CSAR

FOCUS

high performance computing at Manchester ~ taking research into the future

Edition 12
Summer - Autumn 2004

Newton Upgraded to Single 512 Processor System with 1 Terabyte of Memory

www.csar.cfs.ac.uk

MANCHESTER 1824
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

Deputy Director, Manchester Computing, University of Manchester

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

Dr Mike Pettipher

Head of Scientific Computing Services, Manchester Computing, University of Manchester

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

Carl Ward

Editor of CSAR Focus, Research Support Services, Manchester Computing, University of Manchester

Telephone: 0161 275 6822
Fax: 0161 275 6800
Email: carl.ward@manchester.ac.uk

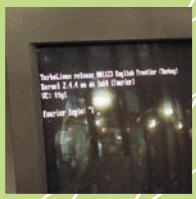


Disclaimer

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

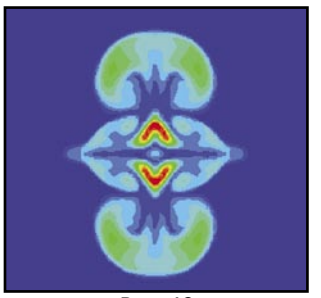

Page 12

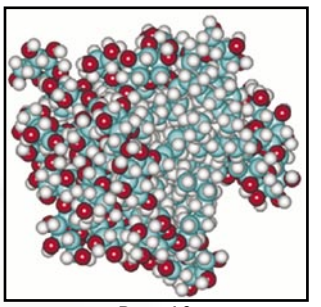

Page 10

Editorial

This edition of CSAR Focus is pleased to report on the expansion of Newton to form a 512 processor system with 1 Terabyte of memory. The article details the development of CSAR's SGI Altix facility highlighting the current configuration as well as the proud plans for the future. Elsewhere in this issue we detail more good news as Manchester wins the bid to host the Access Grid Support Centre.

Various members of the Research Support Services team have reported back from a number of events attended including VECPAR '04, ISC 2004, The SGI User Group and NEC's 10th Anniversary Research Forum, details can be found in Recent Events. A further contingent of the team is now getting ready to attend SC 2004 and a report will follow in the next issue.

A number of articles have been kindly supplied from CSAR Users including an article on Quantum Monte Carlo Techniques that benefitted from the large amounts of memory that the CSAR Service can provide. The article on computational fluids and combustion took advantage of the specialist support provided by CSAR staff particularly when porting and optimising code.

We also detail a new Class 3 project investigating molecular dynamics simulations of glycopids. Other features include a comparison of the performance of a number of HPC machines and commodity clusters and an article entitled "Three Dimensional Simulation of Rotor-Stator Cavity Flow" written by members of the former UMIST now united with the University of Manchester.

Other news to mention is the New Seminar series which will be taking place on Fridays 2-3pm (starting 26th November) on the Access Grid, any CSAR users that have access are welcome to attend. Please see http://www.mrccs.man.ac.uk for more details.

As a final note I hope you will enjoy reading this edition and I'd like to thank everybody who contributed features I would also welcome any future articles that relate to the CSAR Service so please feel free to contact me with any articles you may have.

Carl Ward
Editor, CSAR Focus

The Growth of Newton

Neil Stringfellow
Research Support Services, Manchester Computing, University of Manchester

The Altix Arrives

At the beginning of 2002, the demand for supercomputing resources from the UK academic community led to an agreement between the CfS board and EPSRC to extend the CSAR contract until June 2006. This extension was combined with an announcement of a second technology refresh, following on from the introduction of the SGI Origin 3800 machine named Green, and it was decided that the new technology to be introduced would be one of the first in SGI's new line of Altix supercomputers, based on Intel chips running the Linux operating system and utilising a high performance NUMAlink interconnect.

