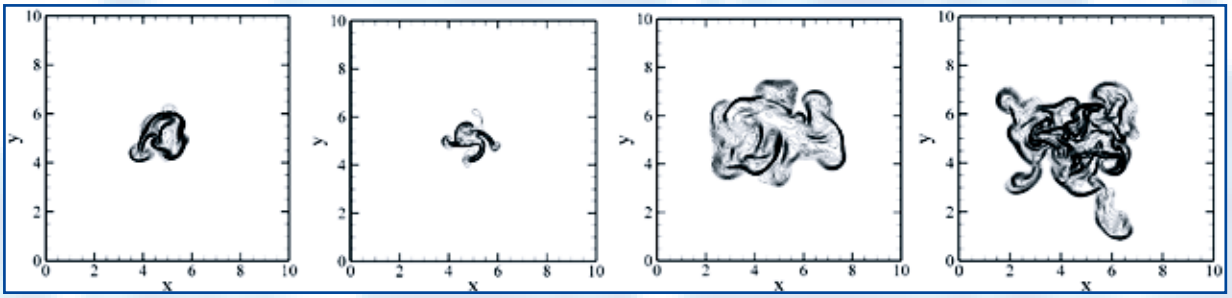

(a) $z = 4$ plane $\rho_l / \rho_g = 2.5$ (b) $z = 4$ plane $\rho_l / \rho_g = 5.0$ (c) $z = 8$ plane $\rho_l / \rho_g = 2.5$ (d) $z = 8$ plane $\rho_l / \rho_g = 5.0$

Figure 1 (top). Instantaneous isosurfaces of enstrophy and x-vorticity of the annular swirling gas-liquid two-phase jets (Siamas & Jiang, 2008).

Figure 2. Instantaneous liquid volume fraction contours in two streamwise locations of the annular swirling gas-liquid two-phase jets (Siamas & Jiang, 2008).

Challenging gas-liquid two-phase jet flows have been simulated in Dr Xi Jiang's group at Brunel University using the HPCx, supported by the UK turbulence consortium (e01), which is led by Dr Gary Coleman at University of Southampton. The gas-liquid two-phase jet flow is a multiphysics problem involving multiple time and length scales, which has practical relevance to liquid atomisation and spray processes. A direct numerical simulation (DNS) of such a flow provides a possibility to gain insight into the underlying flow physics by resolving the relevant time and length scales in the flow. In addition, the DNS database can also be used to assess and develop atomisation and spray models for CFD applications of gas-liquid two-phase flow systems.

DNS and proper orthogonal decomposition have been used to examine the near-field dynamics of annular swirling gas-liquid two-phase jets. Based on an Eulerian approach with mixed fluid treatment, combined with an adapted volume of fluid method and a continuum surface force model, a mathematical formulation for the flow system was developed. The swirl introduced at the jet nozzle exit is based on analytical inflow conditions. Highly-accurate numerical methods have been utilized for the solution of the unsteady, three-dimensional Navier-Stokes equations. Several computational cases have been performed to investigate the effects of swirl, liquid-to-gas density ratio on the flow field. Parallel computations have been performed on HPCx using a mesh of around 134 million (512×512×512) cells with each case costing around 300,000 AUs on the HPCx.

Fig. 1 shows the instantaneous enstrophy and x-vorticity distributions of two annular swirling gas-liquid two-phase jets with different liquid-to-gas density ratios, while Fig. 2 shows the instantaneous liquid volume fraction contours in two streamwise locations of these two cases. In both cases the flow becomes more vortical at downstream locations. Unsteady vortical flow characteristics are observed . It was found that the higher density ratio case is more vortical with larger spatial distribution of the liquid at downstream locations, in agreement with linear theories. Proper orthogonal decomposition analysis revealed that more modes are of importance at the higher density ratio, indicating a more unstable flow field. In the lower density ratio case, both a central and a geometrical recirculation zone are captured while only one central recirculation zone is evident at the higher density ratio. The results also revealed the formation of a precessing vortex core at the high density ratio, indicating that the precessing vortex core development is dependable on the liquid-to-gas density ratio of the two-phase flow, which is different from the dependence on the swirl number alone in single-phase gas jets.

Vive la difference! HPCx reaches out to life scientists

Eilidh Grant, Jon Hill, Nicola McDonnell and Alisdair Tullo (EPCC, University of Edinburgh); Marco Grzegorczyk and (CSBE, University of Edinburgh); Adrian Mulholland and Christopher Woods (University of Bristol).

No one would argue that high-performance computing is traditionally the domain of chemists or physicists, mathematicians or meteorologists, but life scientists? While a number of successful life science projects have burnt hours on HPCx, they tended to be in relatively traditional fields, for example molecular modelling. A new initiative was laid out in the '2008 HPCx Annual Plan' [1]to reach out to key UK life scientists, looking for novel users of HPCx as part of its complementary role and identifying and overcoming the barriers they face when accessing HPC. The resulting HPCx outreach team, which has expertise in life sciences and HPC alike, offers applications support dedicated to novel life sciences projects.

Many scientists working in the life sciences use programming languages not traditionally associated with HPC, including Perl, Python, R and MATLAB. The reasons for using these are many and varied. In the case of R and MATLAB, there are existing toolkits which make these particularly attractive to anyone requiring advanced mathematical and statistical capabilities. Some of these are quite specialised. For example, the open-source Bioconductor [2] suite, for the analysis and comprehension of genomic data, extends R with a set of highly useful open source tools for biologists working with genomic data.

With Perl and Python, mathematical and scientific libraries are available but there are also other advantages. These higher-level, interpreted languages often enable faster development through compact, concise syntax, and standard libraries for other less specialised tasks such as text processing (in the Python community, the latter is referred to as the "batteries included" approach). It's common to use languages such as Perl and Python to drive a compiled code. A program in a language like Python might be used to generate sets of parameters for different runs of a compiled C code. Some of these are already available on HPCx – for example, Python is already installed and in use. In other cases some work will be necessary to make these run in parallel on HPCx.

The science and technology of several of the on-going collaborations between EPCC and life scientists are detailed below.

Gene regulatory network inference

The ultimate objective of systems biology is the elucidation of the regulatory networks and signalling pathways of the cell. These pathways where the expression of certain genes or activation of proteins regulates other genes and proteins can be represented as a graphical model. [3]

There is an increasingly large amount of microarray (gene expression) and proteomics data available from which to infer these networks. Bayesian network methodology can be seen as a marriage of graph and probability theory and provides the most promising and flexible modelling tool for reverse engineering gene regulatory networks from expression data. Bayesian networks are based on directed (acyclic) graphs which assert independence assumptions among the variables. The regulatory relationships are explained by conditional probability distributions.

Two biological systems being investigated using this methodology by CSBE[4] Research Fellow Marco Grzegorczyk are the response of the immune system to infection and the internal clock of the plant Arabidopsis thaliana.

Interferons (IFNs) play a pivotal role in the innate and adaptive mammalian immune response against infection, and central research efforts therefore aim to elucidate their regulatory interactions. Gene expression time-series data was obtained from macrophages either infected with cytomegalovirus, treated with interferon gamma, or both infected and treated. This was then analysed to derive the network of interactions between interferon regulatory factors which are the key regulators in the response of the macrophage cell to pathogens. They mediate the cellular signalling that leads to a transcriptional response to the initial binding events on the surface of the cell.

Plants respond to the day/night cycle by changing their gene expression. They have an internal clock that regulates this circadian rhythm. The method has also been applied to two gene expression time series from A. thaliana cells, from plants entrained to either a 10h:10h light/dark-cycle or a 14h:14h light/dark-cycle[5] . The analysis is focusing on 9 genes which are known to be involved in circadian regulation.

