Capability Computing

The newsletter of the HPCx community

[ISSUE 7, SUMMER 2006]

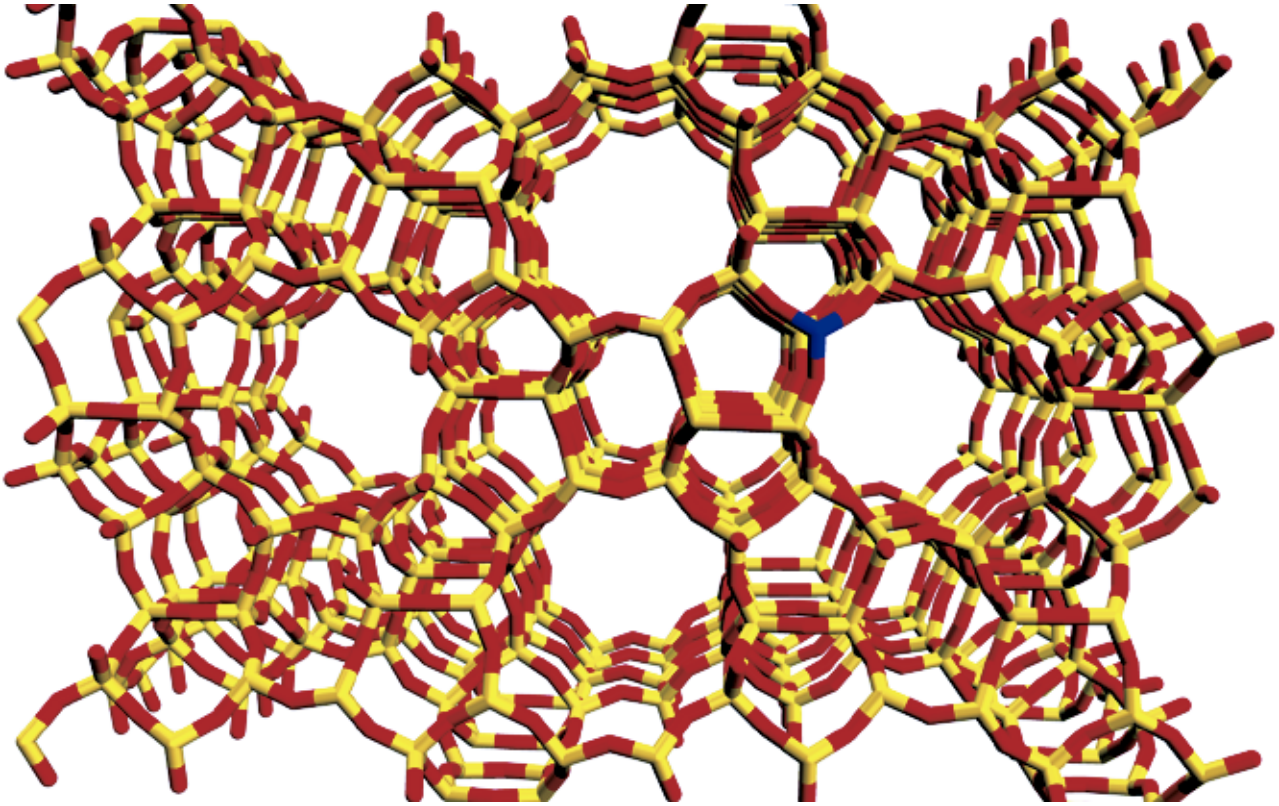

HPCx: More science per second

The optimised structure of TS-1, a microporous titanium silicate, which is employed in the chemical industry as a catalyst for partial oxidation reactions under mild conditions. See article on p12.

CONTENTS

Editorial

Ian Bush, HPCx Terascaling team

It gives me great pleasure to welcome you all to the seventh edition of the newsletter for the HPCx community, *Capability Computing.* As ever the service moves on. As reported in the last issue the hardware has again been upgraded, moving from a POWER4-based machine to POWER5. This has a number of benefits, both to you, the user, and us, the service provider, and three articles in this edition focus on the upgrade and the new possibilities afforded.

However the main focus of the service should not be the service itself, but the science it allows you to perform. It is amazing the diversity of length scales that are studied on modern supercomputers, and this issue includes articles describing calculations on systems that range from the smallest of molecular species to large-scale climate systems, managing to fit in the heart on the way!

I hope you enjoy this edition, and I am sure that with the recent upgrade and Phase3 not far away the HPCx community has an interesting few years ahead.

Less is more – moving HPCx to POWER5

Steve Andrews, HPCx Operations and Systems Group

As forecast in the last issue of *Capability Computing*, November 2005 saw the interim upgrade of HPCx from POWER4+ to POWER5 technology to become the Phase 2A system. At the operational level the success of this transition was due to the hard work and close collaboration between the HPCx systems group, applications team and IBM.

The move to POWER5 keeps us in step with IBM's development of this product line, so that we can get the benefit of their software development. It also means that when we come to upgrade the service to Phase 3 towards the end of the year, adding more nodes to increase the system's performance to 12Tflops will be comparatively simple. The new machine offers several advantages:

- at least the same performance (on Linpack) as the POWER4+ system
- early sight of what will be the Phase 3 machine later this year (see below)
- a doubling of the memory to 2Gb per processor
- a footprint reduction of 75 per cent.

Plans for this transition were first proposed last May, culminating in the agreed 48-hour interruption of service in November. The hardware was delivered on 3rd October to Daresbury Laboratory and the system was built alongside the existing Phase 2 machine. This enabled us to perform all the necessary pre-work on the new system without affecting the user service:

- Cabling
- Power-up
- System set-up
- Interim GPFS – since it was necessary to re-use the Phase 2 disk subsystem a temporary independent filesystem was created to allow configuration of the application packages.
- LoadLeveler – reconfigured and tested for the new 16-cpu lpar layout
- Benchmarking

The Phase 2 service closed on 23rd November to enable us to:

- Freeze the GPFS filesystems and perform a full backup
- Perform the GPFS cutover and validation. This involved removing the filesystem whilst leaving the data on the disks, reconnecting these to the new Phase 2A system and installing a new version of GPFS on top. All the data wonderfully reappeared.

The standard measure of establishing performance on such machines is the Linpack benchmark. Phase 2A achieved an Rmax figure of 7.395Tflops compared with 6.188Tflops from the POWER4 Phase 2 system and well ahead of the contractual requirement of 6Tflops.

Ten minutes early, at 0950 on 25th November, user service resumed on the new machine. Phase 2 had won the reputation of being the most reliable large HPC cluster in the world and although teething troubles are to be expected, we intend to repeat this achievement with the current system and Phase3 later this year.

So what is new? From a systems viewpoint, operation of the new hardware is considerably different from the old p690-based service, requiring considerable changes to be made to our procedures, monitoring and reporting of the system status. In addition to the two login and I/O server frames there are 12 compute frames each containing eight p575 servers. Each of these servers in turn contains eight dual-core POWER5 1.5GHz processors giving a total of 1536 cpus. Details of the architecture and initial performance results are in an HPCx Technical Report: www.hpcx.ac.uk/research/hpc/technical_reports/HPCxTR0602.pdf

The combination of AIX5.3 and POWER5 offers the potential for running jobs in Simultaneous Multithreaded (SMT) mode where the 16 physical processors on a node can masquerade as 32 virtual ones. For a technical description see: www.research.ibm.com/journal/rd/494/mathis.html.
We have seen good results for a number of applications, especially those that saw little benefit from Phase 2A such as molecular dynamics codes, and intend to implement this as a user-controlled option in the next few weeks.

November 2006 will see the opening of the 12Tflops Phase 3 system and planning for this began in February to ensure that the next transition will be as smooth and unobtrusive as the last.

Application performance on POWER5

Ian Bush, Mike Ashworth, Martyn Guest, Joachim Hein, Martin Plummer, Fiona Reid, Lorna Smith and Andy Sunderland, HPCx Terascaling team

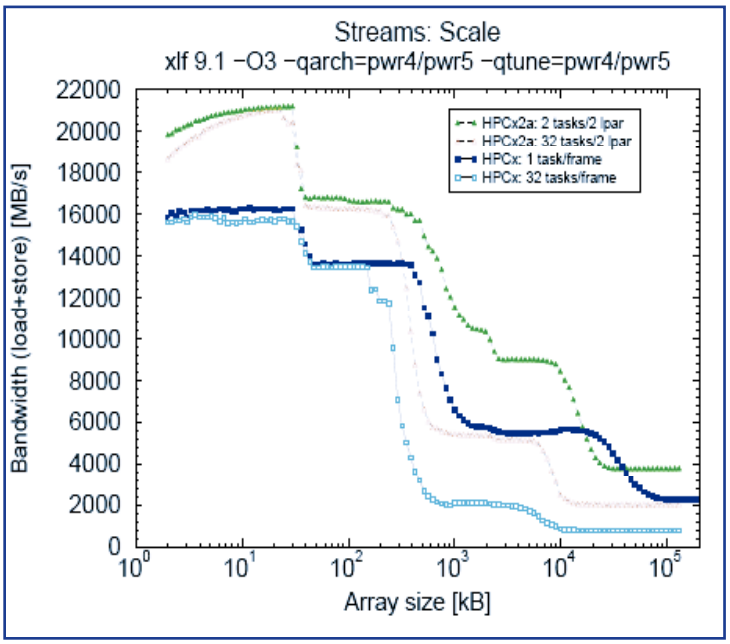

Figure 1.

As reported in the article opposite, the upgrade to Phase 2A last November has bought a number of benefits, not least the doubling of the memory available to applications. However, how do applications perform on the new system? Does the decrease in the clock rate from 1.7 to 1.5GHz impact adversely on the performance? Or does the new memory subsystem enhance it? Do you get more science per second?