On October 1st 2003 a new SGI Altix facility was introduced into service to complement the SGI Origin machines and to act as a more powerful replacement for the Cray T3E 1200 which was due to be retired. The Altix was given the name Newton and consisted of 256 Intel Itanium 2 processors, each clocked at 1.3 GHz, together with 384 Gigabytes of memory. Within a few months Newton was being heavily used and was able to provide a concrete demonstration of the demand for high performance computing facilities in the academic community.

The high utilisation and continued demand for access to the facilities of both CSAR and HPCx, led to the decision being taken to further expand Newton to twice its initial compute capacity and to increase the available memory. This expansion was to take place in three stages in order to introduce extra resources without adversely affecting the continued operation of the service.

A Staged Expansion

In the first stage the original 256 processor Altix had it's memory increased to 512 Gigabytes, and an additional 128 processor Altix machine was added to the service, also with 512 Gigabytes of memory. While the two machines functioned as separate entities, they enabled a higher throughput of work by allowing common job submission to the batch system, with the scheduler able to select the most appropriate resource on which to execute the code.

The expanded 512 processor Newton

The second stage involved the addition of a further 128 processors to the newer machine, further increasing the computational capacity of the Altix service to 512 processors. These additional processors were 1.5 GHz Intel Itanium 2 chips with 6 Megabytes of tertiary cache compared to the original set of chips which were 1.3 GHz processors with 3 Megabytes of tertiary cache. This stage also saw the final configuration of the memory available for jobs being set at 2 Gigabytes per processor.

This is the current configuration of Newton which is already able to provide a high throughput of work for a variety of jobs utilising up to 228 processors in a single run, and with larger per-process memory requirements than have been available on any previous CSAR machine. The memory has increased with each technology refresh from a standard 256 Megabytes per processor on the Cray T3E, to 1 Gigabyte per processor on the

Origin machines and now to 2 Gigabytes per processor on the Altix. The fast processors and high memory are allowing scientific research to be attempted which was not possible on earlier CSAR machines, but the expansion of Newton is not yet complete.

Into the Future

The final stage in the expansion will be the rehousing of the two Altix machines in a new building and the combination of the separate entities which make up the newton cluster to form a single 512 processor system with 1 Terabyte of memory. This final configuration will allow jobs with high processor counts and large memory requirements to run across the entire system, further extending the range of scientific research which will be possible on the CSAR systems.

Further possible improvements to the throughput of research on Newton will not be confined to the advances in hardware, but should also be provided by potential faster program execution for software generated using the next release of the Intel compiler suite, and through the continuing optimisation effort of the CSAR application support team in fine-tuning specific codes to take advantage of the Itanium 2 processor and the NUMAlink interconnect.

The development of Newton will enable the CfS consortium to maintain the commitment of providing an excellent supercomputing facility up to the end of the current service contract, and provides a high quality resource for the UK academic community to continue to lead the way in carrying out world class scientific research.

Recent Events

NEC Europe Ltd. C&C Research Laboratories 10th Anniversary Research Forum: "Next Generation Computing Systems and Environments", Bonn, Germany, 15th-16th July 2004

Jon Gibson, Research Support Services, Manchester Computing, University of Manchester

This research forum was convened to celebrate the 10th anniversary of C&C Research Laboratories (CCRLE) of Bonn, Germany, part of NEC Europe Ltd, who undertake research and development work in the field of high performance and grid computing. Over the past decade the lab has developed algorithms for a full MPI-2 implementation and worked on HPF+ (an extended High Performance Fortran); it has developed support tools in areas such as parallel finite element analysis, numerical solvers and coupled ocean-atmosphere simulation and application software for medical use and computational fluid dynamics. It is now looking to the future with collaborations on a number of grid projects, such as GEMSS, Grid-Enabled Medical Simulation Services.

