CSAR FOCUS

Edition 11
Winter - Spring 2004

high performance computing at Manchester ~ taking research into the future

www.csar.cfs.ac.uk

Igniting Innovation Across the Globe
SC Global 2003

Also in this edition ... Research Activities at CSAR including

Turbulent flow over groups of urban-like obstacles
Numerical modelling of the flame front dynamics
Large Eddy Simulation of Separation from Continuous Surfaces

THE UNIVERSITY of MANCHESTER

CSC

sgi

Contacts

CSAR Helpdesk

Telephone: 0161 275 5997 / 275 6824
Fax: 0161 275 6800
Email: csar-advice@cfs.ac.uk

Terry Hewitt

Head of SVE, Manchester Computing, University of Manchester

Telephone: 0161 275 6095
Fax: 0161 275 6800
Email: w.t.hewitt@man.ac.uk

Dr Mike Pettipher

Head of HPC Team, Manchester Computing, University of Manchester

Telephone: 0161 275 6063
Fax: 0161 275 6800
Email: m.pettipher@man.ac.uk

Claire Green

Editor of CSAR Focus, HPC Team , Manchester Computing, University of Manchester

Telephone: 0161 275 6822
Fax: 0161 275 6800
Email: claire.l.green@man.ac.uk

Copyright

All pictures are reproduced by permission. Any further reproduction will require the owner's permission.

Disclaimer

Articles are the responsibility of the individual author. The Editor, CSAR and the University of Manchester assume no responsibility or liability for articles or information appearing in this publication. Opinions expressed in some articles may not reflect those of CSAR.

Contents

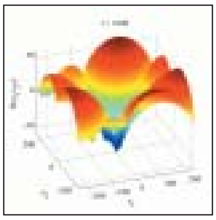

6: Numerical modelling of the flame-front dynamics

18: Op3D-Grid

Editorial

Welcome to this edition of *CSAR Focus* which focuses on the research activities carried out on the CSAR systems. Collaborators from the Universities of Reading and Southampton share their experiences so far in using our new system Newton, in the field of Urban Meteorology, to accurately model flow and dispersion in urban areas. Please see page 4 for further information regarding the work carried out within this particular area. Also included are reports on research activities in large eddy simulation, flame front dynamics, computational chemistry, aerodynamics and aeroacoustics. Thank you to all those who contributed.

Since the last publication of *CSAR Focus* CSAR and SVE staff have been involved in several important events, including the MRCCS/NSF Summer School and SC2003 which are both featured in this edition. SC2003 was an extraordinary success story for Manchester Computing - Penny Richardson details the highlights of this event in her article on page 22. SC Global is the distributed element of the conference, using Access Grid technology to enable a huge number of sites to take part simultaneously. Speakers at SC Global participated from their home institutions in seven different countries, five sovereign tribal nations, distributed across four continents. Michael Daw of the SVE team and the Technical Director of SC Global 2003 outlines the technology behind the event and its success- turn to page 20 to find out more.

I hope you enjoy this issue of *CSAR Focus* which in addition to those items already mentioned above, also contains articles on a diverse range of topics from a look at High Performance Computing in Germany (specifically the HLRN system) to the life of Newton (the man not the machine). If you would like to contribute to the next edition please contact me via the Helpdesk (csar-advice@cfs.ac.uk).

Claire Green
Editor, *CSAR Focus*

Turbulent flow over groups of urban-like obstacles

Omduth Coceal and Stephen E. Belcher (Department of Meteorology, University of Reading); T. Glyn Thomas and Ian P. Castro (School of Engineering Sciences, University of Southampton)

The ability to compute complex turbulent flows around bluff bodies is important for many applications, particularly with regard to flow control in engineering. More recently, stringent demands for urban air quality control and increased concern about the release of hazardous materials in cities have heightened the need to accurately model flow and dispersion in urban areas.

Understanding the dynamics of mixing and transport in urban areas also enables parameterisations of urban areas to be developed for incorporation into numerical weather prediction models, and thereby increases the accuracy of weather forecasts.

The present work is a collaboration between the Department of Meteorology, University of Reading, and the School of Engineering Sciences, University of Southampton. Highly resolved Large Eddy Simulations (LES) and Direct Numerical Simulations (DNS) are performed to simulate turbulent flow around arrays of cubical, urban-like, obstacles. The arrays are arranged in different configurations and at different packing densities to investigate the effects of building proximity and layout on the flow. Initial aims are to compute time and spatial averages of flow quantities, from which parameterisations of drag and turbulence can be extracted. Another important aim is to identify organised structures in the flow and to understand their dynamics. Future work will model dispersion of tracers placed in the flow.

The model used is the Multiblock LES code, which has been developed by Glyn Thomas at Southampton University. The code is parallelised using MPI, and until recently has been run on a Linux cluster at Southampton as well as on the CSAR Cray T3E (Turing). In October 2003, Neil Stringfellow at CSAR ported and optimised the code for the new SGI Altix (Newton). The performance of the code on the Altix has been astounding, giving speed-ups of up to an order of magnitude as compared to the T3E, and of two orders of magnitude as compared to the Linux cluster at Southampton! This has compressed run times from weeks to days, and has enabled research possibilities that could not be envisaged before.

The substantial computational resources make it possible to compute highly resolved turbulent flows at sufficiently high Reynolds number and for sufficiently long to allow the calculation of converged and stable statistics. This is important to establish the credibility of the simulations. Results from the simulations are in very good agreement with published wind tunnel data. These results indicate the potential of the present modelling approach for these important problems in urban meteorology.

Acknowledgements

We are very pleased to acknowledge the expert help offered by Neil Stringfellow of CSAR in porting and optimising the code for the Altix, and the kind assistance of CSAR staff.

Omduth Coceal is funded by UWERN under a NERC grant. This work forms part of the UWERN Urban Meteorology Programme

http://www.met.rdg.ac.uk/Research/urb_met

Snapshot of velocity vector field over a staggered array of cubes: (a) vertical slice through middle of cube, (b) horizontal slice at half cube height.

Each cube is resolved with 32 x 32 x 32 gridpoints.

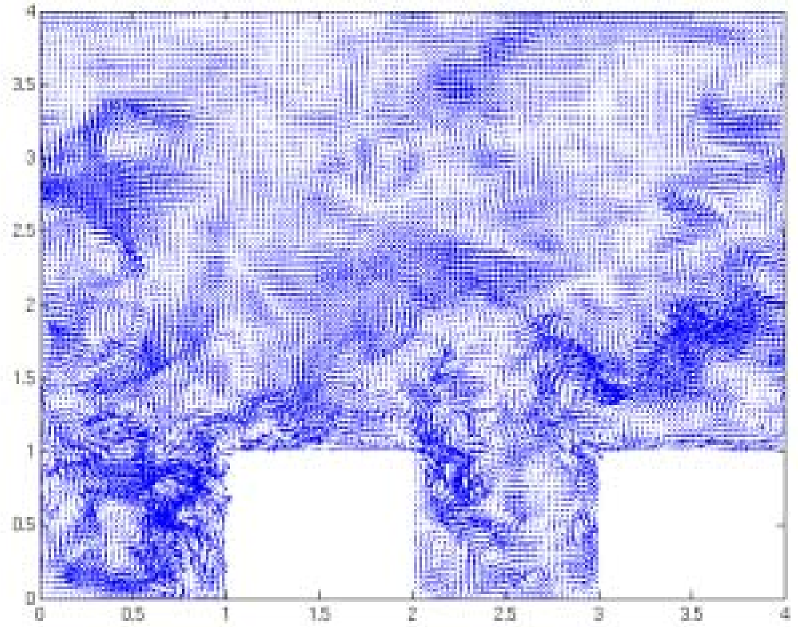

Figure 1: vertical slice through middle of cube.

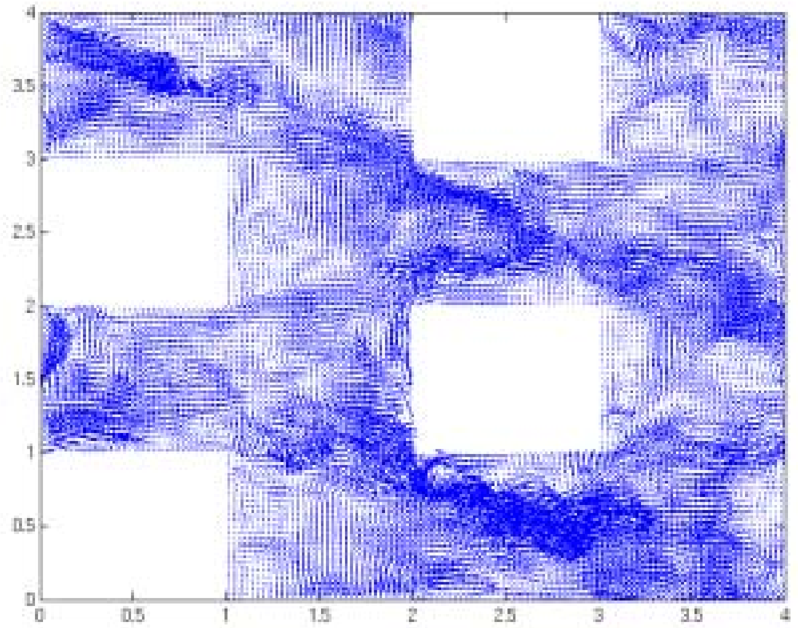

Figure 2 horizontal slice at half cube height.

Numerical modelling of the flame front dynamics

V. Karlin and J. Mai
University of Central Lancashire, Preston, UK

It was observed in experiments that expanding flame fronts do not remain smooth given enough time. Instead, a structure of cells is formed on their surface, and the flame fronts can even experience substantial acceleration. This phenomenon was linked with the intrinsic long wave instability of flame fronts which is also known as the Darrieus-Landau or the hydrodynamic flame instability. Huge variety of time and space scales involved makes it inefficient to study this phenomenon on the base of the general Navier-Stokes system. A hierarchy of asymptotic mathematical models of cellularization of hydrodynamically unstable flames was obtained as an alternative approach to the problem. For example, the Sivashinsky equation

$$\frac{\partial \Phi}{\partial t} - \frac{1}{2}\left|\nabla \Phi\right|^2 = \Delta\Phi + \frac{\gamma}{2}(-\Delta)^{1/2}\,\Phi, \quad \chi \in R^2, \quad t > 0, \quad (1)$$

which governs evolution of the perturbation of the propagating plane flame front, was derived in [1]. A more sophisticated asymptotic model of expanding spherical flames was suggested in [2].

In spite of the fact that the asymptotic models of cellular flame dynamics are much simpler than the original Navier-Stokes system, they are still very difficult to solve. Accurate numerical resolution of the surface patterns requires substantial computing resources and the situation is worsened by the very high sensitivity of these patterns to noise [3]. The formation of a cell on the surface of a plane flame front and of a cellular structure on the surface of an expanding flame is illustrated in Figures 1 and 2 respectively. The calculations were carried out on the Cray T3E 1200 computer (Turing) in the Manchester Computing Centre. Details of the algorithm and its parallel implementation are given in [4].

So far, our investigations have demonstrated that the formation of the cellular structures is triggered by noise, which is always present both in numerical and physical experiments in one or another form. These essentially nonlinear cellular structures are generated through the huge linear transient nonmodal amplification of perturbations of noisy origins. Their final appearance and dynamics on the flame surface is governed by essentially nonlinear mechanisms intrinsic to the physics of the flame front dynamics. As the size of the flame increases, the nonmodal amplification of noise grows. Hence, the number of cells on the flame surface grows as well, contributing into its total area and, eventually, in the acceleration.