The Linpack benchmark is around 1.2 times quicker on the new system compared to the old. Though this is a synthetic benchmark it does depend strongly on the cache architecture of the machine in question, and so this suggests that the improvements in memory subsystem of POWER5 compared to POWER4 could outweigh the small decrease in clock frequency. This is supported by the Streams benchmark (figure 1) where it can be seen that for nearly all sizes of dataset the memory bandwidth for the new system is markedly superior to the old.

For real applications the observation that the improvement in the memory subsystem is more important than the small decrease in clock frequency still holds in the majority of cases. Both applications that have substantial data re-use, such as CASTEP, CPMD, H2MOL and AIMPRO, and those that do not, e.g. PCHAN and ANGUS, are, in general, quicker, sometimes substantially so. For instance figure 2 shows the improvement for PCHAN, the IBM p690+ being Phase 2 and p5-575 being Phase 2A. While this is an extreme case, improvements of slightly bigger than the Linpack figure of 1.2 being more typical, it does show what the new system is capable of.

So it's all good news? Just submit it on the new machine and get scientifically rich quick? Of the applications that the terascaling team have tried there seems to be one class of application that either shows no improvement on the new architecture, or in some cases runs slower. That class of application is classical molecular dynamics (MD). DL_POLY, NAMD, AMBER and LAMMPS all show little or no improvement on the new architecture. For instance figure 3 shows the performance of DL_POLY for a simulation of 216,000 ions of Sodium Chloride.

It can be seen that unlike PCHAN there is little or no improvement, and the reason for this is, at present, far from clear. However, even here there is a silver lining to the cloud – symmetric multithreading (mentioned in the previous article) seems to be particularly effective for classical MD codes, and it seems likely that this will bring this class of application into line with the others.

So, in summary, many applications on the new machine run quicker despite the small drop in clock speed, which is, of course, good news for the user. There is one important exception, however: classical molecular dynamics. However, it does seem that once symmetric multithreading is available to the user even this class of application will run faster on Phase 2A.

Left: Figure 2.
Below: Figure 3.

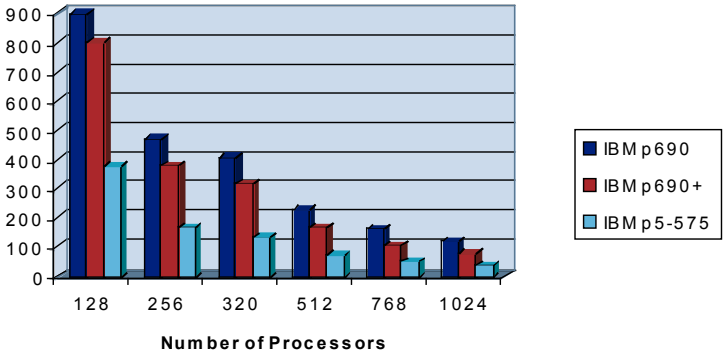

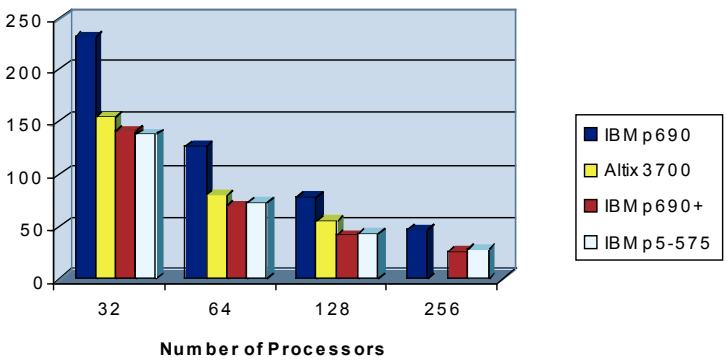

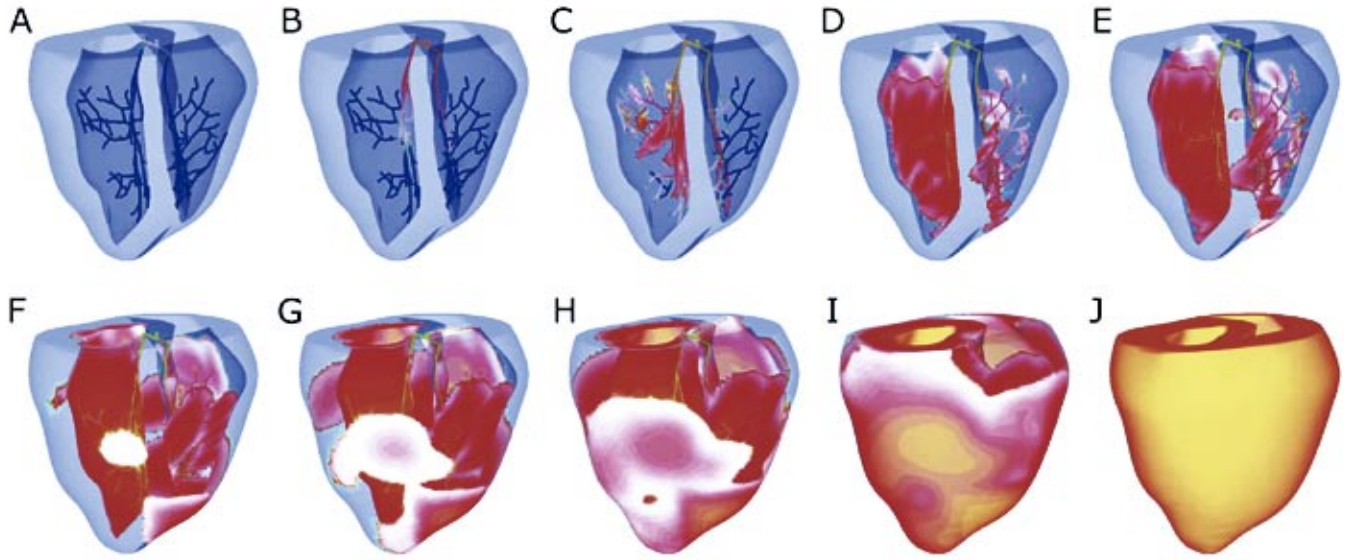

Figure 1. Normal Activation Sequence of the Ventricles.
(A-B): The cardiac impulse propagates from the atria (smaller chambers sitting on top of the ventricles, not shown) via the specialized conduction system into the ventricles. The specialized conduction system starts as a single strand on top (A), bifurcates into a left and right branch (B) and then further bifurcates into a mesh-like structure. The impulse reaches the left ventricular endocardium (the inner surface forming the left cavity) first (C), then the right ventricular endocardium (D-E). Subsequently, activation occurs within the heart's walls until breakthroughs are observed at the epicardium (F). Seen from the epicardial surface, breakthrough sites act as focal activations which join to form an epicardial wavefront (G-H). Finally, the entire ventricles are in excited state (I-J).

Integrative biology: modelling heart attacks with supercomputers

Gernot Plank (University of Graz), Richard Clayton (University of Sheffield),
David Boyd (University of Oxford/CCLRC), Ed Vigmond (University of Calgary)

Trying to understand what kills one third of the UK population is a challenge that needs the most powerful supercomputers available.

Heart disease and sudden cardiac death

Although the incidence of cardiovascular disease continues to fall, it remains an important cause of death in the industrialised world. Recent data for England and Wales for 2004 [1] indicate that diseases of the circulatory system accounted for around 190,000 deaths, 37% of the total, with almost half of these being directly attributed to ischaemic heart disease. In many of these cases, the lethal event is likely to have been a catastrophic arrhythmia called ventricular fibrillation (VF). Despite its immense social and economic impact, the mechanisms underlying formation of arrhythmias (arrhythmogenesis) are still debated and remedial therapies are, at best, suboptimal. In the case of life-threatening arrhythmias like VF, electrical defibrillation, the application of an electric shock within minutes of VF onset, is now recognized as the only effective therapeutic option.

The heart is an electromechanical pump. During each beat a wave of electrical excitation originates in the heart's pacemaker, the sinus node, located in the right atria (smaller chambers sitting on top of the ventricles). During a normal beat, the cardiac impulse reaches the ventricle via the only electrically excitable connection, the AV node, and subsequently is carried quickly to the working myocardium, acting as a signal for the muscle to contract (figure 1). During VF, a state of electrical anarchy exists, with normal electrical activity being replaced by rotating waves which stop the heart beating properly. VF is usually associated with underlying heart disease, but the mechanisms that initiate a particular episode of VF are poorly understood.

Despite the critical role that defibrillation therapy plays in saving human life, our understanding of the mechanisms by which electric shocks halt life-threatening arrhythmias remains incomplete. While recent advances in experimental methodology have