Talks were given by a range of invited speakers from around Europe, as well as by a number of representatives of NEC and CCRLE itself. One talk of particular relevance to many CSAR users was given by Mick Carter of the Met Office, entitled "Software Architecture for Climate and Weather Forecast Models". Here, he discussed the Unified Model, the current climate and weather forecasting code, and how despite having many good points, it is perhaps let down by a lack of software engineering. The Flexible Unified Modelling Environment (FLUME) was introduced as a project leading to a replacement, providing much needed flexibility and maintainability. There is also a collaborative European project called PRISM (Program for Integrated Earth System Modelling) aiming to develop a European climate modelling system. It was suggested that some co-ordination between the projects, perhaps leading to a single system, would be the best strategy.

Another highlight was the talk on cardiovascular simulation by Dr Hose of Sheffield University, where a number of medical exhibits, including an artificial heart, were passed around the audience. Not too long after this slightly gruesome spectacle was the intriguingly-billed "German-Japanese dinner", an opportunity to sample the finest national cuisine that these two countries have to offer – a treat indeed! I'm hoping that the next ten years are just as fruitful for CCRLE so that I may experience this culinary extravaganza once again.

SGI User Group, Orlando, Florida

Andrew Jones, Research Support Services, Manchester Computing, University of Manchester

The first annual conference of the SGI Users Group (SGIUG) was held in Orlando, Florida in May this year. A range of tutorials, talks and demos were given across the areas of HPC, grid, visualization and storage. The conference was felt to be a great success by attendees. Several talks were given by staff from Manchester covering applications, benchmarking, grid and more, and also one from CSC highlighting early operational experiences of Newton. Manchester and CSC are also key members of the board of the SGIUG, and this conference program was primarily planned from Manchester. The 2005 conference will be held in Munich, Germany in June. More details are available at http://www.sgiug.org.

Robert Haines working on the booth at ISC 2004

ISC 2004, Heidelberg, Germany

Pen Richardson, Research Support Services, Manchester Computing, University of Manchester

Mike Pettipher, Robert Haines and Penny Richardson hosted a booth at the 19th International Supercomputing Conference in June 2004. This was the third year of attendance for the University of Manchester, with the main focus being to illustrate work done by CSAR and it's users, as well as visualization and e-Science projects. The theme of this year's conference was "Applications, Architectures and Trends" with Mike Pettipher attending the tutorial on "Performance in HPC: Evaluation, Modelling, Benchmarking and Prediction". Stephen Pickles presented a plenary talk about the TeraGyroid experiments that had taken place at SC2003 and received the ISC award for the best paper on "Integrated Data and Information Management".

VECPAR '04, Valencia, Spain

Kevin Roy, Research Support Services, Manchester Computing, University of Manchester

The 6th International VECPAR meeting was hosted by the Universidad Politécnica de Valencia in Valencia in Spain. This biennial multidisciplinary conference was the first to be hosted outside of its native Portugal. Attendees came from all over the world with a significant presence from Brazil, China, France, Spain and USA.

The papers given ranged from tradititional areas such as parallel and distributed computing, numerical methods and large scale simulations to newer topics to HPC such as biosciences, scientific visualization and cluster and grid computing.

This year the talks were dominated by stuidies of numerical problems, in particular there was a significant presence in finite element codes, and the increasingly popular cluster and grid papers.

There were six very interesting invited speakers from a healthy range of disciplines. Tetsuya Sato form the Earth Simulator Center in Japan and Fabrizio Gagliardi from the EGEE European Grid Infrastructure Project, both talked on how their orgainisations hoped to challenge people's perceptions of what problems should be and could be solved with ever increasing powerful systems.

Mike Shelley and Michael Heath both from the USA talked about the research that their groups were involved with, that of flexible bodies in fluid flows and coupling of multicomponent systems respectively. The other two invited speakers were from the Patrick Valdiurez and Vincent Breton, both working in France, who talked about two projects who seek to use the grid to improve computational use and reliability within the medical community.