This research is supported by the EPSRC grant GR/R66692.

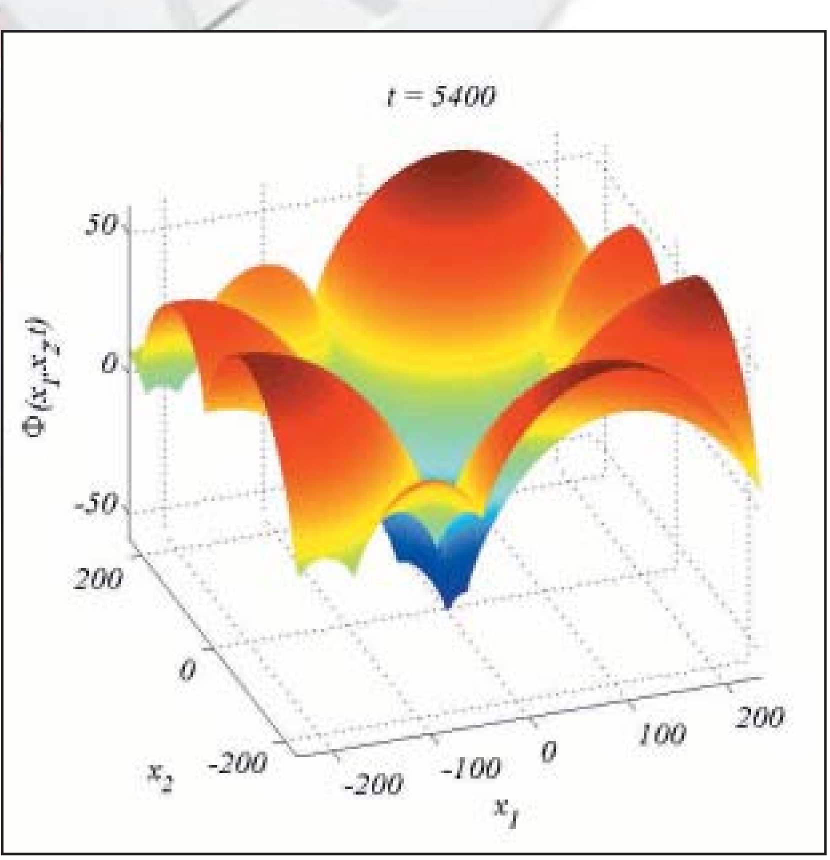

Figure 1: Formation of a cell on the surface of a plane flame front.

References

[1] G.I. Sivashinsky. *Acta Astronautica*, 4:1177–1206, 1977.
[2] Y. D'Angelo, G. Joulin, and G. Boury. *Combustion Theory Modeling*, 4:1–22, 2000.
[3] V. Karlin. *Proc. Combustion Institute*, 29(2):1537–1542, 2002.
[4] V. Karlin, V. Maz'ya, and G. Schmidt. *J. Computational Physics*, 188(1):209–231, 2003.

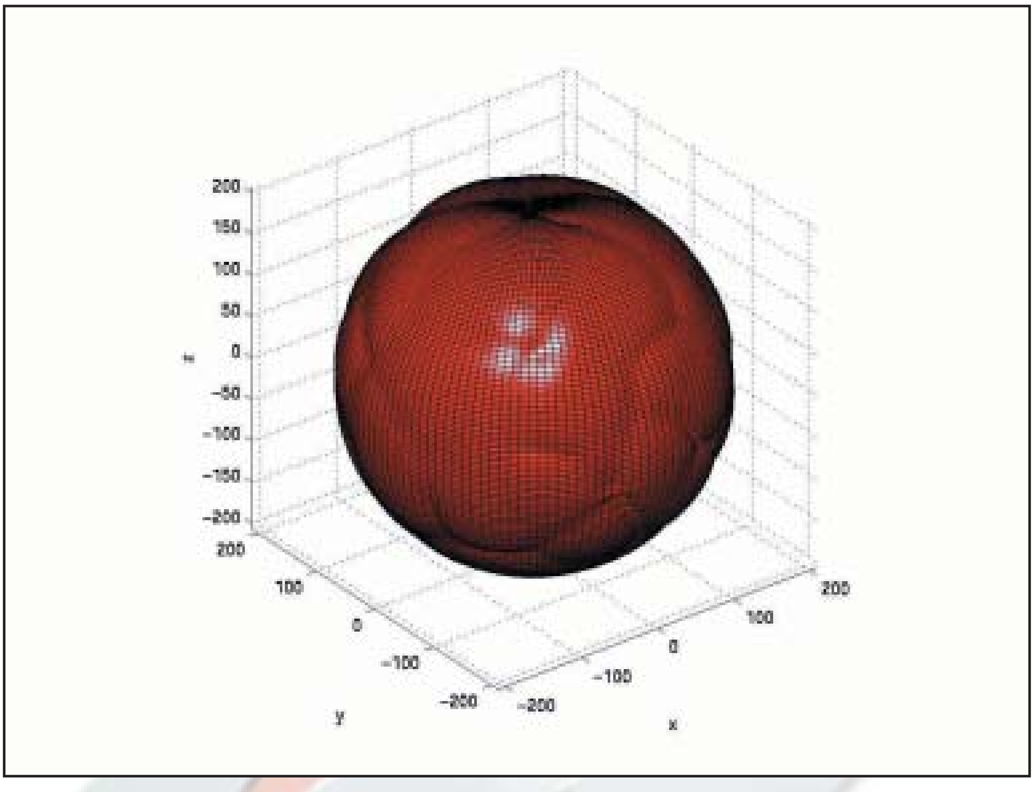

Figure 2: Formation of a cellular structure on the surface of a spherical flame front as a result of a stochastic perturbation.

SGIUG 2004 Conference

May 23-27, Orlando, Florida, USA
Radisson WorldGate Hotel (just one mile from Walt Disney World®)

This year's must attend conference promises to be an exciting event with users assembling from around the globe to present and demonstrate the very latest advancements in:

- Application Optimization and Tuning
- Data Fusion and Data Mining
- Digital Media and High-definition Processing
- Grid and Collaborative Computing
- High Performance Computing
- Scaling Linux®
- Security
- Storage Consolidation and Optimization
- Visualization

and more

Register today at www.sgiug.org
or call 1.800.850.2307 (toll free) or 1.508.486.9176 (outside USA)

Large Eddy Simulation of Separation from Continuous Surfaces

Lionel Temmerman and Michael A. Leschziner,
Aeronautics Department, Imperial College London

Turbulence is an inherent feature of the large majority of engineering and environmental flows. It stems from inertial instabilities that arise when any fluid is strained at sufficiently high rates – say, by shearing. In simple terms, the motion in a turbulent flow may be thought of as consisting of a mean component with a superimposed field of unsteady, non-repeatable interacting eddies having size and time-scales ratios that span, typically, 3-4 orders of magnitude – that is, the smallest eddies may be as small as 0.01% of the global linear 'dimension' of the flow. The practical importance of turbulence arises from the strong mixing it causes – a process that is central to dispersion, chemical reaction, combustion, frictional losses, drag, and the general behaviour of flows around streamlined and bluff bodies and within passages.

Simulating turbulent flows in realistic conditions is a formidable computational task. It entails the numerical solution of the coupled, non-linear set of (spatially) three-dimensional Navier-Stokes equations that describe the time evolution of the flow of interest over a mesh of nodes, volumes or elements covering the flow domain. Because of the wide range of scales involved, the grids required for most practical flows would need to have of order 10^9-10^{11} nodes to resolve all details, and the extraction of practically important statistical properties would require an integration over 10^6-10^7 time steps. This is not regarded as a tenable approach, either today or in the future.

The alternative route thus taken is one that assimilates all or a part of the unsteady turbulence dynamics into a statistical model. The former option starts with *Reynolds averaging* the Navier Stokes (RANS) equations, to yield equations for the time- or ensemble-mean quantities. A *turbulence model* is then required to determine unknown but crucially important correlations of turbulent velocity fluctuations. The quality of predictions then depends critically on the quality of the turbulence model, and the development of general models is a highly challenging area into which many hundreds of man-years of research have been invested over the past 4 decades (see reference [1]).

The compromise route in which only part of the turbulence is represented statistically is called Large Eddy Simulation (LES). This is based on the notion that all major dynamic effects may be captured by resolving eddies not smaller than about 1% of the relevant global dimension of the flow, while the effects of smaller eddies may be represented by a *subgrid-scale turbulence model.* Although this approach is much more economical than a full-resolving simulation, it is still costly – around 50-100 times higher than computations based on Reynolds-averaged formulations, typically involving 10^6-10^7 nodes and 10^5 time steps and requiring computing times of order 10^4 CPU hours on present multi-processor machines, such as those operated by the CSAR in Manchester, UK.

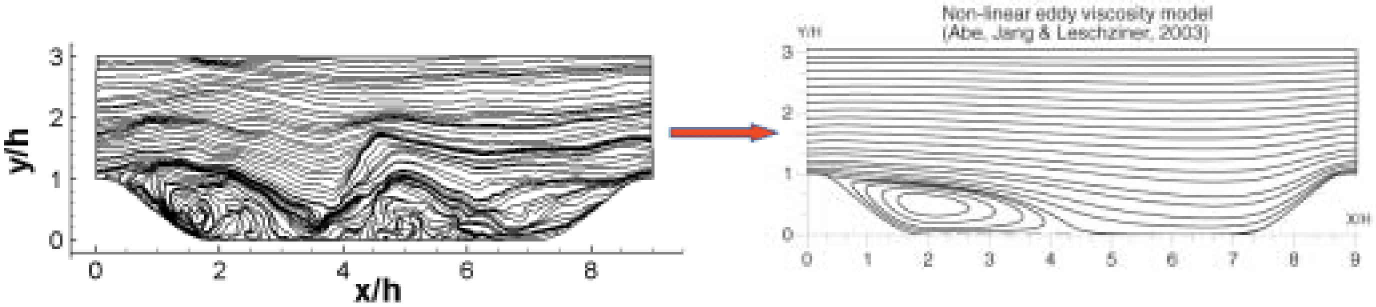

(a) Simulation

(b) Statistical modelling

Figure 1: Instantaneous and time-averaged view of a flow separating from a duct constriction

While LES allows the unsteady turbulence dynamics to be resolved, extreme events to be captured and the mechanisms leading to particular statistical properties to be understood, it faces a whole range of difficult challenges and limitations in practical conditions, apart from the cost. Thus,

- LES does not tolerate poor grid quality (high skewness, gradation and aspect ratio) ;
- LES requires high numerical accuracy, and does not, in particular, tolerate numerical dissipation;
- LES requires the spectral content of boundary conditions to be specified;
- The quality of the simulation can depend sensitively on sub-grid modelling, an area far less developed than Reynolds-averaged modelling;
- LES is very sensitive to the resolution at walls in conditions in which the gross behaviour of the flow depends on the structure of the wall boundary layer.

The last issue poses a particularly serious problem in many important engineering flows at high Reynolds numbers. As the wall is approached, the large scales progressively diminish in size, eventually approaching the scale at which energy is dissipated by viscous friction. In addition, the near-wall structure is highly anisotropic, characterised (in shear flow) by elongated vortical structures. Hence, the grid supporting the near-wall layer must approach one that would normally be used in a full simulation. The requirement that the grid-aspect ratio be constrained to accommodate the above structural features of the near-wall motions quickly leads to economically untenable grids as the turbulence Reynolds number increases beyond 1000.

A case in point is shown in Figure 2. This is a high-lift aerofoil operating at a chord Reynolds number of 2.1 million and at 13.3° incidence, at which marginal stall sets in on the rear end of the suction side. Apart from the need to capture separation, it is important to resolve the laminar-to-turbulent transition which is clearly visible on the front part of the suction side.