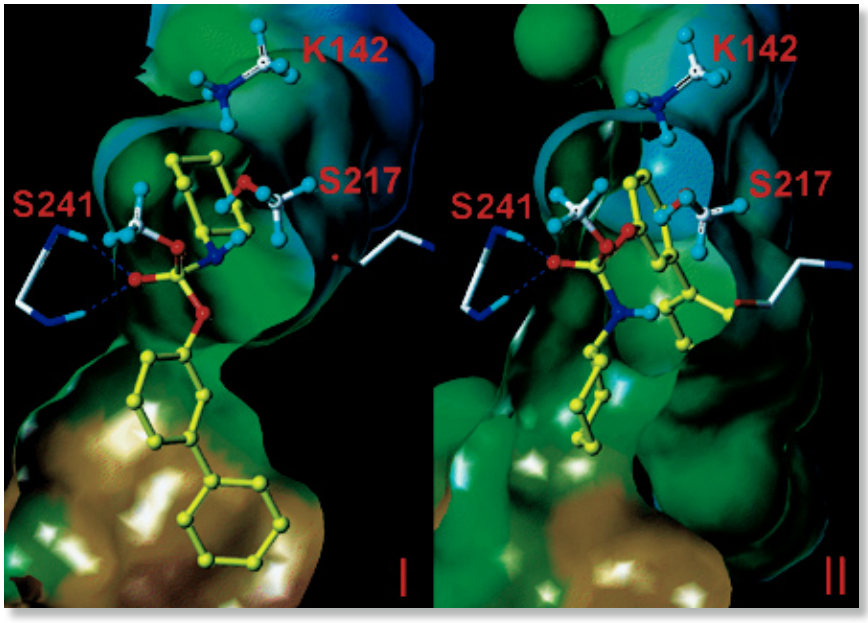

Figure 1: Representation of the tetrahedral intermediate for binding orientations I (left) and II (right) of the FAAH ligand URB524. Carbon atoms of FAAH are coloured in white, and those of URB524 in yellow. The surface of the pocket is shown as a Connolly channel and is colour-coded according to lipophilicity (brown: lipophilic; cyan: hydrophilic). H-bonds between the carbamic oxygen and the oxyanion hole are shown with dotted blue lines.

Marco Grzegorczyk said: "MCMC based Bayesian network inference is computationally expensive so that our current applications are restricted to small domains with few genes only. The implementation of the Metropolis-coupled Markov chain Monte Carlo algorithms (on HPCx) will allow us to infer more interesting networks with much more genes."

The code for this project is currently entirely written using MATLAB® [4]. In technical terms, EPCC's aim will be to develop a parallel code on HPCx which is functionally equivalent to this serial MATLAB code. Unfortunately, MATLAB is no longer available for the AIX operating system, so the simplest option – to parallelise the code in MATLAB on HPCx – is not possible.

As part of this project, several options will be investigated. Firstly, Octave[6] is a free software mathematical language which is largely compatible with MATLAB, so it may be possible to run (and parallelise) the original code using this. Secondly, Python has a toolkit with a similar interface (and similar capabilities) to MATLAB – namely NumPy [7]. It might be possible to quickly port the code to Python using this library. Python also has MPI libraries for parallelisation. Finally, the code could be rewritten in C or FORTRAN. Some crossover between these options is possible – we could write performance-critical parts of the code in C or FORTRAN, and call this from Octave or Python.

Modelling biological catalysts

HPC resources are increasingly helping to illuminate and analyse the fundamental mechanisms of biological 'molecular machines'. An example is enzyme catalysis. Enzymes are very efficient natural catalysts. Understanding how they work is a vital first step to harnessing their power for industrial and pharmaceutical applications.

For example, many drugs work by stopping enzymes from functioning. Endocannabinoids [8] can reduce pain and anxiety. These are molecules produced by our own bodies that are similar to the active ingredient of cannabis. The enzyme fatty acid amide hydrolase (FAAH) catalyses the break down of the endocannabinoid anandamide. Blocking the activity of FAAH is therefore a promising target for drugs designed to treat pain, anxiety and depression.

Adrian Mulholland's team at Bristol University has been using quantum mechanics/molecular mechanics (QM/MM) modelling to understand how inhibitors of FAAH bind to its active site [9]. The traditional molecular modelling method that tries to fit the shape of the ligand into the active site came up with two possible orientations for the ligand. It was necessary to model the interactions at a quantum scale to distinguish which of these orientations has a lower energy and so fits with the experimental data. This understanding will help with the design of drugs that can inhibit FAAH.

Atomically detailed computer models of enzyme-catalysed reactions provide an insight into the source of an enzyme's power [9] [10]. Due to the large size of biological molecules, simplified classical models of atomic interactions are used. These molecular mechanics (MM) models have been used successfully to understand the molecular dynamics of proteins.

However, MM can provide only a low-quality model of a chemical reaction, as electrons are represented implicitly. The best quality chemical models are provided by quantum mechanics (QM). QM calculations are highly computationally expensive, so it would be challenging to solve a QM model of an entire enzyme system. One solution is to use multiscale methods [11] that embed a QM representation of the reactive region of the enzyme within an MM

Continues on next page.

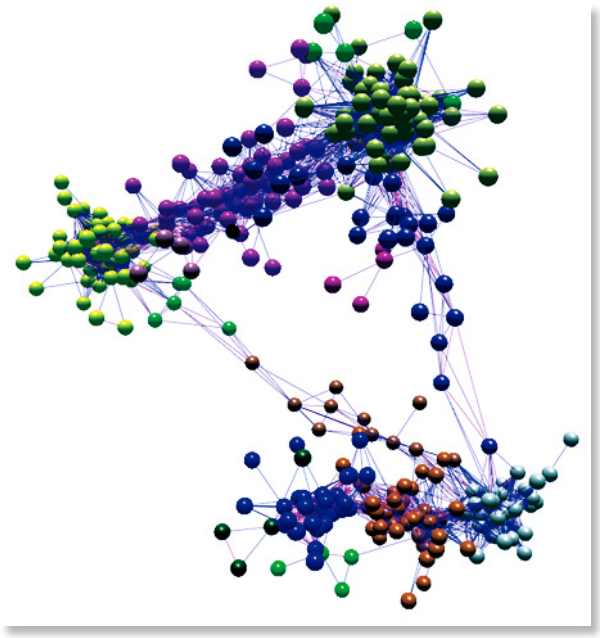

Fig. 2: The expression of many genes across different samples, clustered into groups with similar expression levels (indicated by colour).

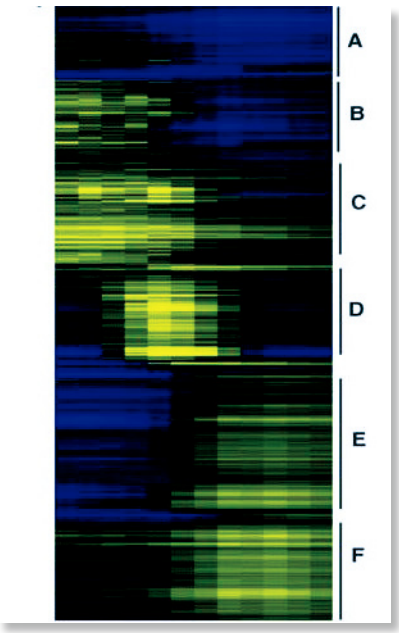

Fig. 3: A gene network based on the similarity of genes when measured across multiple samples and conditions. The layout places genes that correlate highly, closer to each other.

Created with Biolayout: www.biolayout.org

model of the rest of the system. Multilevel simulations of biological systems scale poorly over the many processors available on an HPC resource. New multiscale modelling methods [12] that split a single calculation into an ensemble of loosely-coupled simulations, are therefore a promising new direction to utilize maximum computing power. The aim is to make best use of the large numbers of processors by effectively coupling multiple individual simulations into a single supra-simulation. This method, applied on an HPC resource, promises to lead to a step change in the quality of the modeling of enzyme-catalysed reactions, and will provide new insights into these remarkable biological molecules.

This project is technically interesting in that it uses two codes, a quantum level code called Molpro and a molecular level code called Sire. "Molpro is a complete system of ab initio programs for molecular electronic structure calculations." [13] "[Sire] is designed to provide a complete framework for the rapid development of new molecular simulation algorithms." [14] Sire instances communicate using MPI, whereas Molpro uses OpenMP. The workflow controlling these two codes is a Python program. This project is ideally suited to HPCx, as it can offer the ability to run Python and mixed MPI/OpenMP jobs.

R

Microarrays measure the expression of genes. Genes produce RNA and this binds to spots of DNA on the microarray. Each spot binds specifically to the RNA from a particular gene. This creates a snapshot of the gene expression. Taking several snapshots over time, it's possible to work out which genes are affecting each other.