Overall this was a very interesting conference giving a great opportunity to learn more about the research that is happening in many different areas across the globe.

http://www.vecpar.fe.up.pt

New Avenues for Quantum Monte Carlo Techniques

D.Alfè and M. J. Gillan†, *Department of Earth Sciences, †Department of Physics and Astronomy, University College London*
N. D. Drummond, M. D. Towler and R. J. Needs, Cavendish Laboratory, University of Cambridge

The Holy Grail of condensed matter physics is the ability to solve the 78 year old Schrödinger equation for "real-life" systems. The great difficulty is that the motions of the electrons in molecules and crystals are correlated because of the strong repulsion between their negative charges. The solution of the many-electron Schrödinger equation (the wavefunction) is therefore a complicated function of 3*N* variables, where *N* is the number of electrons in the system. A boost to the field was given by Hohenberg and Kohn (HK) 40 years ago [1], with the introduction of density functional theory (DFT). They showed that it was possible to reformulate quantum mechanics in such a way that the important physical quantity is the electron density, rather than the electronic wavefunction, reducing the complexity of the problem from 3*N* to 3. Of course, free lunches are seldom available, and the price to pay for the great simplification of HK was the introduction of a new quantity, called exchange-correlation (XC) energy, which is an (as yet) unknown functional of the electron density. To make the theory work, HK suggested a simple form for the XC functional, known as the local-density approximation (LDA). The LDA is exact in a system with a homogeneous electron density, and is only an approximation in (real-life) inhomogeneous systems. Forty years later the LDA is still widely used, and has been the main factor for the great success of DFT, however, there are a number of cases where more accuracy is needed than the LDA can provide. Physicists have struggled for decades to find better approximations for the XC energy, and come up with a number of improvements to the LDA, but there are still several 'difficult' cases for which no available approximation for the XC is really satisfactory.

A completely different approach to the problem is offered by so called quantum Monte Carlo (QMC) techniques. Although their demands on computer power are much greater than those of widely used techniques such as DFT, their accuracy is also much greater for most systems. QMC traditionally embraces two different techniques: variational Monte Carlo (VMC) and diffusion Monte Carlo (DMC). For a detailed descriptions of VMC and DMC see for example Ref. [2], here we only outline rather briefly the main ideas for the two techniques.

The VMC method gives an upper bound on the exact ground-state energy E_0, but the results depend totally on the quality of the chosen trial wavefunction. Given a normalized trial wavefunction $\Psi_T(\boldsymbol{R})$, where $\boldsymbol{R} = (\boldsymbol{r}_1, \boldsymbol{r}_2 ..., \boldsymbol{r}_N)$ is a $3N$-dimensional vector representing the positions of N electrons, and denoting by $\hat{H}$ the many-electron Hamiltonian, the variational energy $E_v \equiv \langle \Psi_T | \hat{H} | \Psi_T \rangle \geq E_0$ is estimated by sampling the value of the local energy $E_L(\boldsymbol{R}) \equiv \Psi_T^{-1}(\boldsymbol{R}) \hat{H} \Psi_T(\boldsymbol{R})$ with configurations $\boldsymbol{R}$, distributed according to the probability density $\Psi_T(\boldsymbol{R})^2$. In our calculations the trial wavefunction $\Psi_T(\boldsymbol{r}_1, ... \boldsymbol{r}_N)$ consists of a Slater determinant D of single-electron orbitals $\psi_n(\boldsymbol{r}_i)$ multiplied by a parameterized Jastrow correlation factor $J(\boldsymbol{r}_1, ... \boldsymbol{r}_N)$ (in the pseudopotential-based QMC of interest here, the $\psi_n(\boldsymbol{r}_i)$ are commonly taken from a plane-wave pseudopotential DFT calculation), and J is "optimized" by varying it's parameters so as to reduce the variance of the "local energy" $\Psi_T^{-1}(\hat{H} \Psi_T)$, where $\hat{H}$ is the many-electron Hamiltonian.