The computation shown in Figure 2 (see reference [2]) encompassed only a spanwise segment of 12% of chord, i.e. a very small portion of a practical wing, but required over 5 million nodes, and still was found not to represent the flow especially well. Here, as in other flows in which separation occurs from a continuous curved surface, the challenge is near-wall resolution and the quality with which the thin boundary layer is captured.

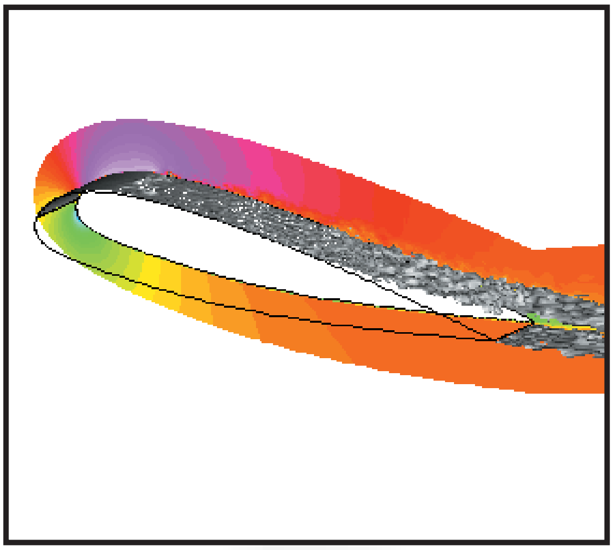

Figure 2: Large Eddy Simulation of a separated flow on a high-lift aerofoil

Another example is given in Figure 3. This shows the results of three simulations by Temmerman et al [3] with successively fine grids, ranging from 0.6 million to 4.8 million nodes. The sensitivity primarily reflects the resolution of the near-wall region, especially to the location at which the flow separates from the curved surface.

The resolution of separation from three-dimensional curved surfaces is even more demanding than that of the above (statistically) two-dimensional process. Engineering applications in which such separation is of major importance include fuselages, ship hulls and streamlined road vehicles. A generic case is shown in Figure 4. This is the flow around a three-dimensional hill-shaped obstruction in a duct, and the figures convey the flow topology on the hill surface by way of 'skin-friction lines'. Experimental data is available for a Reynolds number which is much too high for the flow to be simulated at acceptable costs, again because of the need to ensure adequate near-wall resolution. Hence, the simulation shown in Figure 4 has been performed at a lower Reynolds number (Temmerman et al [4]). The result shown in Figure 4(a) is a short-time average of the flow, thus conveying the complex turbulent motion in the separated, highly vortical wake, while Figure 4(b) show the statistically (long-time-averaged) state. Figure 4(c) is the result obtained for the same flow conditions with an steady-state RANS

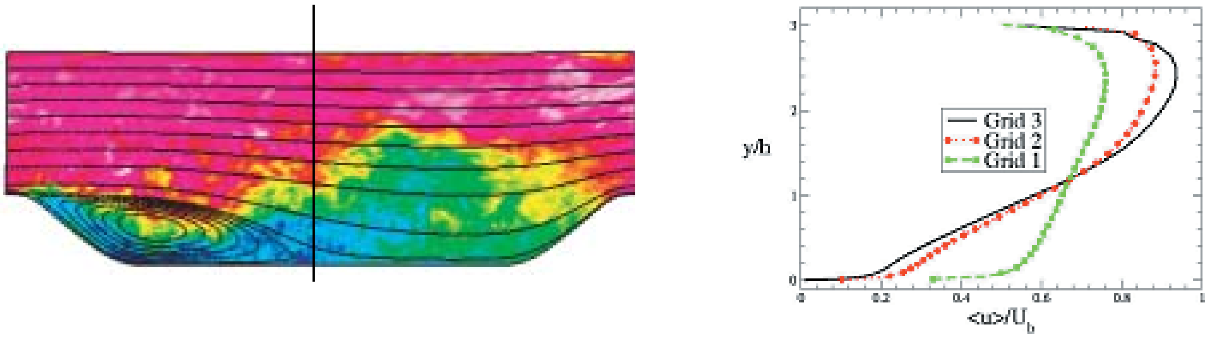

Figure 3: Sensitivity of time-averaged velocity field to grid density in a LES of a separated flow.

scheme operating in conjunction with an advanced turbulence model based on 'second-moment closure'. This case thus illustrates one important objective pursued by the writers when performing costly simulations, namely to assess the predictive realism obtained with statistical turbulence models that are designed to allow the time-averaged flow to be obtained without the costly computation of the details of the time-dependent turbulence field.

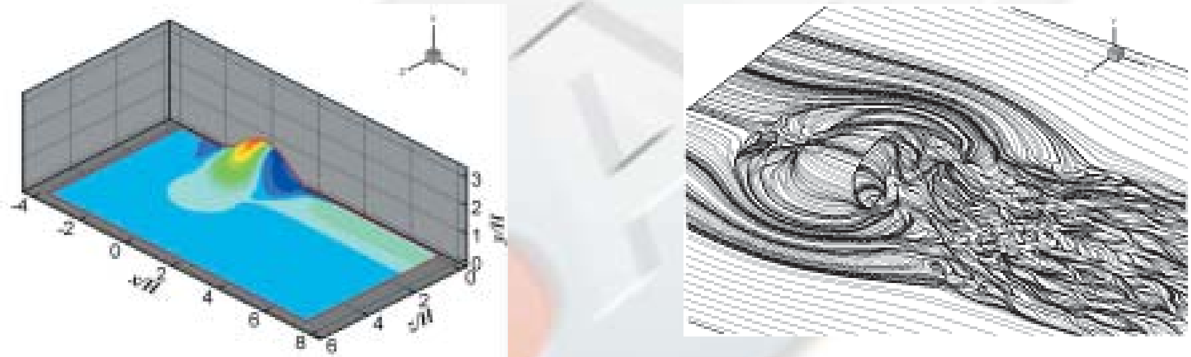

(a) Geometry & pressure contours

(b) LES, short-time average

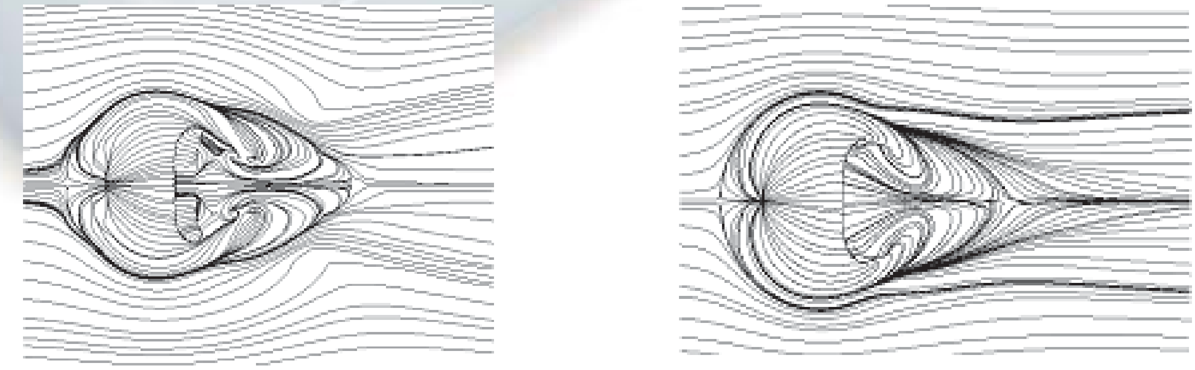

(c) LES, long-time average

(d) Steady RANS computation

Figure 4: Turbulent separated flow around a three-dimensional 'hill' in a duct – flow topology on hill surface.

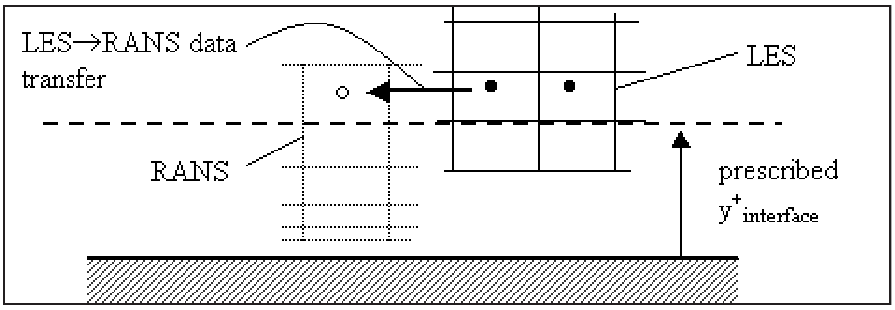

Figure 5: Coupled hybrid RAN-LES strategy.

Against the background of the severe limitations imposed by near-wall turbulence on LES for high Reynolds-number flow in realistic engineering conditions, much research is currently in progress on combining Reynolds-averaging (RANS) in the near-wall region with simulation away from the wall. This idea is rooted in the observation that satisfactory solutions for near-wall flow can be secured with the former strategy even for low-quality grids having very high aspect ratios. The principles of one such an approach – referred to as hybrid RANS-LES – are conveyed in Figure 5.

One crucial issue in this hybrid method is how to achieve compatibility of turbulence-related quantities across the interface, and this is the subject of much debate. As usual, the most challenging environment is one involving separation from continuous surfaces. Thus, Figure 6 shows an application of a hybrid method by Temmerman et al [5] to the flow shown in Figure 3. In this, the RANS-LES interface can be chosen arbitrarily and has here been chosen at grid planes which are well within the turbulent regime. The reference computation, on the l.h.s. of Figure 5 was performed with 5 million nodes, while the grid used for RANS-LES computation contained only 0.6 million nodes. As seen, agreement, although credible, is far from perfect, and further research is in progress in efforts to improve the method.

References

1. Jang, Y.J., Leschziner, M.A., Abe, K. and Temmerman, L., Investigation of anisotropy-resolving turbulence models by reference to highly-resolved LES data for separated flow, *Flow, Turbulence and Combustion*, Vol 69 (2002), pp. 161-203.
2. Davidson, L., Cokljat, D., Froehlich, J., Leschziner, M.A., Mellen, C., Rodi, W. (Eds.), LESFOIL: Large Eddy Simulation of Flow Around a High Lift Airfoil, *Notes on Numerical Fluid Mechanics and Multidisciplinary Design*, Vol 83, Springer (2002).
3. Temmerman, L., Leschziner, M., Mellen, C. and Froehlich J., Investigation of subgrid-scale models and wall-function approximations in Large Eddy Simulation of separated flow in a channel with streamwise periodic constrictions, *Int. J. Heat Fluid Flow*, Vol 24 (2003), pp. 157-180.
4. Temmerman, L., Chen, W. and Leschziner, M.A. (2004), A comparative study of separation from a three-dimensional hill using LES and second-moment-closure RANS modeling, *4th European Congress on Computational Methods in Applied Sciences and Engineering, ECCOMAS 2004*, to appear.
5. Temmerman, L., Leschziner, M.A. and Hanjalic, K., A combined RANS-LES strategy with arbitrary interface location for near-wall flows, *Proc. 3rd Int. Symp. on Turbulence and Shear Flow Phenomena* (2003), pp. 929-934.