Microarrays can also be used to investigate gene expression under different conditions. Bioinformaticians perform statistical analyses on the data sets produced by these microarrays and other similar technologies. The number and size of the datasets are increasing as these technologies become cheaper. In genomic analysis the collected datasets are currently of the order of 1010 data points, 105 sequences by 105 samples, but are expected to get much bigger. Such large datasets can outgrow the available memory on the computer systems used to perform the statistical analyses. In addition these datasets can be complex to process, leading to excessive compute times. Whilst a small dataset may only require 8Mb of memory, the analysis of such a dataset results in an output of 800Mb. Memory sizes in excess of 10Gb are required to store the largest current datasets even for an analysis on a small subset of the data with terabytes of storage required for a complete analysis.

Bioinformaticians use the open-source package R for these statistical analyses. They are constrained by the memory size of their desktop computers. By installing R on HPCx, scientists from the Division of Pathway Medicine in the University of Edinburgh will be able to perform analyses requiring up to 128GB of RAM.

The SFC-funded edikt2 project, Parallel R [15] has designed, built and tested a prototype system that will allow statistical analyses to be done in parallel on standard HPC machines. This prototype will help investigate how to allow statistical analyses to be performed in parallel with minimal editing of existing R scripts and minimal

Continues on next page.

Forthcoming events

26-27 Nov 2008
High Performance Reconfigurable Computing: using FPGAs for HPC; EPCC, Edinburgh.

4 Dec 2008
Parallel I/O Workshop; STFC Daresbury Laboratory.

5 Dec 2008
Course: Parallel IO using MPI-I/O; STFC Daresbury Laboratory.

For more details and a registration form see: www.hpcx.ac.uk/about/events/.

Attendance at all events is free for academics.

STFC Daresbury and EPCC at SC08 Nov 15–21, Austin, Texas

Music is in the air at this year's International Conference for High Performance Computing, Networking, Storage and Analysis; SC08. The conference is being held in Austin, Texas: billed as the "live music capital of the world". A number of extra activities have been added to the conference schedule to encourage attendees join in the musical theme, including: ViSCiTunes, a call for people to add music to their scientific visualisations and submit them to a conference wide playlist; a competition to try to identify songs generated by a computer and those composed with real instruments; and a live music room for attendees to perform in.

Away from the musical theme, SC08 marks the 20th anniversary of the first conference, then known as Supercomputing, held in Orlando, Florida in 1988. A number of events are planned to celebrate this event through the duration of the conference. The conference committee is also looking for contributions from people who have attended all 20 SC conferences. The 20th SC conference also sees a broading of the scope of the program with two Technology Thrusts, one focusing on renewable energy and energy efficiency research, the second focusing on biomedical informatics including computational and systems biology.

If you are at SC08, you can find EPCC at booth number 2329, alongside STFC Daresbury Laboratory in booth 2325.

www.sc08.supercomputing.org

knowledge of HPC systems. An application has been made for a Welcome Trust grant to allow the completion of the Parallel R project and make it available on HPCx.

Conclusion

Life scientists have an ever-increasing need for HPC. They have special requirements which are different to those of standard HPC user such as large memory or Perl, Python, R and MATLAB languages. HPCx, in its complementary role, has an outreach program which offers support to life scientists migrating to the system. All this work will pave the way for the life scientists of the future.

Footnotes

1. Centre for Systems Biology, University of Edinburgh

2. Data were collected by the Institute of Molecular Plant Sciences, School of Biological Sciences, University of Edinburgh.

References

1. Annual Plan for Science Support and Capability Computing, 2008

2. Bioconductor, http://www.bioconductor.org/

3. Werhli, A.V.; Grzegorczyk, M.; and Husmeier, D. (2006): Comparative evaluation of reverse engineering gene regulatory networks with relevance networks, graphical Gaussian models and Bayesian networks, Bioinformatics, 22, 2523-2531

4. Grzegorczyk, M.; Husmeier, D.; Edwards, D.E.; Ghazal, P.;and Millar, A.J. (2008): Modelling non-stationary gene regulatory processes with a non-homogeneous Bayesian network and the allocation sampler. Bioinformatics, 24, 2071-2078.

5. MATLAB®: http://www.mathworks.com/products/matlab/

6. Octave: http://www.gnu.org/software/octave/

7. NumPy: http://numpy.scipy.org/

8. Endocannabinoid: http://www.nature.com/nrd/journal/v2/n2/full/nrd1024.html

9. Alessio Lodola; Marco Mor; Silvia Rivara; Christo Christov; Giorgio Tarzia; Daniele Piomelli and Adrian J. Mulholland: Identification of productive inhibitor binding orientation in fatty acid amide hydrolase (FAAH) by QM/MM mechanistic modeling, http://www.rsc.org/delivery/_ArticleLinking/ArticleLinking.asp?JournalCode=CC&Year=2008&ManuscriptID=b714136j&Iss=2

10. van der Kamp, M.W., Shaw, K.E., Woods, C.J. and Mulholland A.J., J. Royal. Soc. Int., Biomolecular simulation and modelling: status, progress and prospects, http://dx.doi.org/10.1098/rsif.2008.0105.focusMulholland, A.J., Biochem. Soc. Tran., 36, 22-26, 2008

11. Woods, C.J., Mulholland, A.J., Multiscale modelling of biological systems in RSC Special Periodicals Report: Chemical Modelling, Applications and Theory, Volume 5, 2008

12. Woods, C.J., Manby, F.R and Mulholland, A.J., J. Chem. Phys. 123, 014109, 2008

13. Molpro http://www.molpro.net/

14. Sire http://siremol.org/home

15. Parallel-R: http://www.edikt.org/edikt2/ParallelRActivity

URANS, DES and ILES approaches for a complete helicopter fuselage flow simulations with, and without, actuator disks modelling

Bowen Zhong and Dimitris Drikakis, Department of Aerospace Sciences, Cranfield University, UK

Frederic Le Chuiton, Eurocopter Deutshland (ECD), Germany

Thorsten Schwarz and Jochen Raddatz, Deutsches Zentrum fuer Luft und Raumfahrt (DLR), Germany

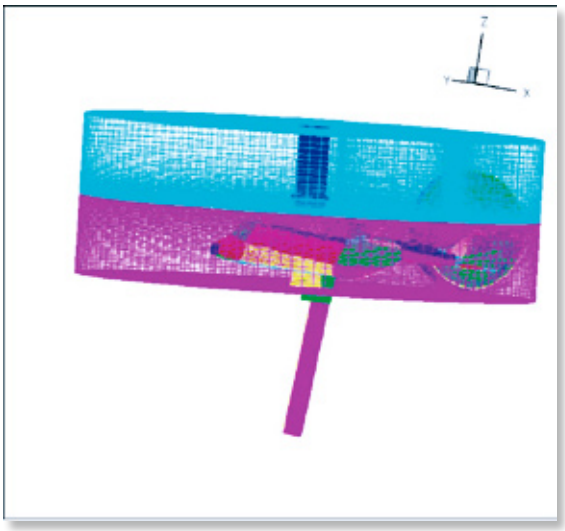

Fig. 1 The fuselage surface grid and the grids for the actuator disks.

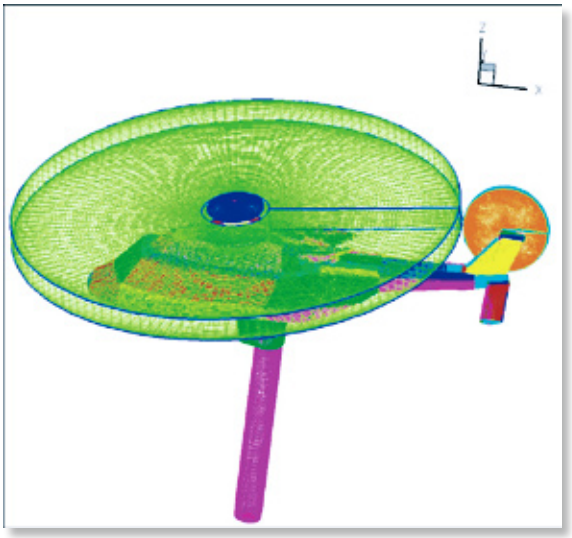

Fig. 2 Hole grids for Chimera approach.

Intensive turbulent separation frequently occurs behind the rotor head, engine fairing, and afterbody of the helicopter fuselage. This poses great challenges to computational fluid dynamics (CFD) simulations with respect to both accuracy and efficiency.