A much more accurate method is DMC, a scheme which in principle yields the exact ground state energy. The main idea of DMC is to solve the Schrödinger equation in imaginary time. This is done by recognizing that this is equivalent to a diffusion equation which includes a potential term, and therefore describes the stochastic motion of Brownian particles (walkers), with the effect of the potential being equivalent to a branching term which has the effect of increasing the number of walkers in regions of low potential and decrease it in regions of high potential. Since there is a one-to-one correspondence between the density of walkers and the value of the wavefunction in any region of space, one can compute the value of the wavefunction by following the motion of these walkers in configuration space. To make the calculations practicable it is necessary to "guide" the walkers using a "trial" wavefunction. The

trial wavefunction should be as close as possible to the real many-body wavefunction for optimal efficiency. Moreover, because of the fermionic nature of the electrons, the many-body wavefunction necessarily has regions of space where it is positive and regions where it is negative. This change of sign is a problem for the diffusion of the walkers. The way this problem is usually addressed is by fixing the surface which separates regions of space where the wavefunction changes sign (nodal surface). The presence of the trial wavefunction allows an easy implementation of this so called "fixed-node" constraint.

With QMC now being applied to large complex systems containing hundreds of atoms, a major issue is the scaling of the required computer effort with system size. The cost of traditional quantum mechanics techniques increases at least with the square of the number of atoms N in the system, and for most of them (including standard QMC) increases as N^3. However, feeding out of ideas about the locality of quantum coherence [3], it has been shown that in fact linear scaling should be achievable, and $O(N)$ techniques have been developed for tight binding (TB) [4], DFT [5, 6] and Hartree-Fock [7]. Recently, a procedure has been suggested [8] for achieving at least partial linear scaling for QMC, based on the idea of "maximally localized Wannier functions" [9]. This is done by performing a unitary transformation which preserves the orthogonality of the single particle orbitals that make up the determinant D, and by using a localized basis set for representing the single particle orbitals. Since the single particle orbitals are mixed by a unitary transformation, the determinant D is unchanged. Linear scaling is achieved then by truncating the orbitals beyond a certain localization region ω. Of course, in order to have accurate calculations, ω must be big enough so that the localization weights of the single particle orbitals, defined as $P = \int_{\omega} d\,\boldsymbol{r} \mid \phi(\boldsymbol{r}) \mid^2$, is as close to 1 as possible.

We have implemented and tested an alternative method which has two important advantages over the method of Ref. [8]. Firstly, we have developed a technique for representing single-particle orbitals in terms of localized B-splines [10]. These localized functions have previously been used in the context of DFT [11], and they are closely related to plane waves. Our implementation [10] therefore provides a natural interface between plane-wave-DFT codes, from which one can obtain good trial wavefunctions, and QMC codes. Secondly, we have shown that the localization of the single-particle orbitals can be greatly enhanced by removing the orthogonality constraint (i.e. by using a non-unitary linear transformation on the set of orbitals) [12].
We have tested our method on the prototypical oxide material MgO in the rock-salt structure at ambient pressure. Because of its large band gap (experimental Eg = 7.7 eV), this kind of material should be particularly suited to O(N) methods. The QMC calculations 2 were performed using the appropriately modified CASINO code [13] and a supercell of 64 atoms.

To compare the efficiency of our method with the earlier technique based on Wannier functions [8], we plot Q = 1 - P against the side-length of the cubic localization region for the two methods in Fig. 1. In Fig. 2 the VMC energies obtained using the two methods are plotted against the side-length of the localization region. The energy obtained with untruncated orbitals (the 'standard' QMC result) is also shown. We see that nonorthogonal orbitals are much more localized than Wannier functions, and that the VMC energy evaluated with these orbitals converges to the correct value much more rapidly when the localization region is expanded. Further details can be found in our recent publication, Ref. [12].