Contact Details

Michael A. Leschziner
Aeronautics Department, Imperial College London, UK
email: mike.Leschziner@imperial.ac.uk

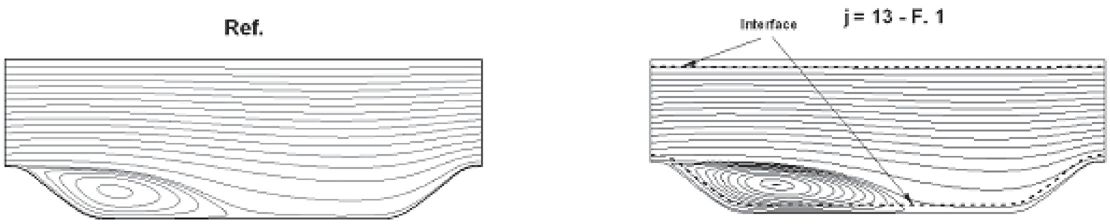

Figure 6: Comparison of solutions obtained with highly-resolved 5-million-node LES (l.h.s.) and under-resolved 0.6-million-node hybrid RANS-LES.

Impulsive Start of 3-D Shear Driven Cavity Flows

Antonio Filippone, Dept. of Mechanical, Aerospace, Manufacturing Engineering, UMIST and Jess A. Michelsen, Department of Mechanical Engineering, Technical University of Denmark

Cavity Flows

In this study we have started an investigation of the flow physics of the vortical structures created by the cavity at a laminar Reynolds number. Deep cavities of the nature described here have been used in the past as resonance systems, and they are commonly referred to as *Helmholtz* resonators. More generally, flows over cavities occur in a wide range of engineering processes. They are a model for re-circulating flows, heat and mass transfer - and related physics.

We considered the impulsive stat of a shear driven cavity flow over a channel partly closed by slits. Compared to the classical lid-driven cavity flow (steady-state, lid-driven, without mass transfer, with well posed boundary conditions) the set-up shown in Figure 1 features formidable complexity - at the expense of minimal increased geometrical details. These flows are also characterized by unsteady three-dimensional effects, with oscillations that are self-sustained. For the general cavity flow to be established, it is necessary that periodic changes of in- and out-flux take place.

Computational Model

The three-dimensional cavity is shown in Figure 1. The spanwise aspect-ratio was 3.0. The lid opening considered was $h/L = 1/4$, and the Reynolds number, based on the free stream U and streamwise length L was $Re = 3{,}000$. The cavity was immersed in a channel flow, simulated over a length of 5 cavities downstream, and 1.5 cavity lengths upstream. The upper surface of the channel was fixed at two cavity lengths above the lid. The Reynolds number affecting the flow inside the lid is based on the reference length L. The flat plate boundary layer at the entrance of the cavity is also important.

The flow was started from rest and simulated in a fully unsteady mode. Only half a model was simulated. Preliminary research showed that the flow is symmetric with respect to the centre plane over a time of several seconds, though this statement may not hold over a long simulation time. Some researchers have pointed the importance of the boundary conditions, and seemed to conclude that correlation of between experimental and computational data may diverge over a long simulation time.

The flow solver is based on the discretised Reynolds-averaged Navier-Stokes equations. It is parallelized with MPI for execution on vector computers using a non overlapping decoupling technique.

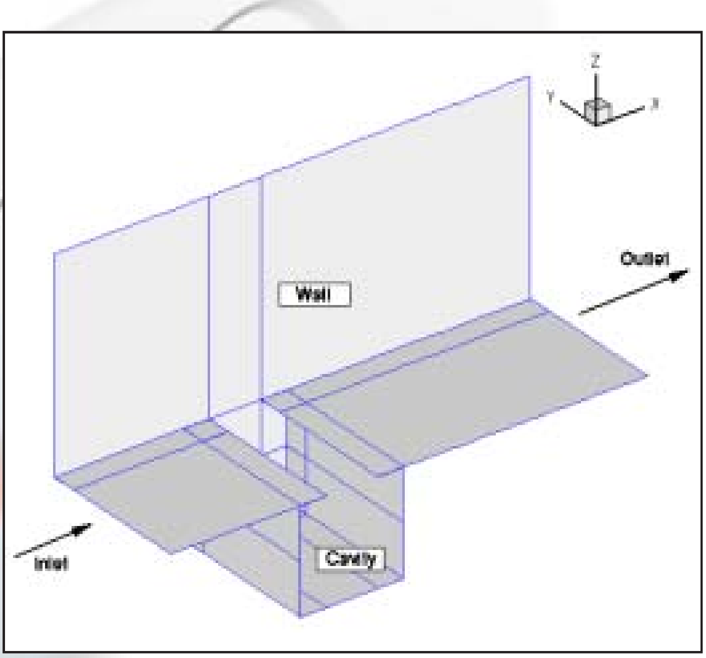

Figure 1: computational model (half-cavity).

Flow Solver Performance

The flow solver was tested on the classical lid-driven cavity flow started from rest. The solution reached an asymptotic value, and the residual history showed that the solver converged to machine accuracy on all the variables, 10^{-16} in about 2,000 iterations.

The results of the runs shown below were obtained on SGI Origin Green computer with 256 nodes. Up to 20 processors were used in our simulations. This computer allocated batches of 4 processors per run, therefore each run had to request a multiple of this number, which is not always optimal; the optimal condition is estimated

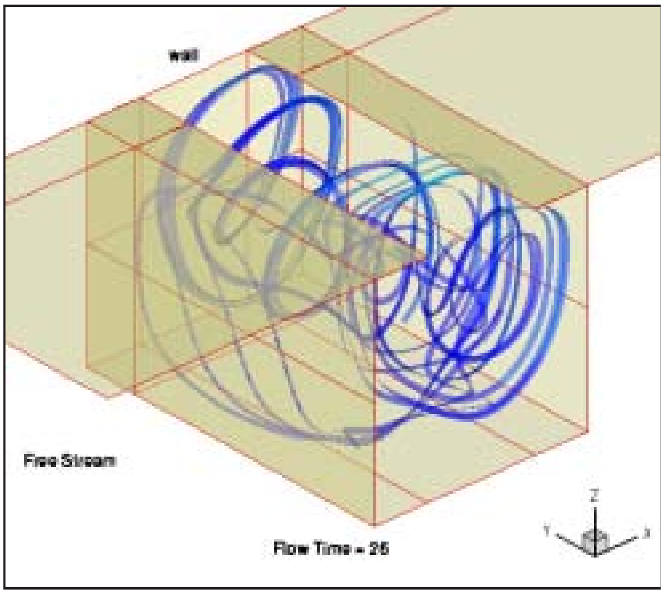

Figure 2: streamtraces at flow time = 26.

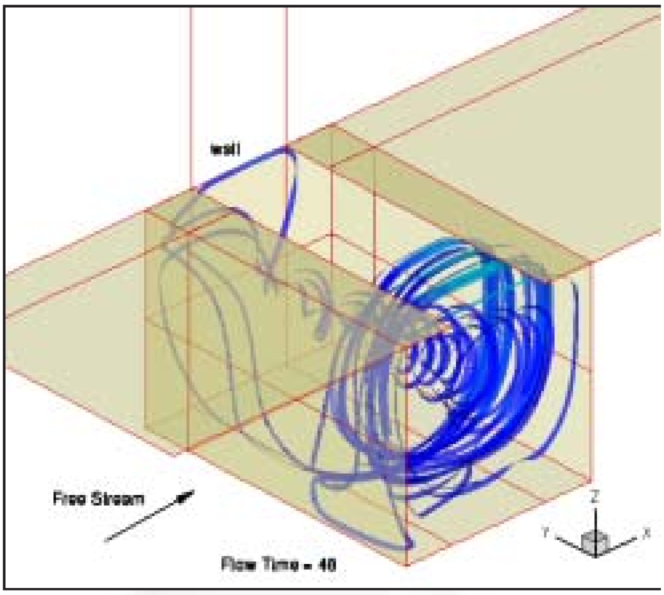

Figure 3: streamtraces at flow time = 48.

to occur when each mesh block is allocated to a single processor, or a multiple of the number of blocks. In spite of this difficulty, the solver was fast. Figure 4 shows the performances resulting from multi-processor runs. Each point was calculated from an average of time 650 iterations.

The computational grid was made of 18 cubic blocks (on the half-model) n^3, where n = 40 was the number of cells in each direction. The number of cells in the cavity was 768,000, out of a total 1,152,000. The simulations shown were carried out up to a normalised flow time equal to 90.

Results

The flow in the cavity is not confined to a two-dimensional pattern. Currents of fluid transfer mass from the end-walls to the centre and back again, see Figures 2 and 3. The end-walls are known to be responsible for this particular feature. The fluid in the primary vortex core tends to spiral towards the cavity's centre, following the pressure gradient.

The fluid is set in motion inside the cavity thanks to the shear layer at the open lid. Since the flow is started from rest, a considerable amount of time is needed to reach conditions such that the entire cavity is affected.

A number of other complex features exist: Taylor-Goertler vortices, hairpin vortices, other secondary vortices of short life-span. We have found that the flow is periodic with a very long wave. Research is undergoing to investigate all these fluid dynamic effects.

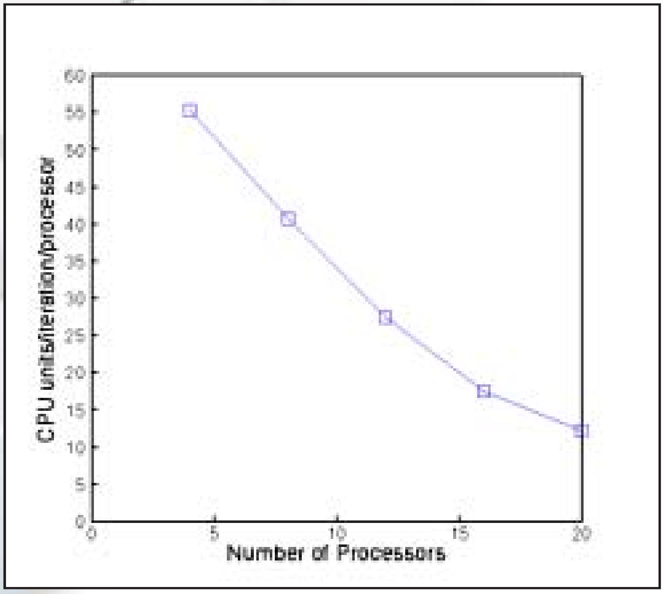

Figure 4: speed up rate on SGI Green.

Contact Details

Dr. A. Filippone
UMIST
Dept. Mechanical Engineering, Manchester M60 1QD
a.filippone@umist.ac.uk
http://www.me.umist.ac.uk/staffpgs/af

Jess A. Michelsen
Department of Mechanical Engineering
Technical University of Denmark
Lyngby DK-2800

Theoretical Chemistry on Supercomputing Facilities

Mark Vincent
Dept. of Chemistry, University of Manchester

Looking at past issues of CSAR Focus one can see that there are frequent references to computational chemistry on its pages. However, these tend to be references to a small part of what computational chemists do on supercomputers. In this article we will attempt to give an overview of what Computational Chemists are doing and why it is important to use state of the art computational facilities. I shall mainly concentrate on the work of the 'Theoretical and Computational Chemistry' group at The University of Manchester.

Dynamics – studying the movement of atoms.

One of the simplest chemical reactions is symbolically of the form

A + BC => AB +C

i.e. an atom reacting with a diatomic molecule giving a diatomic and an atom. Despite the simplicity of this reaction, complex dynamical processes are taking place. The vibrational and rotational states of the diatomic molecule (BC) have to be coupled to the translational motion of the atom A. For example, A can collide with BC head on or in a glancing blow (possibly leading to no reaction but energy transfer). To study such reactions requires an accurate scattering solution of the Schrodinger equation.