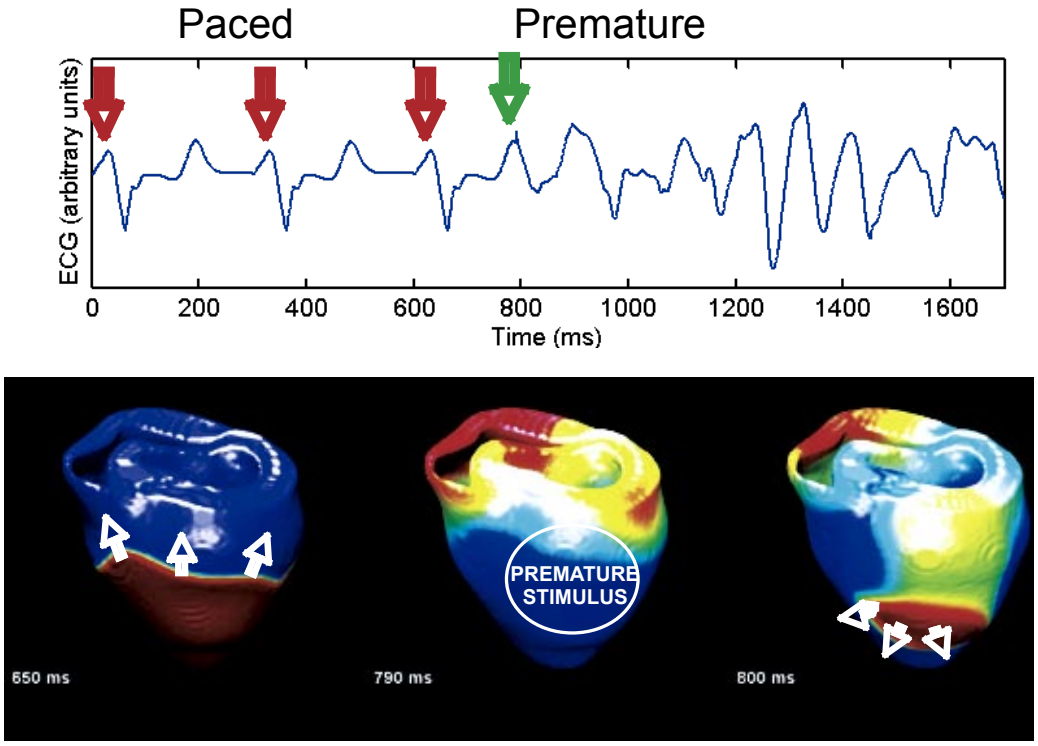

Figure 2. Simulation of the initiation of fibrillation in the ventricles
Top panel: simulated electrocardiogram (ECG). Red arrows indicate three stimuli to the apex (bottom) of the heart that result in normal placed beats at intervals of 300ms. Green arrow indicates premature stimulus delivered to the heart wall. The premature stimulus results in the onset of fibrillation, shown by rapid and self-sustained activity in the simulated ECG. Bottom panel shows snapshots of the simulation, where electrical activation is colour coded with blue indicating resting tissue and red indicating active tissue. The first frame (650 ms) shows the propagation of the third paced beat shortly after the stimulus has been applied to the apex (bottom) of the heart. The second frame (790 ms) shows the state of the heart just before the premature stimulus is applied over the region shown. Part of this region has yet to recover, shown by the yellow and light blue regions. The third frame shows the effect of the premature stimulus. The activity resulting from the stimulus can only propagate outwards and downwards because it cannot propagate into areas that are still recovering.

provided new characterizations of tissue responses to electrical shocks, inquiry into the success and failure of defibrillation is hampered by the inability of current experimental techniques to resolve, with sufficient accuracy, electrical behaviour confined to the depth of the ventricles during and after the shock. Realistic computational models of the heart, sometimes referred to as 'virtual heart simulators', are currently the only viable approach to allowing us to observe all parameters of interest at sufficient spatial and temporal resolution.

The Integrative Biology project

Integrative Biology [2] is a second generation e-Science project funded by EPSRC which is developing a Grid-based framework for computational biology driven primarily by the need to understand the two largest killers in the UK population, heart disease and cancer. The project is supporting a globally-distributed team of scientists developing multi-scale models which attempt to link behaviour at the molecular and cellular level to the operation of whole organs, enabling us to build realistic virtual heart simulators.

Our aspiration is that, as our models become more detailed, they will enable us to understand these diseases sufficiently well for us to develop therapeutic strategies for treating and possibly in the long term preventing them. The current state of the art in modelling the heart is more advanced than for cancer tumours, although developments there are moving rapidly, and our attempts to develop more detailed heart models are driving development of the computing infrastructure and will stretch the capability of current generation supercomputers such as HPCx.

The Integrative Biology team which we have assembled to tackle this challenge brings in expertise from many disciplines. Scientists from Oxford, Nottingham, Birmingham, Sheffield and Auckland are working with computational scientists from CCLRC, Oxford, UCL, Leeds, Auckland and IBM. Since the project started, partners have joined from Graz (Austria), Utrecht (Netherlands), Tulane, UCSD and UCLA (USA) and Calgary (Canada), and we expect the team to continue to grow.

The computing services we are developing within the project allow researchers to target simulations at the most appropriate computer system depending on the resources and response needed. They provide data and metadata management facilities, using the Storage Resource Broker [3], for looking after the many datasets created in computational experiments, and visualisation tools for examining results and discussing these collaboratively within the consortium. An associated project also being carried out by the Integrative Biology team, funded by the Joint Information Systems Committee, is developing a Virtual Research Environment which is embedding these services into a portal-based framework so they are easily usable by researchers who are not themselves computing experts.

Modelling the heart

VF can be studied experimentally using voltage sensitive fluorescent dyes or electrodes to map the spread of electrical activation in heart tissue but, as we noted above, this approach only yields information about activity on the heart surface. Modelling has an important role, because models of VF in the heart can be used to simulate and study electrical activation patterns that are consistent with what can be observed on the surface of the heart. This approach is similar in principle to studying the sun; it is only possible to observe the sun's surface, yet an understanding of the physics of stellar nuclear reactions allows the mechanisms that produce these surface observations to be understood.

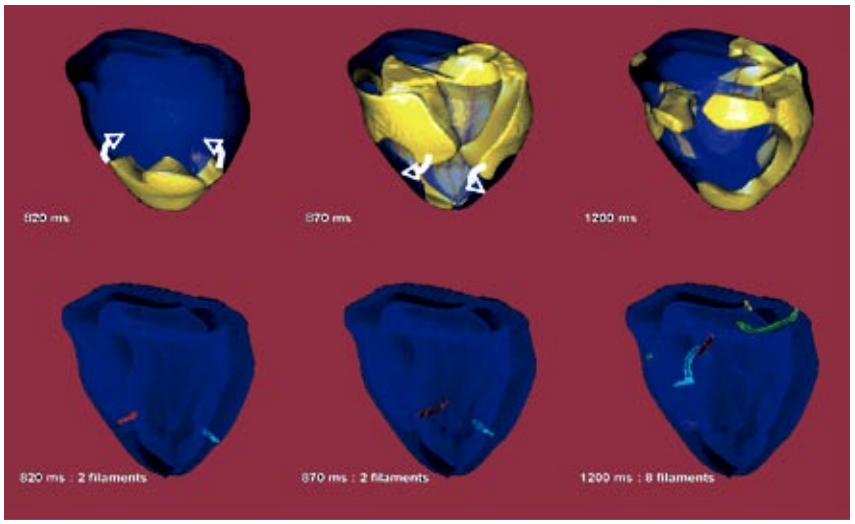

Figure 3. Simulation showing the development of fibrillation in the ventricles. Top panel shows isosurface views, where electrically active regions are enclosed by yellow surfaces. The first frame (820 ms) shows how the downward propagating activity shown in Figure 2 is beginning to curl around as the tissue recovers from the paced beat, and forms a pair of counter-rotating scroll waves. In the second frame (870 ms) these scroll waves have rotated by about 180 degrees, but the activation pattern is unstable, and by the third frame (1200 ms) the initial scroll wave pair has broken up into multiple interacting waves. Bottom panel shows scroll wave filaments – the lines around which the scroll waves rotate. In the first two frames there are two filaments, one for each of the scroll waves. In the third panel (1000 ms) there are 8 filaments, reflecting the more complex activation pattern resulting from the initial instability.

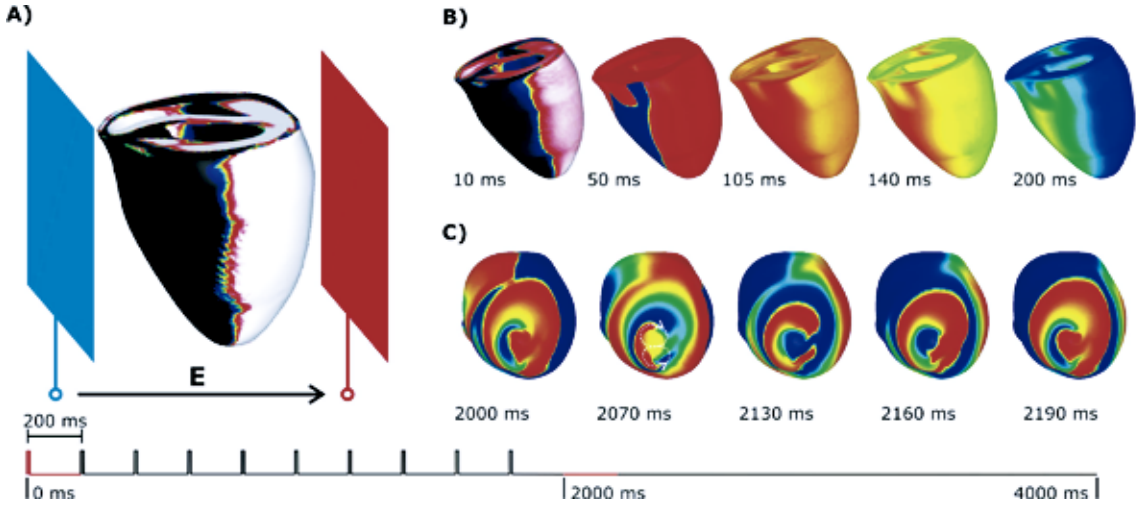

Figure 4. A virtual experiment. Setup for a 'virtual experiment' to induce an arrhythmia in the ventricles by applying an electrical pacing protocol. A) The ventricles are immersed in a conductive fluid and place between two plate electrodes. B) Electrical activation of the ventricles during the stimulation period. C) After the last pacing pulse, a so-called figure-of-eight re-entry (named after the movement of the tips of the wavefronts) ensued and was sustained until the end of the simulation run at 4000 ms.

Models of electrical activity in the heart during VF vary widely in the level of detail that is included. To reduce the computing resources required, pilot simulations often use simplified models of anatomy, which may be a 2D sheet of tissue or a 3D slab representing part of the heart wall, and use simplified representations of the electrical excitability of cardiac cells. However, to properly understand the complexity of a problem like VF, we would ideally like to model the heart as a complete organ.

Mathematically, the virtual heart is based on a set of coupled PDEs which are referred to as bidomain equations. In recent years, most cardiac computer models have used either a computationally significantly less demanding simplification (monodomain) which does not capture electrophysiologically important mechanisms, particularly when defibrillation is under study, or rather coarse discretisation to keep simulations computationally tractable.