In conclusion, we have proposed and tested a new technique for achieving linear scaling in one of the most demanding parts of QMC calculations. This technique is simple and robust, and it appears also to be more efficient than an earlier proposed technique. This new technique already makes it possible to treat large systems that would be out of reach of conventional QMC methods. Research areas where we intend to apply the technique in the near future include defects and surfaces of oxide materials.

These calculations have been performed on both the CSAR machines and the HPCx national services. Despite the larger computational power of HPCx, we note that the CSAR machines were essential for carrying out part of this work, thanks to the much larger memory available. This work has been supported by the Mineral Physics Consortium, the U.K.C.P. consortium and the Materials Chemistry Consortium. DA greatly acknowledges support from the Royal Society, and from the Leverhulme Trust.

References

[1] P. Hohenberg and W. Kohn, Phys. Rev. 136, B864 (1964).
[2] W. M. C. Foulkes, L. Mitas, R. J. Needs and G. Rajagopal, Rev. Mod. Phys. 73, 33
(2001).

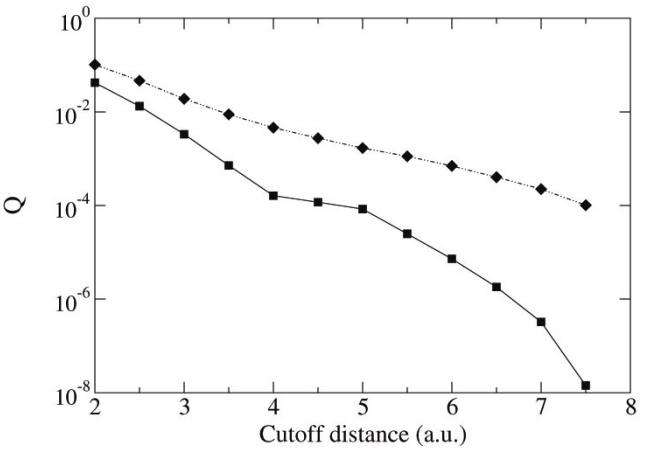

Figure 1: Dependence of localization weight on cut-off distance for bulk MgO. The quantity plotted is $Q \equiv 1 - P$, where P is localization weight defined in the text. Squares: present non-orthogonal orbitals. Diamonds: maximally localized Wannier orbitals.

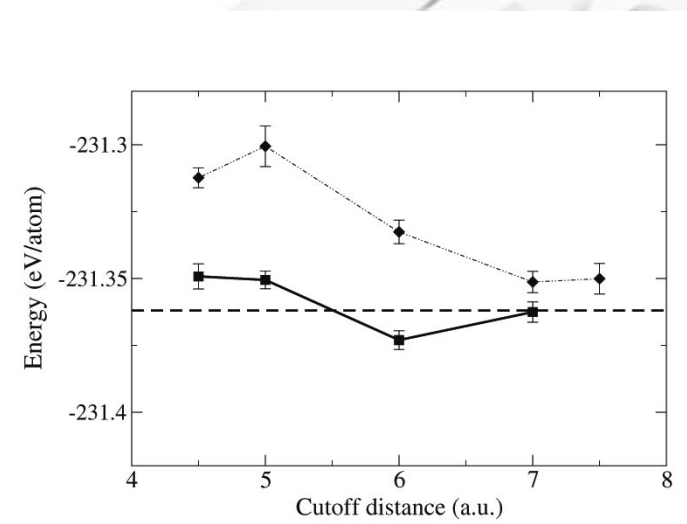

Figure 2: Convergence of linear-scaling VMC total energy per atom to the value obtained with extended orbitals for bulk MgO. Squares: present non-orthogonal orbitals. Diamonds: maximally localized Wannier functions. Horizontal dashed line shows total energy/atom obtained with extended orbitals.