As part of the work of the European 6th Framework project GOAHEAD (Generation of Advanced Helicopter Experimental Aerodynamic Database for CFD code Validation) [1,2], the authors have carried out a comparative study of different simulation strategies and numerical methods for turbulent flows around a complete helicopter fuselage in cruise conditions. The computations have been carried out using the national HPCx facility and Cranfield University's Astral cluster.

The implementation of different methods was performed in the framework of the FLOWer code developed by DLR. The HLLC Riemann solver, a modification of HLL Riemann solver proposed by Toro [3], was implemented in the course of this study and comparative studies were carried out using HLLC Riemann solver in conjunction with various flux limiters [4]. The results were further compared with the Jameson's finite volume central difference scheme. Four different computational strategies were compared, namely RANS, unsteady RANS (URANS), Detached Eddy Simulation (DES) [5] and Implicit Large Eddy Simulation (ILES) [6]. The RANS and URANS simulations have also been performed using the actuator disk helicopter-modelling approach in order to model the effects of the main and tail rotors on the flow field and, consequently, the forces and the moments acting on the fuselage.

The grid consists of over 120 blocks and 10 million grid points. Chimera grids are used for convenience in the calculations of the rotor head, horizontal stabilizer and the strut. In the Actuator Disk (AD) approaches, the main and tail rotors are replaced by actuator disks. The flow simulations are carried out with a forward flight Mach number of 0.204 at an angle of attack -2 degrees.

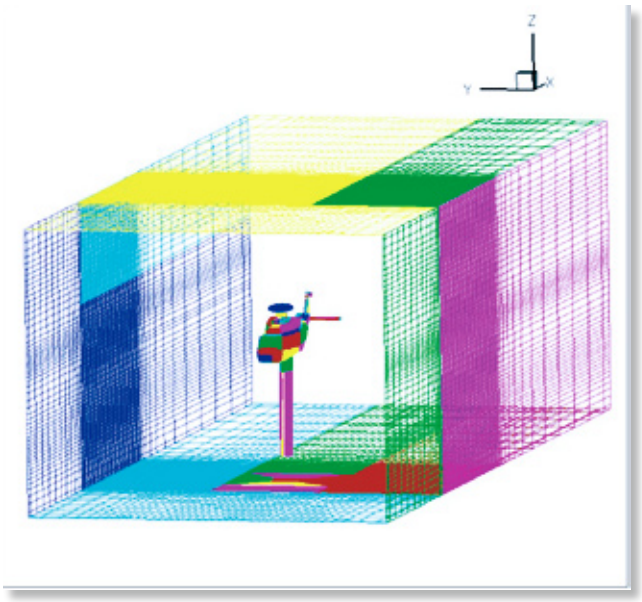

Fig. 3: The fuselage surface grid and the grids on the wind tunnel wall.

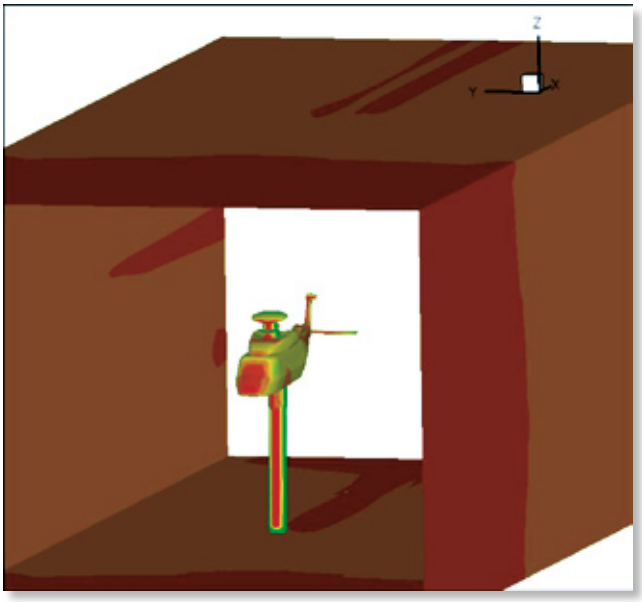

Fig. 4: Pressure contours on fuselage surface.

Figure 1 shows the fuselage surface grid and the Chimera grids used in the AD approach. Figure 2 shows the hole grids used in the Chimera grid approach for the rotor disks. Figure 3 shows the computation set-up of the fuselage surface grids in the relation to the wind tunnel wall.

RANS calculations have been carried out using HLLC Riemann solver and Jameson's central difference scheme. Figure 4 shows the pressure contours on the fuselage surface and the wind tunnel wall obtained using RANS with HLLC scheme. Figure 5 shows the comparison of total forces and moments on the fuselage with and without AD using the HLLC and central difference schemes. The results indicate that Cx is slightly increased and the Cz is decreased compared to the case without AD.

Due to the inherent problem in the formulation of RANS and URANS, it is not possible for RANS or URANS to predicate flows with large separations accurately [7]. To improve the prediction, DES and ILES simulations have also been performed. Figure 6 shows the differences of the vorticity iso-surfaces obtained from RANS, URANS, DES and ILES, respectively. It shows that DES and ILES can capture the instantaneous vortex shedding in the separation region much better than RANS or URANS. ILES shows much finer vortices structures have been captured, which indicates that ILES is less dissipative than DES.

Acknowledgement

This research is funded by the EU 6th framework GOAHEAD project. The support from Drs. Walid Khier, Klausdieter Pahlke, and Prof. Norbert Kroll of DLR are greatly appreciated. The computing time from the national supercomputer HPCx facility awarded under EPSRC grant number EP/F005954/1 for the UK Applied Aerodynamics Consortium 2 (UKAAC2) is very much appreciated.

Continues on next page.

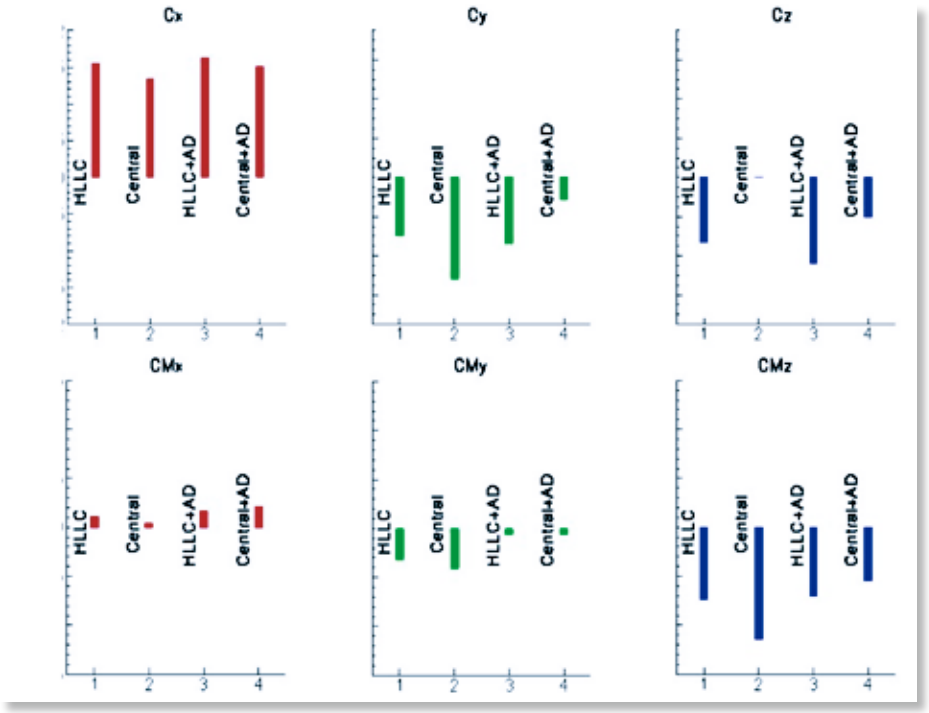

Fig. 5: Total forces and moments on fuselage with/without ADs.