State-of-the-art models are a trade-off. They either capture gross anatomy at limited spatial resolution or represent regions of limited size incorporating microscopical detail. Several questions of major electrophysiological importance, however, require both macro- and microscopic structural details and use of the more complete bidomain formulation. Today, the cardiac modelling community is on the brink of implementing simulators which will allow development of virtual hearts for many different species. The use of high performance computing and state of the art numerical methods to deal with the tremendous computational burden imposed by such simulations is key to success in this endeavour.

Understanding ventricular fibrillation

Current research at Sheffield University supported by the Integrative Biology project is focusing on understanding the mechanisms that initiate and sustain VF. Compute-intensive simulations using whole ventricle detailed anatomy and biophysically detailed models of electrical excitability are run on HPCx, and these build on simulations using simplified models that are run on local HPC resources including the White Rose Grid [4]. A key aspect of this work is to relate the findings to clinical practice. Work on modelling the initiation of VF has already yielded information that could be used to identify patients at risk, and this is the basis of a pilot clinical study about to start in collaboration with clinical colleagues at the Northern General Hospital in Sheffield. Work on the mechanisms that sustain VF is also tightly meshed with clinical and experimental studies, and is one component of a wider project focusing on understanding VF in the human heart that also involves the Universities of Auckland, Utrecht, Oxford, and UCL.

The current version of SCAM (the Sheffield Cardiac Arrhythmia Model) is written in C, uses shared-memory parallelism, and runs on a single frame of HPCx. The code has been optimised for the HPCx architecture, and ongoing development aims to further exploit the mixed-mode parallelism capability of HPCx, so that

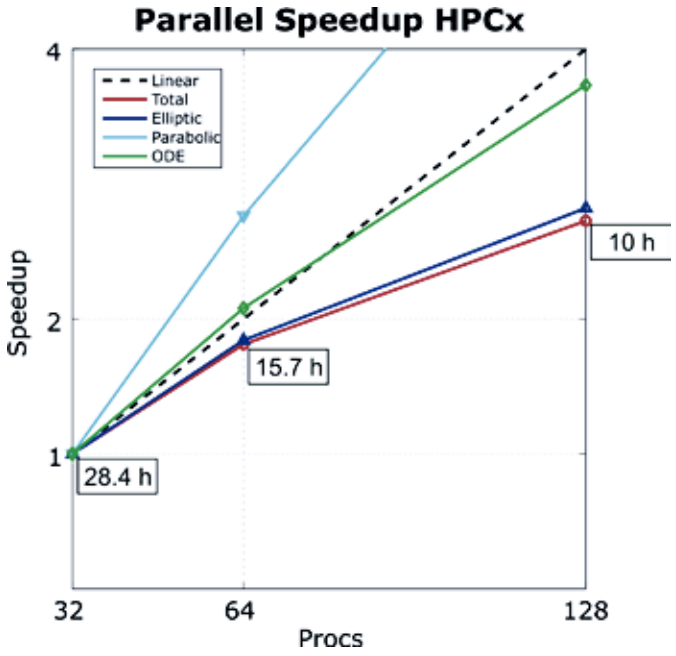

Figure 5. Benchmark results: scaling of different portions of the bidomain computation. As expected, the ODE part scaled linearly (no communication required).The parabolic problem, solved by a simple forward Euler step, basically involved only a matrix-vector product which showed super-linear scaling. Computations were dominated by the elliptic problem which scaled reasonably well, particularly if one takes into account that the problem size is small for the number of CPU's employed in these simulations. Overall execution time in hours is shown as a function of the number of CPUs.

simulations can be run across large numbers of frames. Results of simulations using SCAM show initiation and development of fibrillation in the ventricles (Figures 2 and 3).

Virtual experiments on HPCx

In a preliminary study at the University of Graz using the Integrative Biology framework, the feasibility of carrying out a 'virtual experiment' was tested using HPCx. A computer model, discretised at an average spatial resolution of 200 μm, simulated a ventricle which was immersed in a conductive bath.

At the bath boundaries, two plate electrodes were placed next to the anterior and posterior faces of the ventricle (figure 4A).To test conditions under which an arrhythmia can be induced, a train of 10 pacing pulses of varying basic cycle length was delivered (figure 4B). After the last pacing pulse, 2 seconds of activity were simulated to examine whether an induced arrhythmia was sustained or self-terminated (figure 4C).

Performing this virtual experiment involved the solution of an elliptic PDE (862,515 unknowns), a parabolic PDE (547,680 unknowns) and a set of 21 non-linear ODE's, defined at the same grid as the parabolic PDE. Using a temporal discretisation step of 8μs, the solution scheme had to be repeated 500,000 times to complete the experiment. Preliminary simulations carried out on a Dual Opteron desktop computer suggested that execution times would be around 2 months. Using 128 CPUs of HPCx allowed the execution of a single experiment in only 10 hours. The simulations used CARP, Cardiac Arrhythmia Research Package.

A subset of this simulation was repeated using different numbers of CPUs to demonstrate the scalability of the method (figure 5). The overall computational workload is clearly dominated by the elliptic problem (> 95 per cent of the overall workload). The parabolic PDE, solved by a simple forward Euler integration step, showed super-linear scaling. As expected, the ODE solver scaled linearly since the involved variables do not diffuse and thus no communication is required. The dominating elliptic problem scaled well, although the parallel efficiency decreased slightly when going from 64 to128 CPUs. This problem is rather small for the high number of CPUs, so the scaling efficiency is more than satisfying.

These preliminary results suggest that realistic simulations of a human heart including a torso are feasible on the HPCx platform. In such simulations one has to deal with roughly 20–200 million unknowns (20–200 times larger than in this study). Memory usage and execution times will require the use of more CPUs. It is expected that parallel efficiency will increase significantly thanks to a more favourable ratio between local computational load and communication.

The future – model-based therapy?

We have demonstrated it is feasible to carry out virtual experiments using supercomputers. This leads us to expect that different therapeutic approaches may be studied to investigate possible optimizations. Three avenues of research are particularly interesting:

1) How can the dynamic behaviour of the cardiac membrane be modified to terminate or even prevent arrhythmia? This could lead to development of pharmacological treatments to determine which pharmacological targets are the most promising ones.

2) Can the most effective therapy, namely implantation of a cardioverter/defibrillator, be further optimised? Is it possible to defibrillate in a smarter way to increase the lifetime of the implanted device, to reduce the time to therapy in the case of failed shocks and to minimise tissue damage? For instance, can it be tested how the safety and efficacy of the therapy are affected by geometry and location of the defibrillation electrode, by the instant of shock delivery (relative to the cycle of the arrhythmia) and by the waveform of the delivered shock pulse?

3) The only curative treatment to cardiac arrhythmias is catheter ablation, a procedure by which the conduction pathways within the heart are modified. This treatment is particularly effective in the atria, but can we also find applications in the ventricles (the main pumping chambers)?

Future computational models of the heart have to incorporate all these ingredients, the atria, the ventricles, the specialised conduction system and probably a torso model too to allow direct comparisons with clinically measured parameters like the body surface electrocardiogram (ECG). By combining electrical models

Continued on p9.

A quick and easy way to make your own...

Kevin Stratford, EPCC

New materials by computer simulation

Computer as kitchen

Anyone who has ventured into the kitchen and tried their hand at making a cake will know that the final result can be critically dependent not only on the mixture of initial ingredients, but also on the baking process itself – the exact cycle of heating and cooling to which the mixture is exposed.

The heating and cooling is an example of what the physicists would refer to as a non-equilibrium process (referring to the changing temperature in this context). Such non-equilibrium processes can have important consequences for many manufacturing processes, as it is an entire process history which determines the properties of the final product. This has important consequences which affect flow behaviour and shelf-life of products such as yoghurt, vaccines and paint.

Generic technologies which can create new materials using non-equilibrium processes are therefore of great interest. While laboratory experiments with novel systems can be time-consuming and expensive, with no guarantee of success, computer simulation offers a powerful tool for cooking up new materials and mixtures in relative comfort (and without the washing-up).

Jam today

One particular example of the discovery, via computer simulation, of a new type of material has recently been published in *Science*. It is well known that some pairs of liquids (eg oil and alcohol) do not mix at room temperature, but do mix if heated. The opposite effect, cooling an initially well-mixed pair, often results in droplets of one fluid forming in the other. However, in certain conditions, one sees 'spinodal' decomposition, where continuous domains of each fluid which together form a 'bicontinuous' structure. Here, continuous domains of both phases form and grow with time, typically with complex topological reorganisations and pinch-offs of narrow fluid necks. The characteristic domain size increases in a well-understood way until it is ultimately limited by the fluid container. (In the laboratory, spinodal decomposition is often initiated via a sudden temperate quench, for example, by dipping a small sample of two liquids miscible at room temperature into liquid nitrogen for a few seconds.)

Computer simulations have been used to address the question of what happens if small solid particles are added into the mix. Such colloidal particles (which may be a few nanometres to a few microns in size) can be chosen so that they have a strong preference to be at the interface between the two fluids. In fact, this can be so favourable energetically that, once at the interface, the particles will not be removed by thermal motions.

Particles introduced into the initial miscible mixture will diffuse randomly. However, if spinodal decomposition is initiated, the particles will be captured by newly forming interface. Initially, the interfacial area is large, and the corresponding density of particles per unit area will be small. However, as the domain size increases and the interface shrinks the particles, which cannot escape, are forced together. The simulations show that the particles eventually become jammed together at the interface and so arrest the growth of the fluid domain size. (This process is best appreciated by looking at an animation; see the links below.)

The image shows a small section of a three-dimensional simulation in which solid particles (coloured green) have become jammed at the interface between two immiscible fluids (the interface is coloured differently on each side to represent the different fluids).

It is this jamming of particles into what is apparently a hard glassy state which forms a scaffolding for the two continuous fluid domains which is particularly interesting. The resulting structure appears to be essentially solid (a gel), and may have novel properties and uses. For example, immiscible liquids between which a favourable chemical reaction occurs could be pumped through such a structure, and brought into continuous contact (perhaps with the addition of catalytic particles).