The other end of the size spectrum of molecules from a diatomic is a protein. The importance of understanding how proteins function so that drugs, pesticides or antibiotics can be developed that destroy the function of the enzyme, is important for human health and agriculture. To do this often requires knowledge of the dynamics of the protein. While it is true that the mode of operation of some enzymes can be relatively easily determined from their static X-ray structures, others catalyse reactions by changing from their ground state structure to a higher energy structure or by undergoing large amplitude motions. To determine if these large structural changes take place prior or during reactions in the protein requires the sampling of a huge part of the conformational space of the protein. One other form of dynamics is Car-Parrinello which we will mention later.
[Programs: CPMD, AMBER, NAMD, own codes]

Solids – The study of systems that are of large or infinite spatial extent.

As a large number of industrial catalysts are solids, the understanding of the way these function and the design of new catalysts is an important area of research. In addition, the study of the binding of pollutants and undesirable species to the surfaces of solids as a method of removing them from the environment is vital. Another area of interest is defects on the surface of solids, as these can be places that are highly reactive. The potentially large number of atoms involved in these studies requires massive computing resources. [Programs used: CRYSTAL, VASP, GULP, QM/MM]

Liquids – Models of solvation and the study of pure substances.

Many chemical reactions take place in the liquid phase with the reacting species being different from the solvent that constitutes the bulk of the liquid. However, sometimes the solvent actively participates in the reaction but is regenerated or released by the end chemical process and hence gives the superficial impression of just being a medium for the reaction. Modelling reactions with explicit classical or quantum mechanical water molecules present can give detailed insight into the processes taking place. Using water as a solvent at raised temperature and pressure (supercritical water) leads to a solvent medium with very different properties from normal water. One way of studying chemical reactivity in this form of water is by Car-Parrinello dynamics. The program used for these calculations (CPMD) runs very well on massively parallel machines but requires vast amounts of time. Other types of liquid of interest are liquid crystals, used commercially in display devices, which have the property of a fluid, but retain much more order of the molecules with respect to each other than a normal liquid. [Programs used: DL-POLY, CPMD and own codes]

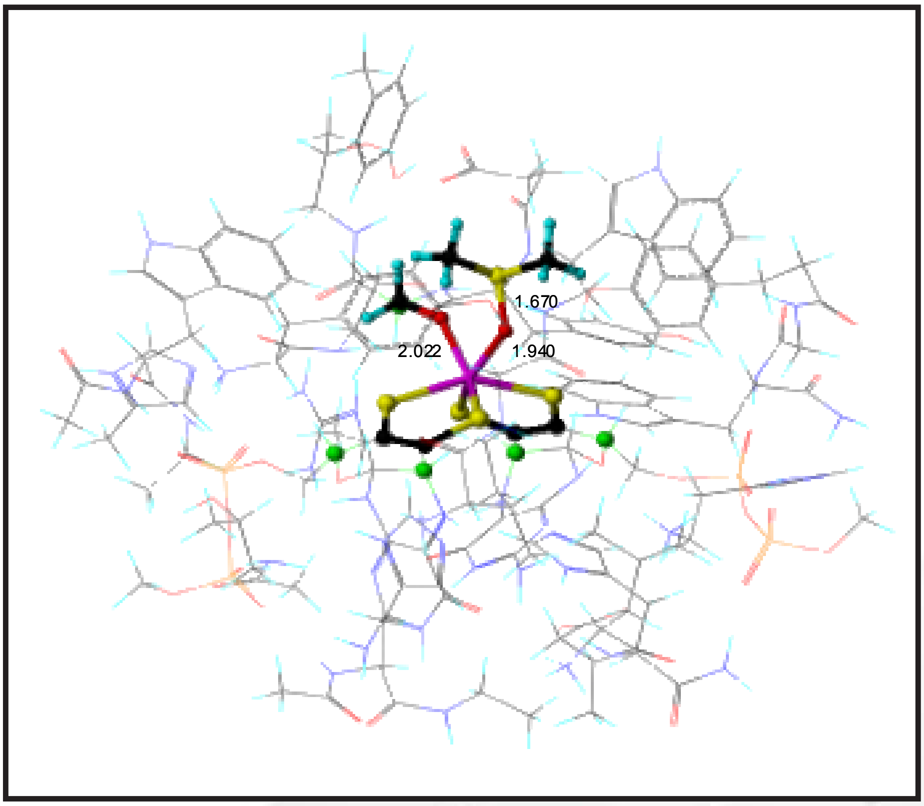

Figure 1: The active site of the enzyme DMSO reductase from the work of Dr Jonathan McNamara, Dr Matthias Mohr and Professor Ian Hillier using the crystal structure of L. J. Stewart et al, J. Mol. Biol., 2000, 299, 593.

Gas Phase – Isolated molecules or clusters of molecules.

Calculations performed on isolated molecules or clusters of molecules in the gas phase are free of complications due to periodic boundary conditions or solvents. A single molecule or a small cluster of molecules can be studied in detail and because they are small in size sophisticated theoretical approaches can be employed. However, these methods often require large amounts of memory and disk space. One use of the information from these studies is to obtain parameters for more approximate levels of theory (the so called semi-empirical theories). A genetic algorithm determines a set of parameters that can describe the electronic structure of a system whose size makes it prohibitive to study by more rigorous methods, but whose reactivity demands a quantum mechanical description. (It should be noted that the faster classical mechanics often cannot describe bond breaking and making processes.) The genetic algorithm code runs well on parallel machines, but requires large amounts of time.

We have also undertaken studies into the nature of the interaction between two fragments of molecule. In this we split the molecule into two parts and see how they interact. This provides information on the nature of the bonding between the fragments. While this does not necessarily require supercomputer facilities, due to the size of systems we have studied parallel computers are needed to yield results in a reasonable time frame. Other studies involved developing methods of obtaining wave functions that include electron correlation and looking at metal complexes that are potential molecular magnets (i.e magnets the size of a single molecule). [Programs used: Gaussian, GAMESS-UK, CADPAC, locally written BOVB and Manchester University Semi-Empirical program (MUSE)].

This brief summary of theoretical chemical activity at The University of Manchester indicates that we study all of the common phases of matter and the chemistry taking part in them. A variety of computer codes are employed in our research, some of which are 'home grown'.

Acknowledgements

I am thankful for comments from Professor Ian Hiller, Professor Jonathan Connor, Dr Joe McDouall and Dr Andrew Masters.

The HLRN System

Peter Endebrock
HLRN / RRZN

The HLRN (Hochleistungsrechner Nord, High Performance Computer North) is a joint effort of six North-German states. The HLRN has two installation sites, one in Hannover at the RRZN, the other in Berlin at the ZIB, but for the user it is intended to look like a single system. Each site is presently equipped with 13 (in the future 16) IBM p690 series 32–way SMP frames, and the two sites are coupled via Gigabit Ethernet. The HLRN offers high-performance computing service and support to the North-German universities and other scientific institutions. A "network of competence" with consultants from the participating institutions provides the support for the users.

The six North German states Berlin, Bremen, Hamburg, Mecklenburg-Vorpommern, Niedersachsen and Schleswig-Holstein decided in the late 1990s to start a joint venture to install a high-performance computer system. The reason was that none of the states was able to finance such a system by itself, but that combining efforts might make it possible to get a high-performance computer system for Northern Germany. That was supported by positive experience from previous cooperation of Berlin, Niedersachsen and Schleswig-Holstein in the NVV, the North German Vector Computer Cooperative.

The idea was (with some hesitation and after a lot of negotiation) approved by the German federal funding and review organizations, and the system was installed in the second quarter of 2002. The concept is that this should be a "single supercomputer at two sites". Identical systems (except for the file archiving system) have been installed in Hannover and Berlin, and for the user they are intended to look like a single system.

The present hardware equipment at each installation site consists of thirteen 32-way IBM p690 frames with Power4 processors connected by the SP2 ("Colony") switch. Three more of these frames will be installed by mid-2004 together with the "next generation" ("Federation") switch. In addition there are eleven p655 servers for I/O and archiving on each site. You can find details of the present configuration of one site in Figure 1.

The summarized data of the two sites are:

- 26 x 32 CPUs with 1.3 GHz, equivalent to a theoretical peak performance of c.4.4 Gflops
- Main memory 22 x 64 MB, 2 x 128 MB and 2 x 256 MB, together c.2.2 TB
- Total mass memory c.52 TB
- Total archive storage c.3 PB

The two sites are coupled by a dedicated fibre network with a bandwidth of 2.4 Gbit/s.

The user view of a "single system image" is still under development, but a lot of features in that direction have already been realized. The system and application software at both complexes are identical, except for minor differences due to licensing restrictions. Each user has a single user name and password at both sites, and has an identical home directory available. The directories for permanent and medium-term files are presently only available directly at the user's primary site, and they have to make file transfers to the secondary site, if necessary, on their own. One batch system (IBM LoadLeveler) manages jobs for both sites, with the possibility of sending jobs to a pre-selected site, or leaving that decision to the system. Actually, as mentioned before, the user has to arrange for some of their files to be available where the job executes, but we hope to improve that to a certain degree. Currently, only test jobs have been allowed to use CPUs from the two sites in one job, but it has been shown that that is possible in principle. As you can imagine, due to latency and transfer capacity, this will not be a sensible method of system usage in general, but it may prove to help with the solution of large loosely-coupled systems.

A challenge for the two installation sites was the cooperation of two computer centres each with a long history of their own, to develop into a single centre with a common policy. A consolidation phase was necessary because of different "cultures", but most of that is history now. Regular video conferences support the cooperation, and the combined manpower of the two sites was and is needed to support the system – one site wouldn't have had enough people to do it by itself.

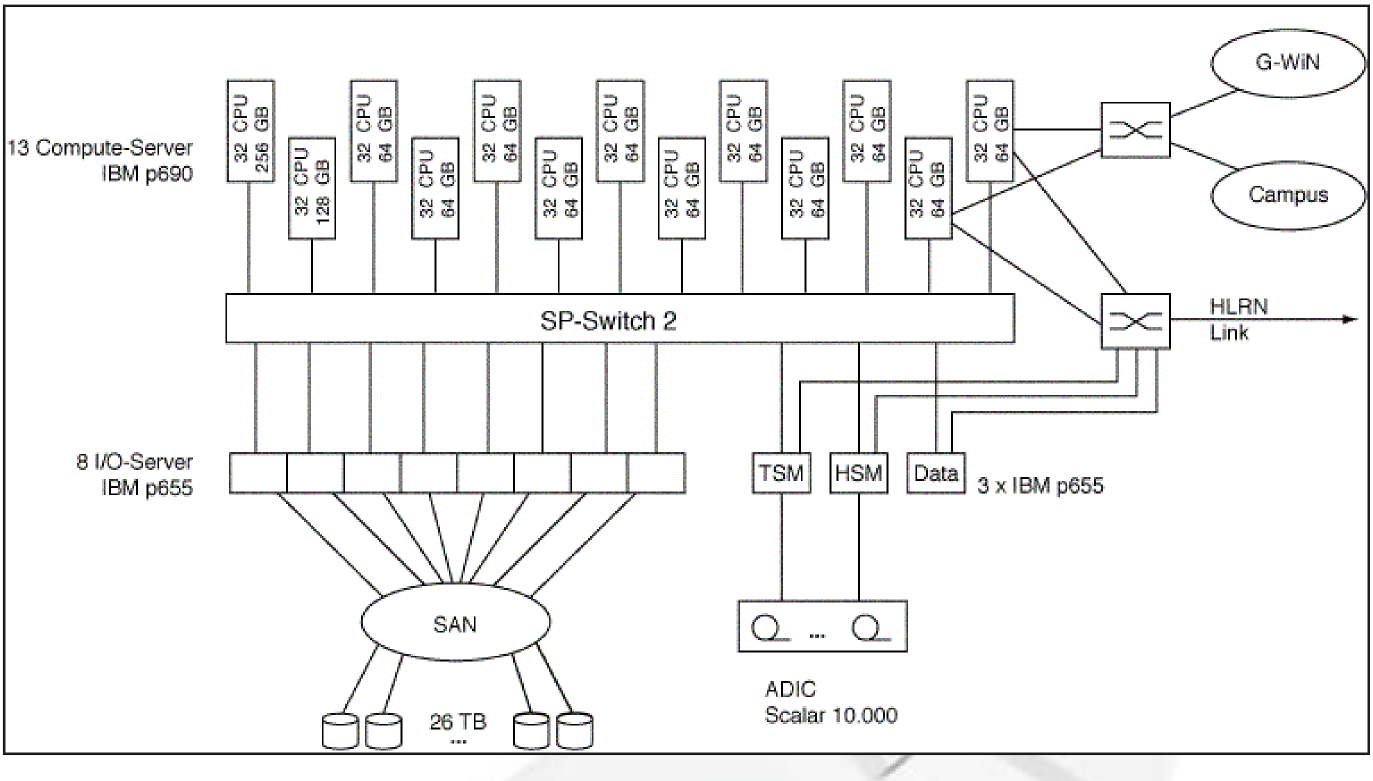

Figure 1: One half of the HLRN system (the one at Hannover; the archive system at Berlin is different)

The main application areas are environmental research, climate and ocean modelling, and various engineering applications like ship building or computational fluid dynamics, as well as fundamental research in physics, chemistry and life sciences. Many users have their own application codes, but standard application packages in different areas are also offered and used.