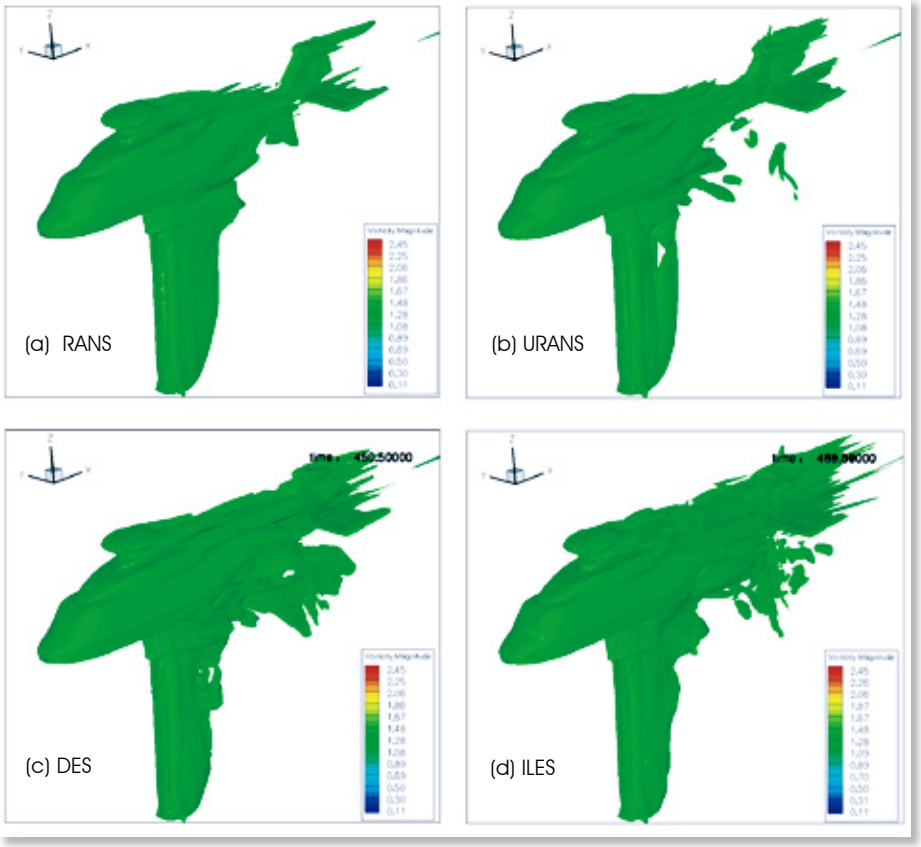

Figure 6. Vorticity iso-surfaces obtained using RANS, URANS, DES and ILES.

References

[1] K. Pahlke, "The GOAHEAD project", Proceedings of 33rd European Rotorcraft Forum, Kazan, Russia, September, 2007

[2] O. J. Boelens, "The blind-test activity of the GOAHEAD project", Proceedings of 33rd European Rotorcraft Forum, Kazan, Russia, September, 2007

[3] E. F. Toro, Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction, Springer, 1999.

[4] D. Drikakis, W.J. Rider, High-Resolution Methods for Incompressible and Low-Speed Flows, Springer, 2005.

[5] Spalart PR et al "Comments on the feasibility of LES for wings, and on a hybrid RANS=LES approach", First AFOSR International Conference on DNS and LES, Ruston, LA, 4–8 August 1997.

[6] Grinstein F. F, C. Fureby, "Monotonically Integrated Large Eddy Simulation of Free Shear Flows", AIAA Journal, Vol.37 no.5, 1999.

[7] B. Zhong and P. G. Tucker, k-l Based Hybrid LES/RANS Approach and Its Application to Heat Transfer Simulation, International Journal for Numerical Methods in Fluids, Vol. 46, No.10, 2004

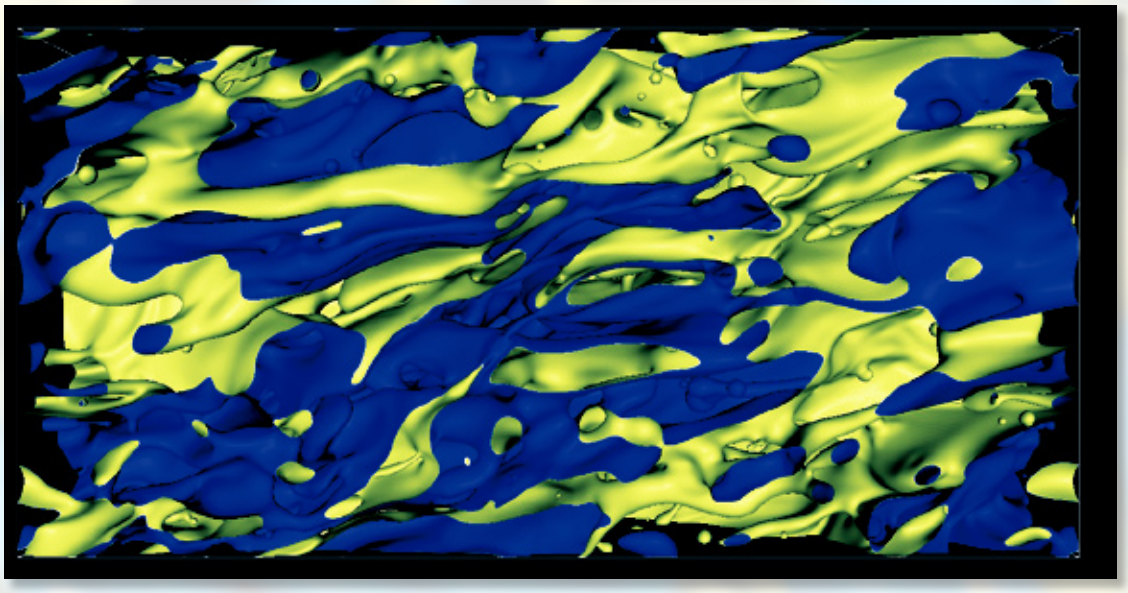

A still image from a lattice Boltzmann simulation of a binary fluid mixture under steady shear. The extreme elongation of the fluid domains in the flow direction (here to the right at the top of the box, and to the left at the bottom) means that large system sizes are required to prevent finite size effects. The rendering shows the interface between the fluids, for which two closely spaced isosurfaces are used – each different colour represents the different fluids.

ParaView – parallel visualisation on HPCx

Iain Bethune and Kevin Stratford, EPCC, University of Edinburgh

Being able to see the results of calculation can be a particularly useful way to detect whether, broadly, things are as you expect. Where dynamical simulations are concerned, an animation can be even better: the human eye is particularly good at spotting physical (or unphysical) features. This is particularly true in subjects such as fluid dynamics, where judgement can be guided by common experience.

Producing high quality animations, however, takes a degree of care, if it can be done at all for the largest simulations. So in this context, any tool that makes life easier is very useful. One such tool is ParaView [1], an open source parallel visualisation packages from Kitware Inc, Sandia National Laboratories, and others.

As part of the move towards Complementary Capability Computing this software has been installed on HPCx and is available to all users. Further information can be found in the HPCx Technical Report "Parallel Visualisation on HPCx" [2].

In recent work as part of the Tera Scale Challenge Project using the HECToR machine by Edinburgh Soft Matter and Statistical Physics Group and EPCC, we have used ParaView to provide animations of results of lattice Boltzmann simulations of fluid mixtures under shear [3]. These simulations are relatively large (e.g., the one shown in Figure 1) is based on a 512x1024x512 lattice. This means rendering a single frame of the animation for the composition field requires around 1GB of data. ParaView's ability to do this in parallel makes this possible in reasonable time. A python script can be used to automate the task of rendering a large number of frames (up to 1000 have been used to provide animation of these systems, representing about 1TB of data). The results can be saved as either a series of still images for further manipulation, or may be automatically used to produce an animation via a key-frame animation capability.

References:

[1] ParaView website, www.paraview.org

[2] Technical Report: "Parallel Visualisation on HPCx", http://www.hpcx.ac.uk/research/hpc/technical_reports/HPCxTR0803.pdf

[3] K. Stratford, J.C. Desplat, P. Stansell, M.E. Cates, "Binary fluids under steady shear in three dimensions" Phys. Rev. E 76 030501(R) (2007).

Figure 1. A) Inner rectangle: rupture area of the 19/09/1985 Ms 8.1 earthquake on surface projection of the 500x600x124 km earth crust volume 3DFD discretization; B) profile P-P´; C) Kinematic slip distribution of the rupture of the 1985 earthquake [1].