Numerical Recipe

These computations are demanding as there are not only many important elements of physics to get right, but they must also be carried out on a large enough scale that finite size artifacts are excluded. The simulation must include full hydrodynamics for the two-component fluid, the hydrodynamic force between fluid and the moving particles, the thermodynamics of wetting at the solid-fluid interface of the particles, and thermal (Brownian) motion of the particles. One tractable method which allows such simulations is lattice Boltzmann (LB), which is used here.

In common with other numerical methods, LB represents the flow domain as a regular discrete lattice. However, it also uses a discrete velocity space. The fluid is then represented by a distribution function at each lattice site which can be thought of as the density of fictitious fluid particles at the site with a given discrete velocity. The velocities are chosen so that the particles from one lattice site move to a neighbour in one discrete time step. It can be shown that the model fluid obeys the Navier-Stokes equations which govern fluid flow.

For a two-component fluid, a second distribution can be introduced which represents the relative composition at each lattice site. Appropriate physics can be introduced by combining

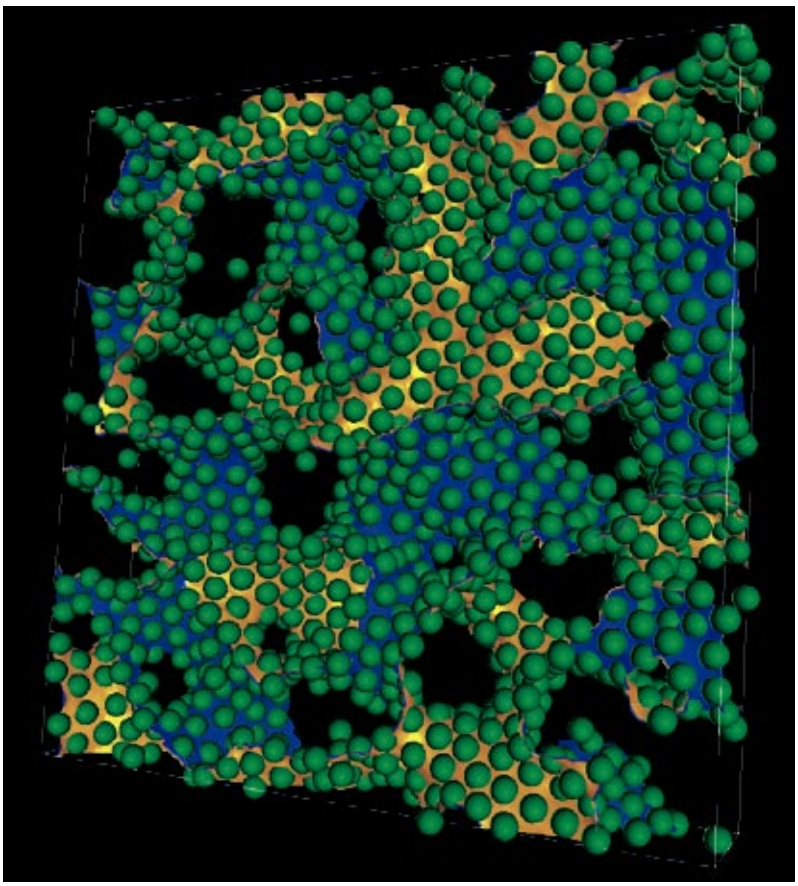

The image shows a small section of a three-dimensional simulation in which solid particles (coloured green) have become jammed at the interface between two immiscible fluids (the interface is coloured differently on each side to represent the different fluids).

information from the two distributions at each lattice site at each time step. This method has a number of advantages: complex topological structure evolves naturally on the lattice, removing the need for complex interface tracking procedures, and continuity of the velocity field at the fluid-fluid interface is automatically satisfied. In addition, in contrast to many other computational fluid dynamic methods, the pressure calculation is entirely local, meaning that parallelisation is almost trivial.

Finally, solid-fluid boundary conditions for moving objects may be introduced by a process known as bounce-back on links. Fluid distributions arriving at a solid-fluid interface can be reflected back into the fluid with a known momentum transfer between solid and fluid. In this way the total hydrodynamic force on the particle from the fluid can be computed (and vice-versa). The contact angle between the fluid-fluid interface and the solid (ie the thermodynamics) can be controlled using a similar procedure for the second, composition, distribution. While the particles add considerable complication vis-a-vis the fluid-only problem they can, remarkably, be implemented with very little impact on overall performance.

Don't try this at home

In all, these calculations have shown how computer simulations can lead the search for new materials. Our experimental colleagues are currently involved in the messy business of trying to make the stuff in the laboratory. But that's part of the fun.

The simulations were performed on HPCx and at Edinburgh Parallel Computing Centre. The work was funded by the EPSRC RealityGrid project (GR/R67699).

References

K. Stratford, R. Adhikari, I. Pagonabarraga, J.-C. Desplat, and M.E. Cates, Colloidal Jamming at Interfaces: A route to fluid-bicontinuous gels, Science, 309 2198 (2005): www.sciencemag.org/cgi/content/full/309/5744/2198.

See also Phillippe Poulin's 'Perspectives' article in the same edition (page 2174): www.sciencemag.org/cgi/content/full/309/5744/2174

Integrative biology continued

with mechanical models we could simulate the mechanical action of the heart and the blood flow driven by contraction and so allow comparison with further clinically important parameters such as cardiac output and ejection fractions. However, integrating these parts into one comprehensive model will substantially increase the computational burden. Parallel environments like HPCx are the only viable way of executing such simulations by providing the required address space and keeping execution times tractable.

In conclusion

The Integrative Biology project is building a computational framework which will enable scientists to develop increasingly realistic virtual heart simulators without having to become computer experts themselves. The potential benefit to the population as a whole from using simulation to help us understand, treat, and possibly prevent, heart disease is a strong motivation driving this work forward.

Acknowledgements

We wish to thank other members of the Integrative Biology consortium for their input into, and comments on, the work described in this paper, in particular the role of the scientific collaborators in determining the requirements for the project. We also thank Michael Holden at EPCC for help with installing and optimising our codes on HPCx. We acknowledge the financial support of the EPSRC (ref no: GR/S72023/01) and IBM for the Integrative Biology project, and of the Joint Information Systems Committee for the Integrative Biology Virtual Research Environment project. R. Clayton acknowledges support from the British Heart Foundation (PG03/102/15852). G. Plank was supported by the Austrian Science Fund FWF (R21-N04).

References

[1] www.statistics.gov.uk/

[2] www.integrativebiology.ac.uk

[3] www.sdsc.edu/srb/

[4] www.wrgrid.org.uk

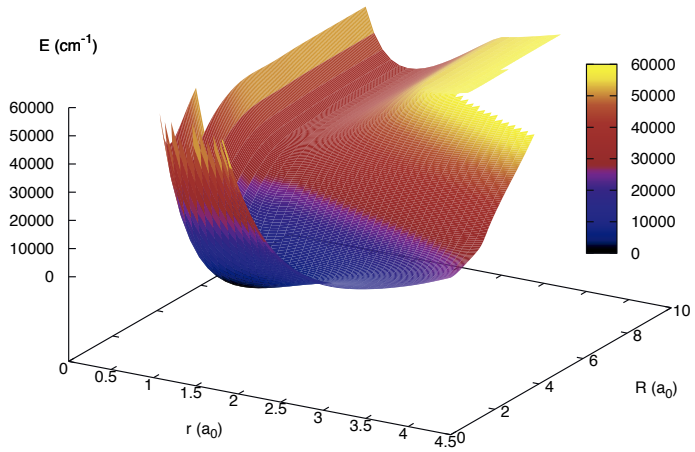

Figure 1: A 2D cut of the potential energy surface of H_3^+ [4].

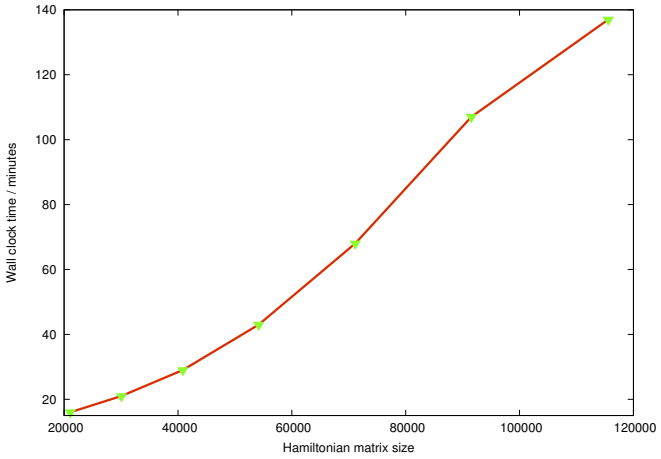

Figure 2: Timings for calculations on D_2H^+ performed on 768 processors for a number of different Hamiltonian sizes.

Capability-enabled calculations on H_3^+ and D_2H^+

James Munro, Bruno C. Silva, Paolo Barletta and Jonathan Tennyson, University College London

H_3^+ is the simplest polyatomic molecule: it is composed of just three protons and two electrons. H_3^+ is usually formed via the reaction

$$H_2 + H_2^+ \rightarrow H_3^+ + H$$

H_3^+ is therefore fundamental to any chemistry involving hydrogen, the most common molecule in the universe, and it has been found to exist in many extreme environments. For instance the discovery of H_3^+ in the atmospheres of Jupiter, Saturn and Uranus has lead to detailed studies of unusual atmospheric and magnetic phenomena, in which it plays a central role, and it is also expected to exist in low temperature stars. It is amazing that such a simple molecule is not yet been fully understood.

In an attempt to fully understand the properties of H_3^+, Carrington, Buttenshaw and Kennedy recorded the breakup of H_3^+ in an infrared laser [1]. Remarkably, the fragments from the breakup often had more energy than the photons from the laser, and the high-resolution detail demonstrated that an extraordinary number of quantum states were playing a part in the molecule's break-up.