The initial user name application is relatively informal and does not require detailed argumentation. You get a small resource allowance to prepare your project application. The projects applying for using the system are reviewed by a Scientific Board that allocates resources on a quarterly basis, for a maximum of four quarters in advance. For the first application period, starting with the second quarter 2003, the total applications amounted to three times the available resources, an indication of how badly the HLRN system had been needed.

The system shows presently a usage of more than 60 percent of the available resources, without calculating downtimes. We consider that to be good for a large parallel processing system with a mixed workload and a preference for large projects.

User support is organized by a network of consultants distributed over the participating states and their local computer centres. Each project is assigned an individual consultant, preferably with experience in the field of the project and geographically close to it. If that is not possible, teams of two consultants are assigned to a project, the "local advisor", and a remote "specialized" consultant.

Figure 2: The installation phase of HLRN at Hannover

In our opinion, the concept of cooperation between several federal states to install and run a "distributed supercomputer" is proving its feasibility and its advantages, and because of the long planning cycles we are already starting to prepare a request for a similar successor system.

If you want, you can find more details about HLRN at http://www.hlrn.de/.

Op3D-Grid

Joanna Leng
SVE Team, Manchester Computing, University of Manchester

Building a Grid-enabled surgical visualization system for use in the operating theatre

Although 3D visualization has been available for some time it has not proliferated into working environments as well as might have been hoped. There are several reasons for this:

- Large data sets need specialist software/hardware to produce images in real-time.
- Manipulating and interacting with the data is difficult and requires practise.
- Integrating the use of a visualization system with other working practises is not always easy.

These problems have all been solved by the Op3D software so that a surgeon can view and interact with 3D images of medical scans in the operating theatre while performing an operation.

Figure 1: Surgeon works at light box in operating theatre.

Figure 2: Surgeon now works with 3D images in operating theatre.

Background

Op3D is a remote visualization system that was developed in collaboration between Manchester Visualization Centre (MVC) and Manchester Royal Infirmary and has been used on many patients throughout 2002 and 2003. MVC was a finalist in the 2003 Computer World Honours 21st Century Achievement Awards for this work.

Figure 3: Medal of Achievement.

The original Op3D project was innately a grid application that relied on the use of SGI's visual area networking technology. More recently a further collaboration, this time including SGI and funded by the ESNW, has resulted in a truly grid enabled version of Op3D.

Technique

The medical data is placed on a high end SGI Onyx system. The Op3D software uses OpenGL Volumizer and OpenGL Performer to produce high quality 3D renderings of the data in real time. OpenGL Vizserver compresses and transfers images to a laptop in the operating theatre across a 100BaseT connection. A projector is used to display the images from the laptop onto the wall. The surgeon interacts with the 3D rendered images throughout the surgery with a joystick that is coated in a sterile plastic bag. Op3D uses a specially defined class of interactions that make its use as intuitive as possible.

Op3D-Grid uses patient specific data, which is stored in DICOM format on a hospital database. Secure access to this data is important so a java based User Interface has been developed that uses globus in the form of the java based cog kit to securely transfer the data to the

Onyx machine. Here the data is anonymised and converted into the appropriate format using the dctk (dicom tool kit) and various scripts. A run script is also produced that sets the Op3D application to run with this data and preferences of the surgeon.

Op3D uses a profile of user interactions. The surgeon can manipulate the data and a cut plane through a series of single button interactions, icons are used to show which interaction is current. The new set of interactions include book marking which allows the surgeon to mark and go to important scenes at the push of a button. More importantly a new type of interaction has been developed which uses "physically" accurate interactions to move the camera and alter the surgeon's view. This type of interaction is still being tested but it is hoped that this will increase the usability of the system and make it easier to locate features in the data.

The graphics hardware on an SGI Onyx is being used extensively to provide real time volume rendering of patient data. Currently the server provides rendering capabilities and little processing. However, Op3D-Grid will provide the infrastructure for more functionality to be delivered to the Operating Room e.g. physiology simulation.

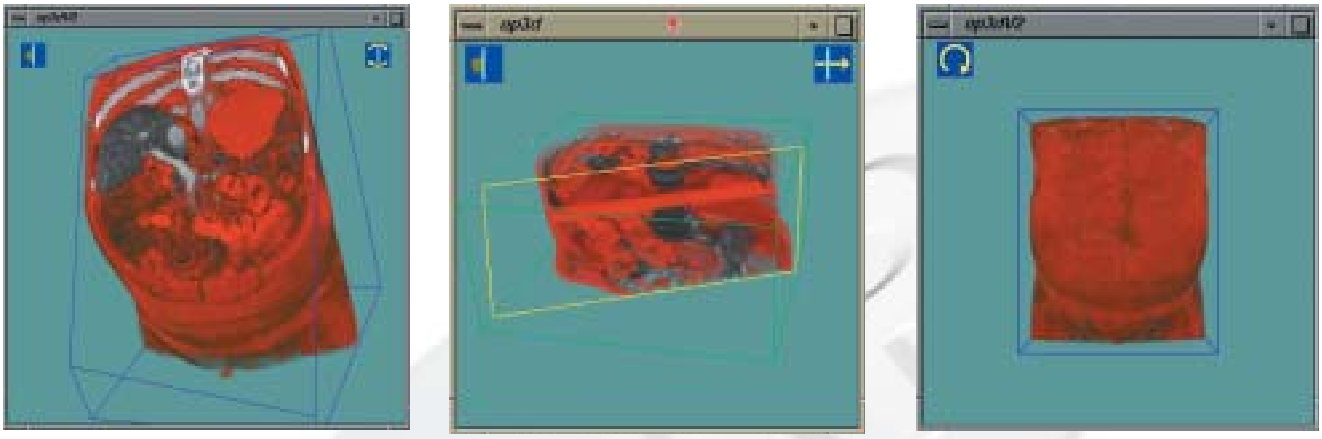

Figure 4: Slices showing 3D views of patient data.

CRAY
Europe
advanced
ECHNICAL
WORKSHOP
2004

2004 Cray Advanced Technical Workshop - Europe -

June 17-18th, 2004
Bologna, Italy

Registration is free!!
To R.S.V.P., please register at:
http://www.cray.com/catw , send email to
catw_europe@cray.com or
call Julie on +44-118-903-3178.

We look forward to seeing you in Italy!

Hosted by CINECA in picturesque Bologna, Italy, please join us at our *2004 European Cray Advanced Technical Workshop* where you will hear directly from users and Cray experts about the scientific advancements currently underway using Cray systems. You will also get a preview into the technical details on current and upcoming Cray products such as the X1 and programs like ìRed Stormî,ìStriderî, and the future OctigaBay offerings.

Cray is focused on designing balanced system architectures - matching fast processors with fast interconnects at high memory bandwidth to support the most advanced HPC applications. This workshop will provide insight into how you can best use the power of Crayís systems to meet your scientific goals.

Featured Speakers already include:
Sandia National Laboratories,
Oak Ridge National Laboratories,
ICM (Poland),
Los Alamos National Laboratories,
INM (Spain),
OctigaBay,
Others -TBD

Igniting Innovation Across the Globe - SC Global 2003

Michael Daw, (Technical Director of SC Global 2003), SVE Team, Manchester Computing, University of Manchester

No longer do conferences have to be limited to a single physical location. SC Global is the distributed element of the Supercomputing (SC) conference, which focuses on high-performance computing and networking. Speakers at SC Global participated from their home institutions in seven different countries, five sovereign tribal nations (these are what used to be known as the American Indian nations - see http://www.bluecorncomics.com/sovreign.htm for an interesting history), distributed across four continents. It was, in the words of Chuck Koelbel, co-Chair of the SC Technical Program, "…clearly one of the most successful parts of the program…" The underlying technology that made all this possible was Access Grid.

Access Grid is videoconferencing on steroids. Typically, Access Grid nodes use a whole wall for display to give the pixel-space to show life-size images of the potentially large number of other participants in any meeting or lecture. They have multiple cameras for close-up views and to support many attendees and have excellent audio that is a pre-requisite for effective communication. But where Access Grid really wins over traditional videoconferencing is the capability for a huge number of sites to take part simultaneously. This gives the technology the potential to support a truly global, distributed conference where participants do not need to travel to take part.

Manchester Computing played a highly significant role in SC Global by taking the lead role in three out of the 11 sessions, ably steered by Kevin Roy at the controls. SVE's Lee Margetts (who joined from Spain's first Access Grid node that he helped to set up) combined with Simon Bee from Salford University to showcase Hydra, a software framework that can deploy and control many interconnected applications, such as Computer Aided Design (CAD) and Finite Element Analysis (FEA) at geographically remote locations. The Manchester Access Grid node also hosted Kelli Dipple, who presented a completely different approach to Access Grid, taking it out of the sometimes staid world of science and showing how the technology may be used as an exciting medium for art.

One of the highlights of the conference was the session "Application Steering in a Collaborative Environment". Three groups, from Jülich, Stuttgart and RealityGrid showcased various techniques for integrating application steering and visualization into the collaborative Access Grid environment. The RealityGrid part featured a live demonstration in which geographically separated scientists migrated running lattice-Boltzmann simulations from systems (CSAR and HPCx) on the UK e-Science Grid to the US TeraGrid, steered them into new regions of parameter space, and monitored their state through real-time visualization using remote systems in Manchester and Phoenix. This demonstration was just one facet of the ambitious TeraGyroid project, which harnessed the combined resources of the US TeraGrid and UK e-Science Grid to produce a major leap forward in soft condensed matter simulation.

The climax of the conference was a demonstration of distributed Karaoke from AIST, Japan. Much to my amazement, there were plenty of volunteers to sing - highlights included a rendition of 'Living on a Prayer' by Bon Jovi, complete with enthusiastic air guitar from revellers in Boston, and a frighteningly lifelike Michael Jackson impression by Donnel Sanders (known as 'Biggie Don') from Winston Salem State University during the group sing-along 'We Are the World'.

The conference as a whole was made possible by a complex global technical infrastructure distributed across many institutions to decrease lags and increase stability and robustness. The major elements of this infrastructure comprised venues servers, multicast-unicast bridges, presentation software servers, Question Tool servers and recording servers.