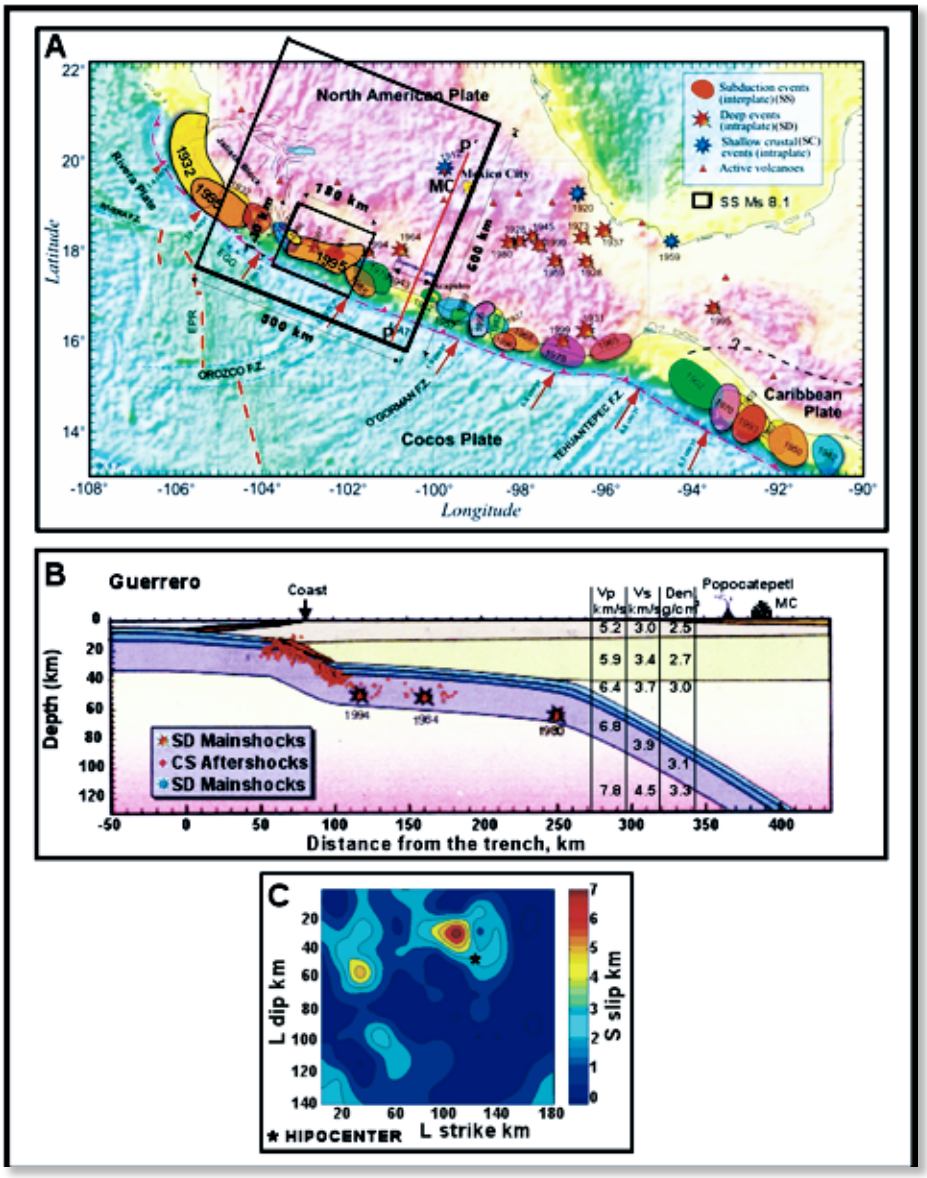

3D parallel modelling of large subduction earthquakes

Mario Chavez (Institute of Engineering) and Eduardo Cabrera (DGSCA) UNAM, Mexico

Realistic 3D modeling of the propagation of large subduction earthquakes, such as 1985's Mexico earthquake (Fig. 1), poses a numerical and computational challenge, particularly as it requires enormous amounts of memory and storage, as well as intensive use of computing resources.

About 30,000 people died in the Mexico event and the economic loss was around $US 7 billion. As the estimated recurrence time for this highly destructive type of event in Mexico is only a few decades, there is seismological, engineering and socio-economic interest in modelling them by using parallel computing [1].

The propagation of the seismic waves generated by earthquakes can be expressed by the velocity-stress form of the 3D elastic wave equation, which on an isotropic medium consists of nine coupled, first-order partial differential hyperbolic equations, for the three particle velocity vector components and the six independent stress tensor components. The medium may be described by the Lame parameters and the mass density, and the seismic source can be represented by a moment source tensor.

A 3D finite difference (3DFD) staggered explicit scheme, 2nd order accurate in time and 4th order accurate in space, was applied to the mentioned equations. Staggered grid storage allows the partial derivates to be approximated by centered finite differences, without doubling the spatial extent of the operators, thus providing more accuracy. Subsets of the dependent variables are stored on different spatial and temporal grids. For example, the three velocity vector components are stored at grid points that are shifted in both space and time by one-half grid interval [1, 2].

We use 3D data parallelism for efficiency: the domain where the seismic waves propagate was decomposed into small subdomains and distributed among a number of processors, using simple partitioning to give equal number of grid points to each processor. This approach is appropriate for the 3DFD wave propagation code, as large problems are too big to fit on a single processor [1].

Message Passing Interface (MPI) was used to parallelize the 3DFD code [1]. In particular MPI_Bcast, MPI_Cart_Shift and MPI_SendRecv instructions were used, the former to communicate the geometry and physical properties of the problem before starting the wave propagation loop, and the latter to update the velocities and stresses calculations at each time step. MPI-2 Parallel I/O was used in the code to read the Earth model data by all processors and to write the velocity seismograms by the processors corresponding to the free surface of the physical domain, which is only a small percentage of the total number of processors. This type of parallel I/O is machine independent, therefore it fitted a benchmark experiment performed on four platforms [7].

As part of the SCAT European Community project, in which UNAM and the STFC Daresbury Laboratory participate, the code was run on four dual-core platforms: KanBalam (KB) [3] in Mexico; BlueGene/P (BGP) [4], HPCx [5]) and HECToR [6] in the UK. Part of the results obtained from this collaboration are reported in [7].

The actual size of the problem is 500 x 600 x 124 Km. The domain and its physical properties are shown in Fig. 1. We used spatial