After nearly 25 years of study, only the most basic theoretical aspects of this problem have been understood. The dynamical aspects of how the molecule breaks up arise from vibrational and rotational motion of its nuclei moving under the influence of the electric forces acting between the protons and electrons. The complete quantum description of this is not easy to visualize, but, classically, it can be seen as the nuclei behaving like marbles moving on the inner surface of a big, wavy, ridged bowl – the potential energy surface (figure 1). The potential energy surface binding the protons together is very deep and therefore supports a large number of quantum states, making calculations difficult.

Similar phenomena have also been recorded for H_2D^+, D_2H^+ and D_3^+, which are like H_3^+ but heavier due to the additional neutrons, resulting in more quantum states, making the problem even harder.

To solve such molecular motion problems, a computational package named DVR3D[2] has been developed that greatly accelerates the calculations by implementing what is called a Discrete Variable Representation (DVR). A DVR represents the solutions as points in a grid spanning the potential. Methods like this can solve a large range of problems in molecular physics. However, calculation of all the quantum states up to and including the break energy remains extremely demanding.

To solve the H_3^+ problem we use a technique common in quantum mechanics: the variational principle. This technique is much like a trial and error system; results are computed from a set of trial functions, designed to satisfy the conditions of the quantum mechanical problem. These functions are trial wave functions and they are constructed from combinations of basis functions, also called a basis set. It turns out that calculating the wave functions and their energy from given a basis set is an eigenproblem, the energies being the eigenvalues of the Hamiltonian matrix, and the wave functions the corresponding eigenvectors. These wave functions describe the probability of finding the molecule in a particular configuration. The more general or complete the basis set is, the more accurate the answer becomes. Ideally, one would use an infinite basis set to solve the problem perfectly. Practically, however, one has to limit the accuracy to the available computational power.

To enable these calculations, DVR3D has been parallelised, harnessing the capability of high performance machines like HPCx.

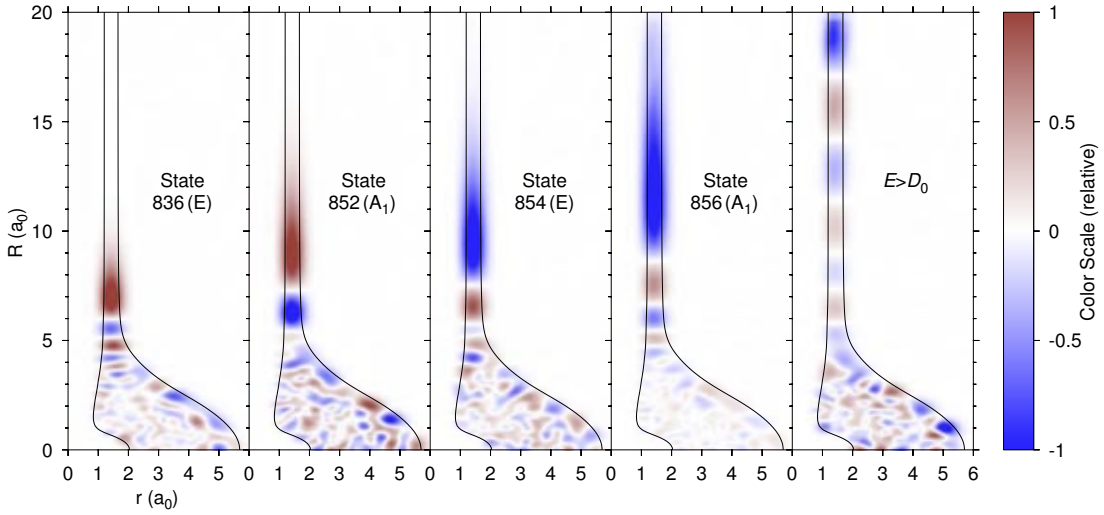

Figure 3: Four bound but long range states in H_3^+ and an eigenvector with an energy larger than the dissociation energy.

PDVR3D: A program for large basis set calculations of nuclear motion

We have implemented PDVR3D using MPI, BLACS, PBLAS and ScaLAPACK [6]. The best routine we found for diagonalising very large, dense matrices was ScaLAPACK's PDSYEVD. PDSYEVD is based upon a recently developed divide and conquer method: the matrix is split up into smaller matrices each of which are solved in a recursive manner. This is ideal for a parallel implementation and proves to be very efficient.

Our chief reason for using supercomputers such as HPCx is that the vast amounts of RAM available enable the construction and diagonalisation of these very large matrices. Our largest calculation (on D_2H^+ [3]) required over 153 Gigabytes for one copy of the matrix. To stress again how demanding this problem is, it required the diagonalisation of an order 143,541 matrix to find all the eigenvalues and vectors, and even routine runs require matrices of the order of tens of thousands. This calculation was only just possible on 1280 processors on HPCx phases 1 and 2.

A disadvantage in the variational approach is that a single calculation does not tell us how accurate the numbers are; several different calculations have to be performed, testing different basis sets. Figure 2 shows the wall clock time for a series of calculations to test the convergence of results for D_2H^+. The fact that the time is increasing almost linearly with matrix size in figure 2 means that PDSYEVD is becoming more efficient as the problem size increases; for a serial diagonalisation the time usually scales as the cube of the matrix size.

Long range bound states in H_3^+ and D_2H^+

The results from our calculations on H_3^+ and D_2H^+ were both surprising and intriguing. Near the energy at which the molecule breaks apart we find that there are a number of bound states confined only by the zero-point energy. This effect emerges in the region above the classical break-up energy (defined by the potential surface's asymptote), but below the quantum break-up energy, the zero-point energy of the H_2 fragment.

The most fascinating aspect of these states is that, classically, there is nothing binding them! The nuclei are bound because quantum mechanics will not allow any state of H_2 below its zero-point energy, therefore keeping H_3^+ bound but with extremely large inter-nuclear separations. Because these states have a behaviour which approaches that of the asymptote, they have been named asymptotic vibrational states (AVS) [5,3]. Figure 3 shows four such AVS. The existence and extent to which such states can span the exit channel has never been determined before.

These four states may not be the whole story though. It is expected that an even larger basis set calculation, with radial distances that extend beyond 20 a_0, will reveal more of these states. While it remains to be seen what effects these long range states have on the spectroscopy of H_3^+, it is clear that they will be important in its complete description.

Another very important aspect of our calculations is the fact we are able to use the eigenvectors corresponding to energies above dissociation. These eigenvectors resemble the wave functions for a free particle in the quantum sense ($E>D_0$ in figure 3). They can form the basis of further work as the starting point for low-energy time-dependant scattering calculations, enabling the study of basic chemical reactions such as

$$D_2+H^+ \rightarrow HD+ H^+$$

Of particular interest is the use of these eigenvectors for the calculation of resonances in the continuum of states above the dissociation energy, which will be of fundamental importance in solving the H_3^+ near dissociation spectrum problem.

It seems that these simple molecular systems will remain a computationally challenging problem in quantum mechanics, and that only with HPC systems like HPCx will we obtain the long sought answers for these fundamental scientific problems.

Acknowledgements

We thank EPSRC for providing the computer time on HPCx as part of the ChemReact Computing Consortium.

Bibliography

[1] A. Carrington and I. R. McNab Acc. Chem. Res.22 218 (1989).

[2] J. Tennyson, M. A. Kostin, P. Barletta, G. J. Harris, O. L. Polyansky, J. Ramanlal & N. F. Zobov Computer Phys. Comm. 163 85 (2004)

[3] P. Barletta, B. C. Silva, J. J. Munro, J. Tennyson Mol. Phys. Submitted (2006)

[4] O. L. Polyansky, R. Prosmiti, W. Klopper and J. Tennyson Mol. Phys. 98 261 (2000).

[5] J.J. Munro, J. Ramanlal and J. Tennyson New J. Phys. 7 196 (2005).

[6] www.netlib.org/scalapack/

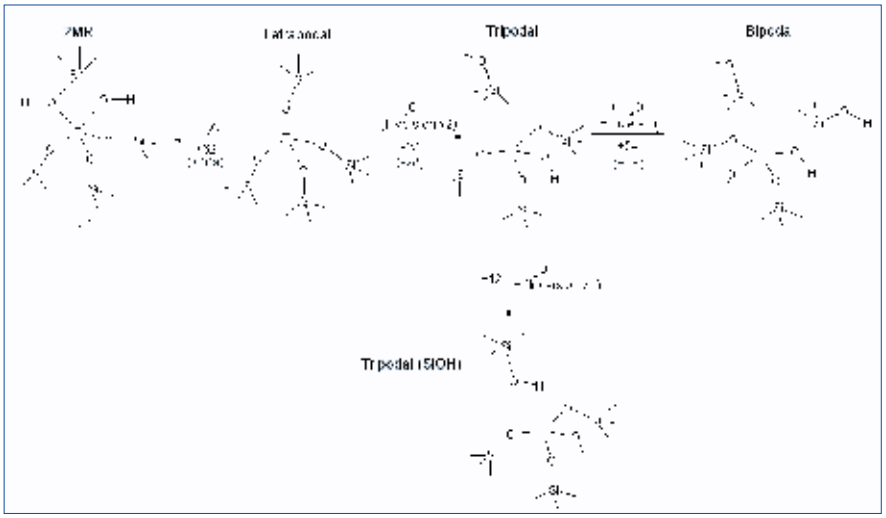

Figure 1. Hydrolysis reaction scheme for Si and Ti sites in TS-1. Calculated formation energies in [kJ mol^{-1}]. The energies indicated are calculated from Ti- and Si-centred clusters (Si values in parentheses).

QM/MM modelling of the TS-1 catalyst

Judy To, Alexey Sokol, Richard Catlow, Samuel French (all Davy Faraday Research Laboratory, the Royal Institution of Great Britain)
Paul Sherwood, Ian Bush, Huub J. J. van Dam and Martyn Guest (all CCLRC Daresbury Laboratory).