There are many software routes that may be taken by sites to utilise Access Grid. The newest release of the research project software is known as AG2. However, AG2 has some stability issues, so the decision was made to allow remote sites to connect to SC Global using any of the available methods - AG2, AG1 (the original, prototype Access Grid software), inSORS (the commercial version) or bypassing the venues servers altogether and connecting directly using vic and rat, the video and audio Mbone tools. Five venues servers were established to support this using hardware sited in different institutions.

Video and audio streams in Access Grid are distributed using multicast, a set of network protocols that allow for efficient utilisation of bandwidth. However, multicast is by no means ubiquitous and in some parts of the network is not reliable. Therefore, bridges are required to support sites that do not have good multicast connectivity. For SC Global 2003, we established 12 such bridges, based in nine institutions in Australia, Korea, the UK and the US.

different locations and putting hands up for questions is not effective. The audience submitted questions via a web interface, which also allowed for submitters who did not have access to a microphone. This software was specially written for SC Global.

The event was recorded using another inSORS product, IG Recorder, that allows the event to be played back over the Access Grid at a later date. A Windows Media stream was shown to allow people who did not have access to an Access Grid node to view the event. This was also recorded.

Figure 1: SC Global 2003 in Phoenix. On the platform (L to R): Jennifer Teig von Hoffman (Chair SCG03), Jim McGraw (Conference Chair), Jackie Kern (Chair SCG04)

Many of the sessions involved presentations of PowerPoint slides. These were distributed to remote sites using an inSORS product known as IG Pix. This gives the capability for the local site to deliver a presentation as normal on a laptop; remote sites can view the slides via a web page that is updated as the different slides are shown.

The Question Tool allowed for question management - always a difficult issue when your audience is in many

SC Global was a thoroughly enjoyable, successful and high profile addition to SC that is sure to continue to make an impact in coming years.

Further Information

For more information, see http://www.sc-conference.org/sc2003/global.html, http://www.insors.com, http://www.accessgrid.org, or you can contact the author at michael.daw@man.ac.uk.

Recent Events

SC2003 Phoenix

Pen Richardson, HPC Team, Manchester Computing, University of Manchester

Each year the Supercomputing conference brings together high performance vendors, research groups and the user community. SC2003 attracted 7,641 registered attendees from 43 countries, with participants from outside the U.S. constituting 16 percent of total attendance. SCGlobal which proved a huge success in 2001 was also back with a large program of events (see pages 20-21 for details).

For the fifth year Manchester Computing had a booth focused on the Supercomputing, Visualization and e-Science group (SVE), within whose responsibility the management and maintenance of CSAR falls. Representatives from SVE manned the booth for the 3 exhibit days discussing all aspects of the group's work with other members of the supercomputing community.

For the group as a whole the conference was an enormous success involving people both on site in Phoenix and back in Manchester. Here are some of the highlights:

- Adrian Tate assisted with a tutorial on Co-array Fortran.
- SGI BOF on visualization.
- Mike Daw was the Technical Director for the SC Global event on collaborative working.
- Lee Margetts led a showcase demonstration on virtual prototyping at SC Global.
- Terry Hewitt was the 'Master of Ceremonies' of two SC Global events. Kaukab Jaffri also chaired a session.
- Presentations on SGI and UK e-Science stands about CSAR, Reality Grid and Triceps.
- Further presentations on ANL, NCSA, PSC and CALTECH stands.
- Major involvement in Teragyroid project (http://www.realitygrid.org/TeraGyroid.html)
- HPC Challenge competition winners in the 'Most Innovative Data-Intensive Application' category with TRICEPS which involved resources from CSAR, HPCx, Pittsburgh, San Diego and Argonne National Labs in Chicago.
- Partners in the second HPC Challenge Category 'Most Geographically Distributed Application' led by HLRS Stuttgart and Craig Stewart of Indiana.

For more information on the Supercomputing Conference visit www.sc-conference.org Next year will be Pittsburgh.

Our next conference will be ISC2004 (www.isc2004.org) in Heidelberg which takes place this June (23rd-25th). The conference will give a comprehensive overview on HPC topics presented by leading, world-renowned scientists and experts. The conference exhibition with more than 50 booths and 30 exhibition hours ranks as the number one European HPC event. Please visit us at exhibition booth E06.

Launch of Newton

Andrew Jones, HPC Team, Manchester Computing, University of Manchester

A dinner event was held at the London Art House on 26th January to celebrate the successful launch of Newton - CSAR's new SGI Altix system.

The event was attended by around 60 people including leading academic users and PIs of the service, commercial users, EPSRC, NASA, and representatives from across the CSAR partnership - SGI, CSC and the University of Manchester.

The event included three after dinner speakers. Professor Peter Coveney (PI of the RealityGrid project at UCL, a major user of both CSAR and HPCx). Dr Eng Lim Goh (Chief Technology Officer at SGI) outlined the highly promising future of SGI's research into next generation supercomputers. Walt Brooks (Head of Advanced Computing Division at NASA Ames), talked about the importance of supercomputing to NASA's mission, and the role played in investigating the recent shuttle disaster. NASA Ames are closely related to CSAR in that their main computing systems are a 1024 processor Origin and a 512 processor Altix - that is the same systems as CSAR, just twice the size!

Overall, the event was felt to be a positive success, and similar events are being discussed for the future.

CSAR User Survey 2003

Claire Green
HPC Team, Manchester Computing, University of Manchester

The CSAR User survey for 2003 was conducted between 26th November and 12th December 2003. The survey was publicised to all CSAR users through the monthly CSAR Bulletin and by an email from Deborah Miller of EPSRC. 7% of CSAR users completed and submitted the survey via an online form with questions on topics ranging from dealings with CSAR staff to the introduction of the new Altix system.

Newton, the new Altix system, was introduced into the service in October 2003. We asked various questions about this technology refresh in our survey. 93% of respondents felt that Newton had been well-advertised and had been advertised at the right time. 23% had tried out Newton or Reynolds as early access users and 72% of those who answered were aware of the free porting assistance provided by CSAR to enable users to get started on Newton.

On using the CSAR systems users were most satisfied with temporary disk space - 93% very or fairly satisfied. The least satisfaction was with job scheduling on the systems with 84% very or fairly satisfied. We also asked if users had attended the CSAR training courses and had found them useful - 100% had. Of those who had not attended training provided by CSAR, 65% had not because it was not required, with the remaining being equally split between taking part in training elsewhere or for other reasons, for example learning from other members of their consortium.

61% of those who answered the questions on code efficiency are aware of how efficiently their code is running. 29% are not aware of their code's efficiency but would be interested in their code being analysed to determine this. The remaining 10% are not aware of how efficiently their code is running and are not interested in having their code analysed. (Please contact csar-advice@cfs.ac.uk if you would like further information regarding this).

The majority of users are satisfied with the way they have been dealt with by CSAR staff. The feedback mechanisms that have been used are very similar to previous years with most users having used the CSAR Helpdesk, followed by the CSAR management team. Users are more satisfied with the responses they have received in 2003 with 100% viewing the response as acceptable as opposed to the 94% who believed they had received an acceptable response in 2002.

Overall 91% rated the overall level of the High Performance Computing service as Very Good or Good. The remaining respondents viewed the CSAR service as Adequate, with no one viewing it as Poor or Very Poor. All of the users who took part in the survey agreed that using CSAR had contributed to advancements in their research.

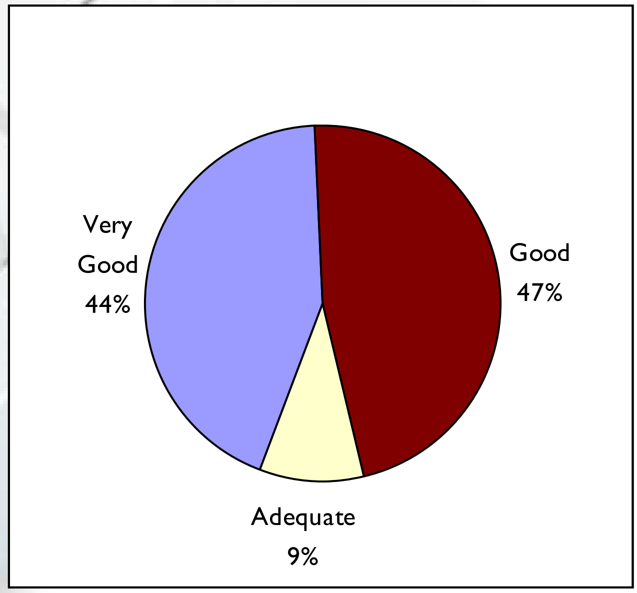

Figure 1: The overall rating of the level of HPC service provided by CSAR.

Further Information

To view the full report from the CSAR User Survey for 2003 please visit http://www.csar.cfs.ac.uk/admin/reports/user_surveys/ Reports from the User Surveys for previous years (dating from 1999) can also be viewed via this webpage.

Please contact csar-advice@cfs.ac.uk if you have any comments regarding the issues mentioned in this article or should you require further information and/or assistance in using the CSAR systems.

Newton: A Short History

Gerry Todd
HPC Services Manager, CSC

….setting my heart on money, learning, and pleasure more than Thee…..

Sir Isaac Newton (1642-1727), mathematician, physicist and politician, one of the foremost scientific intellects of all time. Born in Woolsthorpe, near Grantham in Lincolnshire on Christmas Day 1642, according to the Julian calendar that was the measure of the time; this is however corrected in his biographies to 4th January 1643 by the Gregorian calendar.

Isaac came from a family of farmers but never knew his father, also called Isaac, who died in October 1642, three months before the birth of his son. His father, though a relatively wealthy man owning both property and livestock, was uneducated and could not even sign his own name.

Isaac spent his early years in the care of his grandparents, even though his mother had remarried the local minister, one Barnabas Smith. Isaac grew to resent this treatment and harboured a great resentment towards his mother and stepfather. When examining his sins at the age of nineteen he wrote:-

….Threatening my father and mother Smith to burn them and the house over them….

Figure 1:
Sir Isaac Newton 1642-1727.

In 1653 Isaac began attending the Free Grammar School in Grantham; however he appeared to show little promise in academic work, his teachers describing him as ‘idle’ and ‘inattentive’. At this time his mother, by now a lady of reasonable wealth and property, took Isaac away from school to manage her affairs, a task in which young Isaac showed no talent or interest .

An uncle, William Ayscough, decided after discussion with his mother that Isaac should be made ready to enter university, so in 1660 Isaac returned to the Free Grammar School in Grantham to complete his studies. During this time he lodged with Stokes the headmaster of the school, and began to show academic promise. It is possible that at this point in his academic career Stokes introduced him to Euclid's Elements.

Isaac entered Trinity College Cambridge on the 5th June 1661, as a sizar, this being a student who receives an allowance towards college expenses in return for acting as a servant to other students.

Isaac's aim at Cambridge was a law degree. This involved the study of philosophy, primarily Aristotle, but also Descartes, Gassendi, Hobbes and Boyle. He was also attracted to the astronomy of Galileo and studied Kepler's Optics.

In 1663 Isaac began his interest in mathematics after he bought an astronomy book at a fair in Cambridge, which upon reading he discovered he could not understand the mathematics it contained. So began his passion for mathematics, reading Oughtred’s Clavis Mathematica, Descartes’ La Geometrie as well as other contemporary texts.

Isaac was elected a scholar in April 1664 and received his Bachelor’s degree in April 1665. His scientific genius had still not emerged; however it did so suddenly after the closure of the University due to the plague. This event forced him to return to Lincolnshire where, in a period of less than two years, he began revolutionary advances in mathematics, optics, physics and astronomy. He started to lay down the foundations for differentiation and integral calculus, several years before Liebniz made the same discovery.