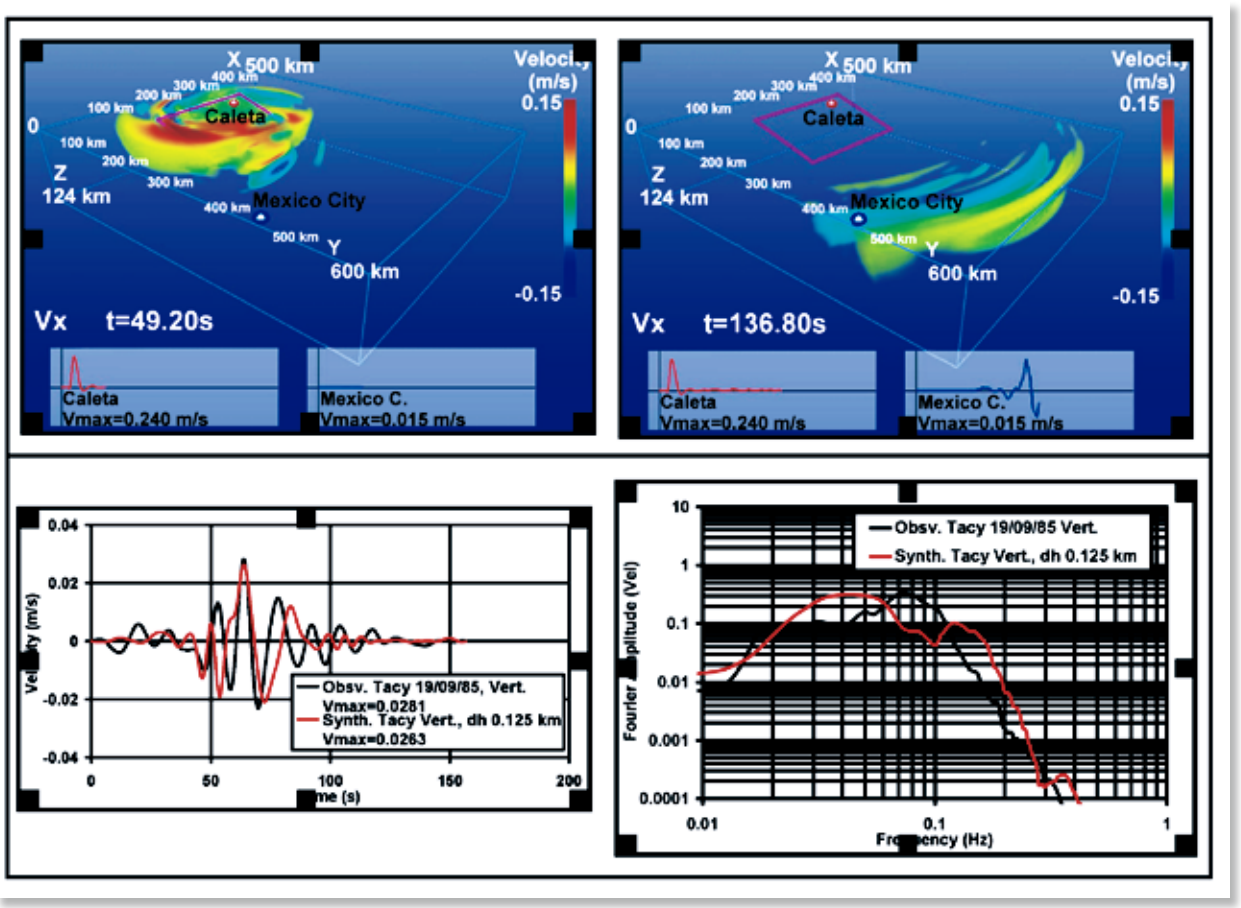

Figure 2. Top: 3D Snapshots of the velocity wave-field in the X direction of propagation for t = 49.2 and 136.8s for the 1985 Mexico earthquake.

Bottom: Left side observed and synthetic seismograms at Mexico City, right side Fourier amplitude spectra.

discretizations (hx = hy = hz) of 1.0, 0.500, 0.250 and 0.125 Km, and the associated time discretizations were 0.03, 0.02, 0.01 and 0.005 s, respectively (to comply with the Courant-Friedrich-Lewy condition). Therefore, Nx=500, 1000, 2000, 4000; Ny=600, 1200, 2400, 4800 and Nz=124, 248, 496, 992 are the model sizes in the X, Y and Z directions, respectively. The number of time steps, Nt used for the experiment was 4000. Speedup and efficiency results for different size models, number of processors (P) from 1- 4096 and with one or two cores, whenever the platform architecture allowed, for KB, HECToR, HPCx and BGP, were reported in [7]. Among other results it was concluded that when large amounts of cores (≥1024) were used per processor, in KB and HECToR, the speedup decreased considerably, due to the very large number of communications required by the 3DFD algorithm [1, 7].

In Figure 2, we present examples of the type of results, that for the 1985 Mexico earthquake (Fig. 1), were obtained in KB and HECToR with the 3DFD parallel MPI code implemented. At the top of Fig 3, the 3D low frequency velocity field patterns in the X direction, and the seismograms obtained at observational points, in the so-called near (Caleta) and far fields (Mexico City), of the wave propagation pattern for times equal to 49.2 and 136.8 s.

The complexity of the propagation pattern at t = 49.2 s, when the seismic source (Fig. 1) is still rupturing, is contrasted by the one for t = 136.8 s, in which packages of coherent, high amplitude, well developed surface waves, are propagating towards Mexico City.

Finally, at the bottom of Figure. 2 we show the observed and synthetic (for a spatial discretization dh = 0.125km) low frequency, North-south velocity seismograms of the 19/09/1985 Ms 8.1 Mexico earthquake, and their corresponding Fourier Amplitude spectra for the firm soil Tacubaya site in Mexico City, i.e. at a far field observational site. Notice in Fig. 3, that the agreement between the observed and the synthetic velocity seismogram is reasonable both in the time and in the frequency domain.

From the seismological, engineering and socio economic side, one of the most relevant results that can be drawn from this HPC experiment, is that the 3D velocity field patterns of the 1985 Mexico simulations obtained for the first time (running in KanBalam and HECToR) shows that, coherent, large amplitude, well developed surface waves, propagating towards Mexico City, were the fundamental cause of the loss of up to 30000 people and about 7 billion US dollars observed for this event.

We would like to thank the support of Genevieve Lucet, José Luis Gordillo, Hector Cuevas, the supercomputing staff and Marco Ambriz, of DGSCA, and the Institute of Engineering, UNAM, respectively. We acknowledge DGSCA, UNAM for the support to use KanBalam, as well as STFC Daresbury Laboratory to use HECToR, HPCx and BlueGene/P.

References

[1] Cabrera E., M. Chavez, R. Madariaga, N. Perea, M. Frisenda. 3D Parallel Elastodynamic Modeling of Large Subduction Earthquakes. F. Capello et al. (eds): Euro PVM/MPI 2007, LNCS 4757, pp. 373-380, 2007, Springer-Verlag Berlin Heidelberg 2007.

[2] S. E. Minkoff. Spatial Parallelism of a 3D Finite Difference Velocity-Stress Elastic Wave Propagation code. SIAM J. Sci. Comput. Vol 24, No 1, 2002, pp 1-19.

[3] http://www.super.unam.mx/index.php?op=eqhw

[4] http://www.cse.scitech.ac.uk/sog/

[5] http://www.hpcx.ac.uk/

[6] http://www.hector.ac.uk/

[7] Chavez M, E. Cabrera, R. Madariaga, N. Perea, Ch. Moulinec, D. Emerson, M. Ashworth, A. Salazar, Benchmark Study of a 3D parallel Code for the Propagation of Large Subduction Earthquakes, Submitted to Euro PVM/MPI 2008.

PRACE surveys the current European HPC ecosystem

Jon Hill, EPCC, University of Edinburgh

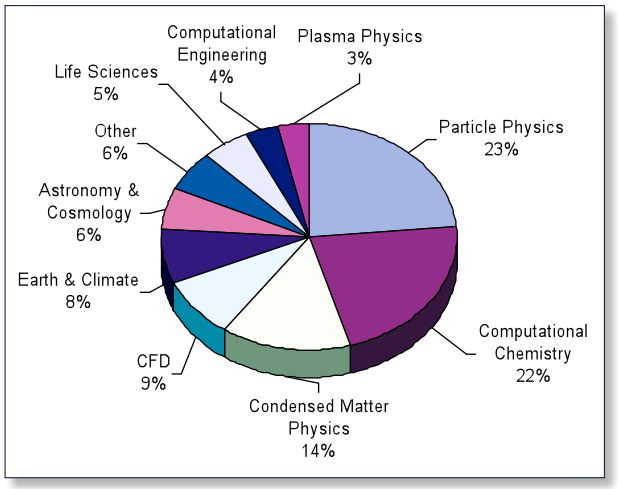

Figure 1: useage of European machines sorted by scientific area.

Area/Dwarf	Dense linear algebra	Spectral methods	Structured grids	Sparse linear algebra	Particle methods	Unstructured grids	Map reduce methods
Astronomy and Cosmology	0	0.62	4.91	3.59	5.98	2.99	0
Computational Chemistry	15.35	26.09	1.80	3.45	7.49	0.53	12.98
Computational Engineering	0	0	0.53	0.53	0	0.53	2.8
Computational Fluid Dynamics	0	1.70	7.37	3.05	0.32	3.00	0
Condensed Matter Physics	9.10	15.07	1.62	0.73	1.76	0.28	5.70
Earth and Climate Science	0	2.03	5.83	1.33	0	0.26	0
Life Science	0	4.72	0.94	0.13	0.94	0.28	3.46
Particle Physics	12.50	0	4.59	0.92	0.10	0	89.27
Plasma Physics	0	0	1.33	1.33	3.55	0.42	0.63
Other	0	0	0	0	0	0	0

Table 1. The utilisation matrix based on the survey results. There are 70 categories, 10 scientific areas and 7 algorithmic 'dwarves'. The figure in each cell is an estimate of the number of Tflop/s burned in each category. White boxes are those with no usage. Orange boxes are those with usage greater than zero, but less than 5 Tflop/s usage. Red boxes signify usage greater than 5 Tflop/s.