In this article we report how HPCx has been used to perform a series of coupled Quantum Mechanical/Molecular Mechanical (QM/MM) studies on the Titanium-substituted zeolite catalyst, TS-1. The periodic framework of the catalyst is shown in the cover image, but in our models we use a finite cluster, typically 2000 atoms, centred on the Ti atom. The QM/MM approach uses a quantum mechanical approach for the active site, coupled to a classical treatment of the environment. In this way quantum mechanical methods (here Density Functional Theory) are extended by the inclusion of the steric and electronic influence of the surroundings. We deployed the ChemShell software [1,2], which incorporates code from the GAMESS-UK[3] and DL_POLY[4] packages to perform the QM and MM parts of the calculation.

TS-1 has an outstanding ability to catalyse various oxidation reactions with aqueous H_2O_2 as the primary oxidant under mild conditions[5,6]. It is commonly believed that the isolated titanium atoms in the framework of TS-1 are the active sites for selective oxidation. However, the formation and structure of the active sites and the influence of water on the stability of the sites have not been definitively resolved. In our studies we have identified the hydrolysis and inversion of Ti-O-Si bridges as the key mechanism, which helps to stabilise Ti within the framework of zeolite. A number of the possible transformations are shown in figure 1, the tetrapodal species is that present in the un-hydrolysed framework structure (see cover image).

The formation energy of the Si ion in the different structural environments within (non-Ti substituted) silicalite indicates that the order of stability of these species decreases as:

tripodal > tetrapodal > bipodal > 2MR

whereas for the Ti silicate clusters, which includes the tripodal (SiOH) intermediate, the order is slightly different:

tripodal > tripodal (SiOH) > tetrapodal > bipodal > 2MR > titanyl.

In both sets of calculations, the 2MR cluster is markedly higher in energy than the other three clusters, probably due to the formation of bridging silanol and/or titanol groups, which repel each other. However, the repulsion of these hydroxyl groups, which are analogous to bridging Si-O(H)-Al groups in acidic zeolites, is released after the inversion of the tetrahedral sites. During the inversion process, the Si and/or Ti atoms (bonded to the hydroxyl groups) move away from their tetrahedral configuration to a new 'external' relaxed position. The hydroxyl groups follow the T-atoms, which breaks the two-membered ring and forms the more stable tripodal structure. The tripodal structure is lower in energy by 12 $kJmol^{-1}$ than the tripodal (SiOH). This result suggests that the silanol inversion process required to transform the tripodal (SiOH) to the tripodal structure can easily occur under experimental conditions, and thus can be considered as an intermediate step in the formation of the tripodal structure.

Based on these results we conclude that under anhydrous conditions a tetrapodal structure will probably constitute the ground state of the substitutional site. However, upon hydration, a dramatic change should take place with the tripodal structure being prevalent. This tripodal structure will be the active site of the TS-1 catalyst. A fuller report on this work has been published [7].

This study illustrates the type of complex chemical problem that can be tackled on the HPCx system. At the moment we can exploit 64–256 processors, and work is underway to extend the scalability further and also to provide access to more demanding quantum mechanical methods, such as correlated wavefunctions and treatments for excited states.

References

[1] P. Sherwood et al., J. Mol. Struct. (Theochem.) 2003, 632, 1. and refs therein.

[2] www.chemshell.org

[3] www.cfs.dl.ac.uk

[4] www.cse.clrc.ac.uk/msi/software/DL_POLY/index.shtml

[5] I. W. C. E. Arends, R. A. Sheldon, Appl. Catal. A 2001, 212, 175, B. Notari, Adv. Cat. 1996, 41, 253.

[6] J. To et al., Angew. Chem., 2006. 118, 1663

[7] J. To et al., .J. Mat. Chem. 2006 in press

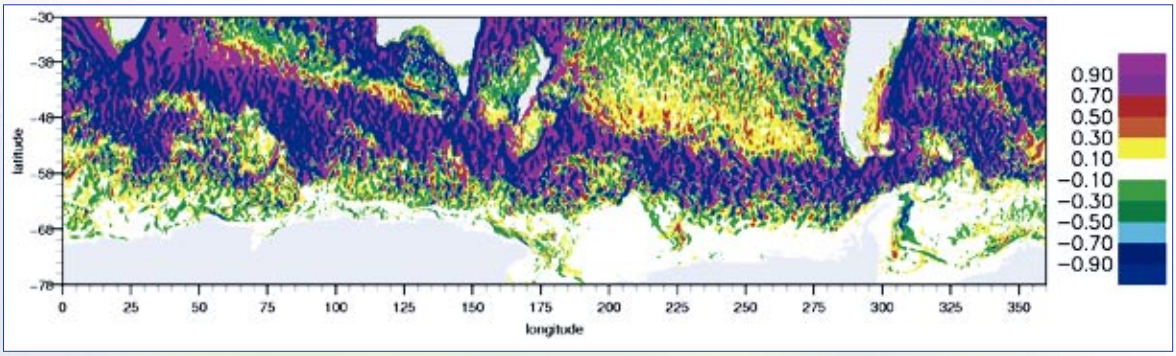

Figure 1 (right). Depth integrated northwards eddy heat transport (100GW per 1/12o x 1/12o gridcell).

Figure 2. Solid line: zonally and vertically integrated eddy northward heat transport (PW). Dotted line: zonally and vertically integrated total ocean heat transport. Dashed line: the heat loss to the atmosphere cumulative from Antarctic to the latitude.

Warming the extremities

Andrew Coward and Mei-Man Lee, National Oceanography Centre, Southampton

To maintain the Earth in near thermal equilibrium there must be a transfer of excess heat from the equator towards the poles. The atmosphere and ocean play roughly equal roles in the physical transport of this heat with the processes influenced by a variety of complex physical, chemical and biogeochemical interactions. In the oceans, the major manifestation of this transport is in the well-known Western Boundary Currents. These flow north in the Northern hemisphere (eg the Gulf Stream and Kuroshio Currents) and south in the Southern hemisphere (eg the Agulhas and East Australia Currents). However, Western Boundary Currents, which result from an intensification of the wind-driven circulation due to the rotation of our planet, can only exist where they have a western boundary to lean against. In the latitude band between the southern tip of South America and the northern tip of the Antarctic Peninsula there is no such western boundary in the top 2000m. This is the domain of the Antarctic Circumpolar Current (ACC), the greatest ocean current, which encircles Antarctica.

Alternative processes must be dominating the poleward transport of heat in this region and understanding those processes is vital to reducing uncertainty in climate change predictions. The oceans are turbulent and full of motions with a wide range of time and space scales. Time scales range from minute to minute variations in the surface wind stresses to thousands of years for the thermohaline ('conveyor belt') circulation. Space scales range from the molecular scale to thousands of kilometers. Turbulent flows are commonly modelled by expressing the turbulent properties in terms of mean values plus a fluctuating component. The fluctuating component will vanish in any averaging of linear expressions containing the component. However, quadratic and higher-order expressions of the fluctuating components cannot be eliminated and approximations (parameterisations) are often introduced based on experimental or observational evidence.

In the Southern Ocean, the mean flow does make some north-south deviations but is dominated by the ACC flowing from west to east. Poleward transport across this current is most likely carried out by the fluctuating components, in particular the correlation between the meridional (North-South) velocity fluctuations (v') and the temperature fluctuations (t'). This correlation is termed the eddy heat transport. The problem for oceanographers is that practical climate models cannot resolve all the time and space scales necessary to resolve the eddy heat transport but collecting observations with which to test parameterisations across the whole range is physically very difficult and prohibitively expensive too.

Facilities such as HPCx provide an alternative solution. Using 512 processors at a time we are able to integrate a global ocean model (OCCAM) which has a spatial resolution of approximately 10km. Furthermore, we can force the surface of the ocean with atmospheric conditions from hindcasts of numerical weather prediction programs provided at a 6 hourly temporal resolution. With this model we can resolve some of the eddy processes and directly diagnose eddy transports. The post-processing of the vast amounts of model output is almost a Grand Challenge problem itself but for the first time we can see how and where the poleward transport of heat occurs in the Southern Ocean. Figure 1 presents a multi-year average and depth integrated view of this transport in units of 100GW (10^{11}W) per model grid cell. To put these figures into context it is worth noting that the wintertime peak demand for electricity in the UK is of the order of 50GW and there are nearly 2.5 million gridcells represented in figure 1. The heat transported south is eventually lost to the atmosphere. The overall effect can be be judged by averaging heat transport around the zonal (longitudinal) bands. The result is shown in figure 2 which shows the zonally integrated northward eddy transport (solid line) in comparison to the zonally integrated total ocean transport (dotted line). Also shown is the heat loss to the atmosphere (dashed line) which accumulates from Antarctica to the latitude concerned. The heat transports are shown in units of PetaWatts (10^{15}W or 1 million GW). The significance of the eddy transport is clear and over some latitude bands it is the major component of the total heat transport. Climate models will need to parameterise this transport better (a goal which remains elusive) or to accept the need to resolve the processes better. Based on such evidence we can expect the next generation of climate models to include higher resolution component models and to place ever greater demands on high performance computing facilities.

A detailed account of this work, *Eddy advective and diffusive transports of heat and salt in the Southern Ocean*, M.M. Lee, A.J.G. Nurser, A.C. Coward and B.A. de Cuevas, has been submitted for review to the Journal of Physical Oceanography.

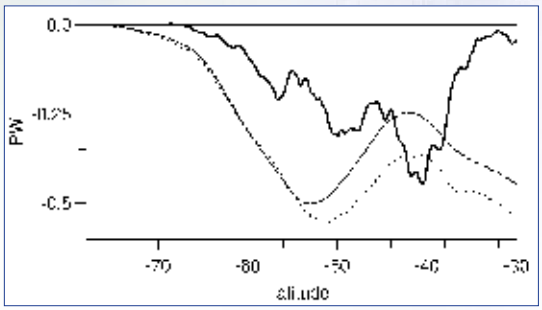

Something for nothing? Simultaneous multithreading (SMT) on HPCx

Mike Ashworth, HPCx Terascaling team

One of the incidental benefits of last year's upgrade to IBM's POWER5 technology is the opportunity to use symmetric multithreading (SMT).