When the University re-opened in 1667 after the plague, Newton put himself forward as a candidate for a fellowship and was awarded a minor fellowship at

Trinity College. After attaining his Master's degree he was elected to a major fellowship in 1668.

In 1669 the then Lucasian Chair, Isaac Barrow, resigned to devote himself to divinity. As Barrow had taken an interest in Newton's progress, he recommended that the 27-year-old fellow be appointed in his place.

Newton's first work after his appointment as Lucasian Professor was optics, the topic of his first lecture course in 1670, and contained his conclusion, reached during the plague years, that light is not a single entity. This discovery was made whilst using a telescope that had a chromatic aberration in the lens; subsequent experiments with a prism noting the spectrum it produced verified this. This went against current scientific thought as it was believed, and had been since Aristotle, that white light is a singular entity. Newton argued that white light is a mixture of many different rays of light, which are refracted at differing angles when passing through a refracting lens. Newton reasoned that a refracting telescope would always suffer a chromatic aberration because of this; he therefore proposed the construction of a reflecting telescope.

Newton's greatest work in physics and celestial mechanics in 1666 was his theory of universal gravitation containing the Three Laws of motion.

In 1672 Newton was elected a fellow of the Royal Society after donating a reflecting telescope, his first scientific paper on light and colour also being published in this year. The paper was in general well received; however Hooke objected to Newton's attempts by experiment alone to prove that light consisted of small particles and not waves. This was the beginning of a bitter feud with Hooke and led in part to Newton's first breakdown in 1678, assisted by a violent exchange of letters over his theory of colours with an English Jesuit and culminating in the death of his mother in 1679.

Newton also had a political career whose roots may be traced back to 1685 when James II became King of England. James became a convert to the Catholic church in 1669 from which point he began placing Roman Catholics in high offices in the army and then as judges and officers of state. He then proceeded to appoint Roman Catholics to Cambridge and Oxford as positions became vacant. Newton, a staunch Protestant, saw this as an attack at the very heart of the University of Cambridge. When James tried to have a Benedictine Monk awarded a degree without taking any examinations or swearing any oaths, Newton wrote to the Vice-Chancellor:-

Be courageous and steady to the laws and you cannot fail.

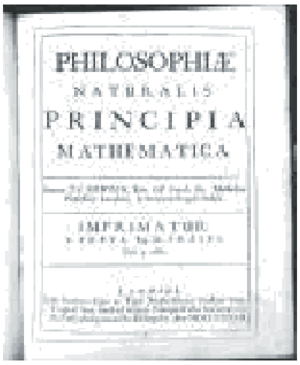
PHILOSOPHIÆ
NATURALIS
PRINCIPIA
MATHEMATICA
IMPRIMATUR

Figure 2: One of Newton's works.

The Vice-Chancellor did and was dismissed. Newton was not discouraged and continued to argue the case strongly, gaining such support that when James fled the country in 1688 after William of Orange had landed, Newton was elected to the convention parliament to represent the university the following year. This allowed Newton to see life outside academia such that after his second breakdown in 1693 he retired from research.

Newton left Cambridge to take up a position in government which led to him becoming Warden of the Royal Mint in 1696 and Master in 1699, but did not resign his position at Cambridge until 1701. To most people, attaining this position would have been seen as reward for their scientific achievements, however Newton continued to actively contribute to the work of the Royal Mint. He worked on introducing measures to prevent counterfeiting of the coinage.

In 1703 he was elected President of the Royal Society and was re-elected every year until his death. Queen Anne knighted him in 1705.

His latter years were spent revising his major works, polishing his studies of ancient history, defending himself against his critics and carrying out his official duties. Newton has been regarded for almost 300 years as the founding father of modern physical science, his achievements in experimental investigation and mathematical research being equalled if not surpassed by the energy and enthusiasm that he plunged in to chemistry (and alchemy), the early history of Western Civilisation, and theology. The study of the form and dimensions, as described in the Bible, of Solomon's Temple in Jerusalem was one of his specialist topics.

Newton never married and was a modest man, generous to his friends and harsh towards his detractors. He was buried with great pomp and ceremony in 1727.

MRCCS/NSF Summer School

Graham Carey
Director, CFDLab, Institute for Computational Engineering and Sciences (ICES), The University of Texas at Austin

During the week of September 1-5, 2003 Manchester Research Centre for Computational Science (MRCCS) hosted a summer school/workshop entitled "High Performance Computing in Finite Element Analysis". The goal was to bring faculty researchers experienced in parallel computing, algorithms finite element methodology and applications together with graduate students in this topical area. The workshop was comprised of three main components: (i) formal lectures by experts on several inter-linked topics covering both fundamental and state-of-the-art developments; (ii) demonstrations and application studies to provide "hands on" experience with tutorials where applicable, and (iii) informal discussion to encourage exchange of ideas among faculty and students.

This proved to be an excellent format since it provided an 'accelerated entry' for beginning graduate students whereas students further along in their research could discuss deeper technical issues with experienced faculty researchers. It also facilitated an interchange between faculty carrying out research and supervising graduate students in this topical area. This interaction at Manchester will no doubt lead to further contact, exchange of papers and possible collaborations between individuals at the various research groups represented.

It seems especially appropriate that this meeting was convened at the University of Manchester. The era of modern computing machines began in the middle of the last century with the development of the first generation of scientific computers in several countries. One of the earliest machines was designed and built at the University of Manchester. The small-scale experimental machine SSEM, or "Baby" as it was known, made its first successful run on June 21st 1948. In the intervening half century, the University of Manchester has maintained a prominent position in scientific computing. This position was formalized in 1969 with the opening of Manchester Computing (MC), and a succession of supercomputer generations have been in use at MC since that time. The September summer school program, mentioned in the opening paragraph, followed in the steps of this tradition.

First, one may reasonably comment that there are only a handful of truly major scientific developments in the past century and the development of electronic digital computers must stand with these few. The continuing revolution in microelectronics and the ability to mass produce the associated products as a commodity has reaped unimagined price-performance benefits so that computing is now pervasive throughout most areas of society. In particular, it has enhanced our simulation and problem solving capability. Previously, experiment and theory were regarded as the two pillars of science but today computer modeling stands as an equal partner in this endeavor.

The term "supercomputers" has been applied to refer to the leading high performance computers in any given 'technology generation'. Generally, supercomputers may be defined as current generation machines that are capable of solving, in a timely manner, problems that were significantly beyond the scope of previous computing technology. Since a computer generation here implies only a few years, this is a "rapidly moving target," and it has certainly been the case that the supercomputer a decade ago is surpassed by the desktop computer today. Consequently, the term supercomputing has recently been replaced by High-Performance Computing or Advanced Computing to better reflect the current state of the art.

Most high-performance computing today relies on parallel systems comprised of inexpensive commodity-off-the-shelf (COTS) processors with very fast communication hardware. However, high-performance computing is much more than the hardware alone. Amongst many other items it requires, in particular: (i) algorithms and software for partitioning very complex problems across parallel architectures containing many networked processors; (ii) methods, algorithms and software for scientific problem solving via numerical analysis, and related tasks; (iii) parallel software infrastructure, libraries and toolkits; (iv) storage for extreme volumes of data with data mining/ manipulation and (v) advanced visualization capability.

These key topics and other related issues were the subject of the summer school program. However, they were specifically framed in the context of Finite Element Analysis, so a few comments on the finite element method are also relevant. In a practical sense, finite elements also began in the engineering community approximately 50 years ago. The method was initially developed for structural analysis using the natural idea of combining contributions of structural members or "elements" into an "assembled" structure. Later, the relation to variational principles, partial differential equations and extension to more general applications was gradually realized. Above all else, it is the integral formulation that allows the use of unstructured and graded meshes of general elements that underlies the applications power of the finite element method. The idea of partitioning the mesh and parallel processing over collections of elements has led to an easy transition of finite element analysis into the parallel high performance computing arena while losing none of the generality of the method. The MC summer school workshop was designed to explore high performance computing in this context. The program of lectures on the subject was (in alphabetic order):

Professor Jacobo Bielak, Carnegie Mellon University, USA - *Parallel FEA*
Professor Graham Carey, University of Texas, USA - *Parallel FEA Algorithms*
Professor Mark Cross - University of Greenwich, UK - *Techniques and Tools for Parallel FEA*
Dr Jon Gibson, University of Manchester, UK - *Introduction to HPC*
Professor Boris Jeremic, University of California, Davis, USA - *Parallel FEA in Geomechanics*
Ms Jo Leng, University of Manchester, UK - *Introduction to Visualisation*
Dr Lee Margetts, University of Manchester, UK - *Parallelisation*
Dr Kengo Nakajima - GEOFEM/Earth Simulator Project, Japan - *GEOFEM / Earth Simulator Project*
Dr Mike Pettipher, University of Manchester, UK - *Performance Measurement and Optimisation*
Professor Olivier Pironneau - University of Paris, France - *Schwarz and Schur Algorithms*
Professor Mark Shephard - Rensslaer P I, New York, USA - *Parallel Automated Adaptive Analysis*
Professor Ian Smith - University of Manchester, UK - *HPC in FEA, Parallel EBE, Practicals*
Professor Nigel Weatherill - University of Wales Swansea, UK - *Parallel Mesh Generation*

The lectures were followed by organized software and applications demonstrations/tutorials/discussions. Details of the lectures and these sessions are available on a DVD. (Please contact Fiona Cook by email - fiona.cook@man.ac.uk - if you would like a copy of the DVD).

The technical exchanges continued into the coffee breaks and after hours: extra-curricular highlights of the meeting were (not necessarily in order of importance): the after- house social hour and discussions; the conference dinner at "The Ox" pub (one recalls that the algebraist Sylvester worked all night on binary forms with the aid of a decanter of port to sustain flagging spirits); and, last but not least, the surprise wake up "discussion session" on the college courtyard before early morning due to an errant fire alarm.

Figure 1: Enjoying the Conference Dinner at "The Ox".

Special thanks are due to the organizers, Professor Ian Smith, Dr Mike Pettipher, Dr Lee Margetts, Dr Jon Gibson, Jo Leng and Fiona Cook at Manchester Computing and to Boris Jeremic at UC Irvine, USA for planning and carrying out an excellent summer school program. Organizers, lecturers and students also thank the US National Science Foundation for providing financial support for the workshop.

CSAR,
Manchester Computing, Kilburn Building,
University of Manchester, Oxford Road,
Manchester. M13 9PL
Tel: +44 161 275 6824/5997
Fax: +44 161 275 6040
E-mail: csar-advice@cfs.ac.uk

CSAR Information

The CSAR Website - http://www.csar.cfs.ac.uk - contains help and information on all aspects of the service, including sections on Software, Training, Registration and Project Management and Accounting.

Additional information, particularly with regards to service developments and other events associated with the service, is also provided via a monthly bulletin issued by email to all users. An archive of these bulletins is available at http://www.csar.cfs.ac.uk/monthly_bulletin

CSAR Focus is published twice a year and is also available to view on the Web - http://www.csar.cfs.ac.uk/general/newsletter.shtml. To change your mailing address please email your new details to the CSAR Helpdesk.

Getting Help

If you require help on any aspect of using the CSAR service you can contact the CSAR Helpdesk team who will deal with your query promptly and efficiently.

Telephone: 0161 275 5997 / 0161 275 6824

Email: csar-advice@cfs.ac.uk

The CSAR Helpdesk is open from 08:30 - 18:00 Monday to Friday, except on Public Holidays.

ISBN 1470-5893