PRACE, the Partnership for Advanced Computing in Europe, has been established to prepare the ground for the creation of a persistent pan-European HPC service. As part of this activity, EPCC is leading Work Package 6 – Software for Petaflop/s Systems. This will involve porting applications to state-of-the-art systems and assessing how well they will scale to the ten or hundreds of thousands of cores needed to take full advantage of a Petaflop/s system. Some key applications will be packaged up into a benchmark suite which will be used in the procurement of future European Petaflop/s system. However, you first need to know which applications to target. The best start for predicting what will be run on a future system is to see what is currently being run today. This is exactly what EPCC undertook in its latest deliverable to the PRACE project.

The approach taken was aimed at understanding current usage and using this to see what might be run on a future system. To gain data on current usage, EPCC carried out a pan-European survey on the major HPC systems managed by the PRACE partners. Data on 24 systems and nearly 70 applications were collected, in what is believed to be the largest survey of its kind to date. The results make for an interesting read.

The systems surveyed represent 14 PRACE partners from 12 countries. The total power of systems is 926 Tflop/s peak, and 675 Tflop/s achieved Linpack, from a total of 169,522 cores. The largest machine in the survey is Jugene at FZJ in Germany. EPCC's HECToR was the third largest in terms of numbers of cores. However, the survey was not just about the machines, but also what the machines were used for. We asked, for each machine, the percentage of time used in ten scientific areas. Overall, the results show that most (45%) of the time is spent in either Computational Chemistry or Particle Physics. Condensed Matter Physics was the next highest scientific area (see Fig. 1).

We also asked about the top applications running on each system and how much of the system they used. The top applications were a number of Lattice QCD applications and chemistry packages, such as VASP, NAMD, DALTON and CPMD. GADGET (a cosmology simulation application) also featured. The kernels in the applications, as well as the libraries they depend upon were all determined from the survey. From this information we were able to build a 'usage matrix' of European HPC, which split the total cycles over algorithms and scientific areas (see table). We therefore have a relatively complete and detailed picture of the current HPC systems and applications.

The results from this survey not only gave a detailed picture of the current European HPC usage but were used to guide the choice of which applications would be suitable for inclusion in a representative application benchmark suite. The next step is to look at these applications in more details, determining their kernels, communication-to-computation ratio, etc. These applications will then be ported to prototype systems by the various PRACE partners in the coming months before undergoing optimisation and scaling. The information obtained from this work will be used to help other users of the Petaflop/s systems scale their own code.

UK scientists to get access to DEISA resources

Gavin J. Pringle, EPCC, University of Edinburgh

DEISA, Europe's HPC infrastructure, is currently evaluating applications from scientists wishing to employ DEISA to achieve groundbreaking science through the DEISA Extreme Computing Initiative (DECI).

The DEISA (Distributed European Infrastructure for Supercomputing Applications) EU-funded Research Infrastructure comprises a number of leading national supercomputers in Europe interconnected with a high bandwidth 10 Gb/s point-to-point network provided by GEANT and the National Research Networks. Selected middleware allows the deployment and operation of a number of services enabling high-performance distributed computing.

DECI, which was launched in early 2005, identifies, enables, deploys and operates flagship applications in selected areas of science and technology. These applications must deal with complex, demanding, innovative simulations that would not be possible without DEISA and which would benefit from the exceptional resources of the Consortium. Projects supported by DECI are chosen on the basis of innovation potential, scientific excellence and relevance criteria. Multi-national proposals are especially encouraged.

A European Call for Extreme Computing Proposals had an excellent response in 2008. Published annually in spring since 2005, each year has seen an increase in the number of proposals received, from 40 proposals in 2005 to 66 proposals this year. Indeed, the requested aggregated total is over 134 million CPU hours. This year there is a total of 8 proposals with UK PIs, requesting an aggregated total of over 31 million CPU hours.

Every year, each DEISA site commits a certain percentage of its platforms (typically 5% or more). This year the total available to DEISA as a whole is over 50 million CPU hours. In the UK, EPSRC has agreed to commit 5% of HPCx and 5% of HECToR, including both the XT4 scalar cluster and the X2 vector platforms. This amounts to over 5 million CPU hours.

This huge over-subscription for capability computing resources in Europe underlines the need for a persistent European HPC ecosystem.

The Applications Task Force, a European team of HPC experts led by Alison Kennedy from EPCC, supports the enabling of the applications in the projects to be used within the heterogeneous DEISA infrastructure and also helps to select the most suitable architecture for each project, depending on its specific requirements. In this way, DEISA is also opening up the most powerful HPC architectures available in Europe for the most challenging projects, mitigating the rapid performance decay of a single national supercomputer within its short lifetime cycle (typically about 5 years, as implied by Moore's law).

Petaflop computing unveiled at ISC2008

Dresden, Germany

Kenton D'Mellow, EPCC, University of Edinburgh

This year's International Supercomputing Conference [1] marked the dawn of a new era in supercomputing – the Petaflop era. The conference began with the traditional announcement of the Top500 [2], with IBM's Roadrunner [3] reaching the No 1 spot. Roadrunner weighs in at 1.026Pflop/s LINPACK benchmark, while running on 17 out of its 18 available frames. With the world's first Petaflop machine, IBM has raised the bar for future machines.

The excitement caused by Roadrunner provoked a debate that ran through the entire technical programme: the move to multi-core heterogeneous architectures, which are an increasingly attractive proposition in terms of both price/performance, and energy/performance. Solutions such as ASIC and FPGA-based accelerators, GPGPUs and of course IBM's Cell processor, are capable of providing unparalleled performance. However, this will require more complex algorithms and further application development. A seminar by Jack Dongarra, University of Tennessee, emphasised the need for more sophisticated high-performance algorithms to take advantage of the new compute power. Machines are growing at a faster rate than our ability to scale useful applications, and a community-wide increased software effort is now needed.

Another recurring theme of the conference was environmentally-aware 'green' computing. The Green500 [4] is an increasingly recognised measure, and terms like 'Flops per Watt' are becoming more common. Interestingly, Roadrunner itself is very green, within the top 3 in the Green500, achieving 750MFlops/Watt. As a result of these two strong trends, HPC's recent plethora of hardware options now seems likely to converge to a more manageable set of architectures.

Attendance at ISC continues to grow. There were over 1350 delegates and 88 exhibitors this year, reflecting the event's increasing international importance. ISC has outgrown the Conference Centre in Dresden and it was announced that ISC2009 will be held in Hamburg.

[*[1] www.supercomp.de/*
[2] www.top500.org/
[3] www-03.ibm.com/press/us/en/pressrelease/24405.wss
[4] www.green500.org/

HPCx Annual Seminar

e-Science Institute, Edinburgh

Chris Maynard and Chris Johnson, EPCC, University of Edinburgh

The sixth Annual HPCx seminar was held at the eSI in Edinburgh in June, with a focus on the science HPCx enables. It was extremely well attended, with more than seventy participants. The four user talks came from a diverse range of scientific disciplines: biochemistry, material chemistry, aerodynamics and condensed matter, all of which were well received. ECMWF speaker, Paul Burton, gave an excellent talk about medium-range weather forecasting and Ralph Warmack of IBM gave a status report on the BlueGene program. Talks about I/O performance and the future plans of HPCx complementary computing completed a balanced program of science and technology.

The 'Novel Parallel Programming Languages for HPC' workshop took place the following day with more than sixty participants. The workshop consisted of talks from vendors, language developers and researchers. The opening speaker was John Reid of Rutherford Appleton Laboratory who gave an excellent talk on Co-array Fortran. This was followed by talks from language developers working on HPC languages initially funded by the DARPA (Defense Advanced Research Projects Agency) High Productivity Computer System Program. This session included talks on Chapel (Cray), Fortress (Sun) and X10 (IBM). In the afternoon there was a talk on applications of UPC (Unified Parallel C) followed by a

talk on the use Object Oriented Languages for HPC. The event was rounded off with an open discussion on all the topics discussed during the day in which all of the speakers present participated. In particular there was some discussion about whether or not Co-array Fortran should be included as part of the Fortran 2008 standard, either as a part of the core standard or as an optional part; a decision is due to be taken on this very soon and workshop participants have been invited to comment.

With the HPCx service now extended until January 2010, these successful events show the continued importance of HPCX.

www.hpcx.ac.uk/about/events/annual2008/
www.hpcx.ac.uk/about/events/hpc_languages/