Multithreading is being increasingly used to increase the efficiency of processor utilisation. By running multiple threads simultaneously SMT allows instructions from one thread to execute when those from the other are stalled (eg while waiting for data to be retrieved from cache or main memory) or when execution units are not being used. This has the potential to reduce the number of idle cycles and increase application code efficiency.

POWER5 processors support SMT with two threads (MPI tasks or OpenMP threads) able simultaneously to access the resources of each physical processor. When running with SMT enabled each 16-processor node appears to have 32 virtual CPUs. This is selectable by the user in the LoadLeveler job submission script.

For example a job script containing the following lines:

```
#@ cpus = 512
#@ tasks_per_node = 32
#@ requirements = ( Feature == 'SMT' )
```

will allocate 512 virtual CPUs with 32 tasks per node. LoadLeveler will therefore assign 16 nodes to the job and place the job in the par256_1 class. Obviously, the amount of memory available per thread (specifically the rss_limit per process) will be half that of an equivalent non-SMT job – you don't get something for nothing! Note that this won't work with shared nodes (interactive, course) or with the serial queue.

We have been testing SMT for a range of application codes. The benefits are dependent on the type of code and the number of processors. The greatest benefit is found at low and medium processor counts, with the largest improvement being a factor of 1.4 for classical molecular dynamics codes. The results of our evaluation are available in an HPCx Technical Report [1]

We encourage you to try out SMT and see whether it brings any benefits for your codes. Please tell us your experiences.

[1] An Investigation of Simultaneous Multithreading on HPCx, HPCx Technical Report HPCxTR0604, A. Gray, J. Hein, M. Plummer, A. Sunderland, L. Smith, A. Simpson and A. Trew

www.hpcx.ac.uk/research/hpc/technical_reports/HPCxTR0604.pdf

Simulations too large for HPCx alone?

Gavin Pringle, HPCx Terascaling team

Distributed European Infrastructure for Supercomputing Applications

Large simulation? Not enough cycles? Perhaps your simulation can run as a Taskfarm where each node is itself a supercomputer, or perhaps you would like to run on alternative architectures or, more simply, you'd just like to have more cycles?

DEISA is a consortium of leading national supercomputing centres that currently deploys and operates a persistent, production quality, distributed supercomputing environment with continental scope.

The purpose of the DEISA Extreme Computing Initiative (DECI) is to enhance the impact of the DEISA research infrastructure on leading European science and technology.

This initiative consists in the identification, enabling, deployment and operation of 'flagship applications' in selected areas of science and technology. These leading, ground breaking applications must deal with complex, demanding, innovative simulations that would not be possible without the DEISA infrastructure, and which would benefit – if accepted – from the exceptional resources of the Consortium.

The 2nd call for challenging projects and applications from computational sciences has been launched. In addition to providing resources from the DEISA infrastructure, support for complex application enabling will be made available. Deadline for proposals will be July 15, 2006.

See: www.deisa.eu/applications/index.php for more details.

UK Applied Aerodynamics Consortium Conference

Andrew Sunderland and Dave Emerson, HPCx terascaling team

Loughborough 5th–6th April

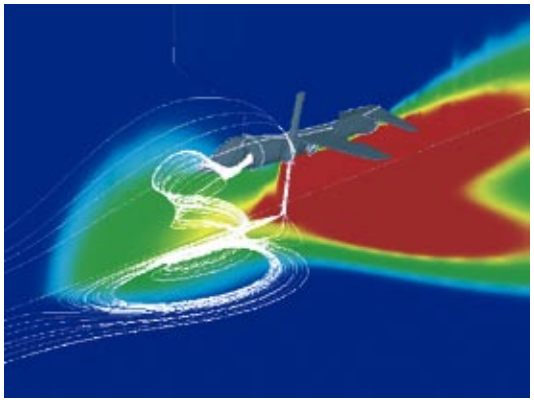

The UK Applied Aerodynamics Consortium (UKAAC) is an EPSRC funded activity (GR/S91130/01) that is led by Prof. Ken Badcock (Liverpool) and enables twelve academic institutions to access HPCx to tackle problems in aerodynamics. In April, the UKAAC held a 2-day conference at the University of Loughborough. The aim of the conference was to present results from research undertaken on HPCx to an audience of leading academic researchers and key representatives from industry, particularly from the Aerodynamics National Advisory Committee (ANAC), and to highlight the work within an international context. The meeting attracted 73 participants with 21 attendees from industry. The delegates represented 17 academic institutions and 10 industrial companies including BAE, MBDA, Rolls-Royce, QinetiQ, Airbus, and ARA.

Keynote presentations included speakers from NASA, U.S. Air Force, and SGI. The sessions over the two days of the conference covered the following 6 themes:

1. High-Fidelity Simulation of Helicopter Interactions
2. Simulation of a free flying flexible aircraft
3. Simulation of Vertical Landing Aircraft
4. Aeroelasticity studies for aero-engine core-compressors
5. Simulation of Internal Air System
6. Simulation of Store Separation

A number of these themes were prompted by the Defence Aerospace and Research Partnerships (DARPS). These are industry-led university-based partnerships that focus on important areas of aeronautics research and are supported by EPSRC, DTI and MOD, and industrial companies. Many consortium partners have research programmes that are funded through DARPS, and areas that could benefit substantially from access to world-class supercomputing facilities are the Programme for Unsteady Aerodynamic Modelling (PUMA) and Rotorcraft DARPS.

Many interesting and challenging applications of aerodynamics were demonstrated, including abrupt wing stall behaviour, uncontrolled aircraft spin, helicopter rotor vortex interactions, space shuttle booster flow simulation, and hot-gas ingestion simulations during vertical take-off and landing of Harrier aircraft. Representatives from EPSRC and CCLRC addressed the audience on computational developments including the HECToR procurement, the current challenges facing HPC exploitation, and an introduction to novel technology solutions.

The final afternoon involved presentations from several of the ANAC sub-committees on national priorities and strategies for aerodynamics. Concluding the conference was a panel discussion on suggested themes for the future UK applied aerodynamics consortium, chaired by its new Principal Investigator, Nick Hills (Surrey). Two important topics for the next proposal are the establishment of an activity in computational hypersonics and the need for more advanced post-processing and visualisation capabilities.

The consortium has one more year to run on HPCx.

Above: flow into the intake of a Harrier (provided by Cambridge University). Right: Nick Hills, UKAAC (United Kingdom Applied Aerodynamics Consortium).

HPCx annual seminars

Andrew Sunderland, HPCx Terascaling Team

The third in the series of HPCx Annual Seminars was held at CCLRC Daresbury Laboratory on Monday 5th December 2005.

The theme of the seminar was 'Capability Science on HPCx' and presentations were made from many areas of scientific research. Speakers included Roger Proctor from the Proudman Oceanography Laboratory, Henning Weber from Deutscher Wetterdienst, Professor Ken Badcock from the University of Liverpool and Shantenu Jha from UCL. Presentations were also made by staff on HPCx related topics including Phase 2A performance analysis, HPCx consortium visits and message passing optimisation techniques. The day was rounded off by an HPCx User Group Discussion chaired by David Henty.

The fourth HPCx annual seminar will be held on the 4th October 2006, at the National e- Science Centre in Edinburgh. See: www.hpcx.ac.uk/about/events/

Copies of all the Annual Seminar presentations are available from the HPCx website at: www.hpcx.ac.uk/about/events/annual2005/

Rise of the machines: EPCC at the Edinburgh International Science Festival 2006

Kenton D'Mellow, HPCx Terascaling team

edinburgh international science festival

After last year's successful participation in the Edinburgh Science Festival, EPCC returned to enthrall a crowd of over 100 interested festival-goers. Dr David Henty delivered one of the festival's 'Cutting Edge' series of talks, entitled Supercomputing: Rise of the Machines.

David started with a whistle-stop tour of supercomputers hosted by EPCC over the years, and took a look at some of science's current 'Grand Challenges' and the modern supercomputers that are addressing these huge computational problems.

He described a wide range of the problems that supercomputers are built to solve, ranging from studies of the tiniest particles of matter and anti-matter that are the building blocks of our world, through topical issues of medicine, climate change and clean energy production, right up to simulations of the evolution of the entire Universe.

The audience was then introduced to many of the issues associated with building and programming massively parallel computers, and some of the steps being taken to reduce the power consumption of these machines. In a series of entertaining demonstrations, David highlighted the challenges of scalability and communication, and the efficient decomposition of calculations and data into manageable parts. Along the way, the audience were presented with a host of interesting facts, like how the supercomputers of 25 years ago had less computing power than a typical mobile phone today. Technology has moved at such a fast pace since then that modern supercomputers like HPCx are almost a million times more powerful than those original machines.

Finally, David gave an idea of what is to come in the future, both in terms of the machines we'll be using, and the science we'll be able to do with them. A thoroughly stimulating event!

www.sciencefestival.co.uk

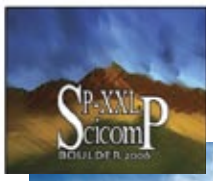

ScicomP 12

Boulder Colorado, July 2006

Image copyright University Corporation for Atmospheric Research

The ScicomP meeting series enables computational scientists and engineers to learn about tools and techniques for developing applications that achieve maximum performance and scalability on IBM HPC systems. Technical presentations will highlight recent results and advanced techniques, and will provide the kind of information, expertise, and experience that scientific and technical application developers need but cannot easily find elsewhere.

ScicomP 12 is hosted by the National Center for Atmospheric Research and will be held alongside the summer meeting of SP-XXL, the user group for system administrators of large IBM installations.

www.cisl.ucar.edu/info/scicomp/