[3] W. Kohn, Phys. Rev. Lett. 76, 3168 (1996).
[4] D. R. Bowler, M. Aoki, C. M. Goringe, A. P. Horsfield and D. G. Pettifor, Modell.
Simul. Mater. Sci. Eng. 5, 199 (1997)
[5] D. R. Bowler, T. Miyazaki and M. J. Gillan, J. Phys.: Condens. Matter 14, 2781
(2002).
[6] J. M. Soler, E. Artacho, J. D. Gale, A. Garcia, J. Junquera, P. Ordejon and D. Sanchez-Portal, J. Phys.: Condens. Matter 14, 2745 (2002).
[7] M. Challacombe, J. Chem. Phys. 110, 2332 (1999).
[8] A. J. Williamson, R. Q. Hood and J. C. Grossman, Phys. Rev. Lett. 87, 246406 (2001).
[9] N. Marzari and D. Vanderbilt, Phys. Rev. B 56, 12847 (1997).
[10] D. Alfè and M. J. Gillan, Phys. Rev. B, Rapid Comm., in press.
[11] E. Hernández, M. J. Gillan and C. M. Goringe, Phys. Rev. B 55 13485, (1997).
[12] D. Alfè and M. J. Gillan, J. Phys.: Condens. Matter 16, L305 (2004).
[13] R. J. Needs, M. D. Towler, N. D. Drummond and P. R. C. Kent, 'CASINO Version 1.7 User Manual', University of Cambridge, Cambridge (2004).

Molecular dynamics simulations of glycolipids using high performance computing

Richard A. Bryce (School of Pharmacy and Pharmaceutical Sciences, University of Manchester), Teoh Chong and Rauzah Hashim (Department of Chemistry, University of Malaya)

Carbohydrates are the most naturally abundant biomolecule. Polysaccharides are involved in a range of functions, from energy storage to structural integrity. In addition to constituting a key component of DNA and RNA, carbohydrates form conjugates with proteins and lipids. With a hydrophilic carbohydrate headgroup and a hydrophobic lipid tail, glycolipids are amphilic molecules, able to dissolve hydrophobic moieties and self-assemble in different ways to form bilayers, vesicles and micelles. These assemblies have nanotechnological potential, with liquid crystal behaviour, in addition to uses as non-ionic detergents, drug delivery agents and as a model to probe aspects of the glycocalyx. The relationship between structural and physical properties of glycolipid assemblies is not fully understood. Coarse-grained models have been employed to simulate these systems, but have struggled to describe the subtlety of their physicochemical behaviour.

Therefore, the research groups of Richard Bryce at the University of Manchester and Rauzah Hashim at the University of Malaya, Malaysia, are employing detailed molecular dynamics (MD) simulations in an effort to understand the structure-property relationships of glycolipids in lamellar phases. Carbohydrates themselves exhibit considerable conformational complexity: even the simple monosaccharide glucose has 729 possible conformers, assuming three rotamers per single rotatable bond. Thus, to equilibrate the multiple flexible carbohydrate headgroups requires long time scale simulations. The computational expense is also increased by the need for a realistic description of (1) the long-tail electrostatics interactions engendered by the polar headgroups and (2) interaction with aqueous solvent. For this, the micelle is simulated in a bath of explicit solvent molecules, using truncated octahedral periodic boundary conditions and the particle-mesh Ewald method (Figure 1).

High performance computing is well-suited to the molecular dynamics technique, where the large number of pairwise nonbonded interaction energies and forces can be distributed over many processors. We use a message-passing version of AMBER where bonded and nonbonded energies and forces, pairlist creation and integration steps are parallelized. Thus, simulation times can be reduced from weeks to days, permitting improved statistics using different initial conditions. Results from the simulations are providing insights into the molecular-level behaviour of glycolipid micelles. For example, an MD snapshot of a micelle of 1-*O*-*n*-octyl β-D-glucopyranoside is shown in Figure 2. The snapshot is taken from a 5 ns simulation at a temperature of 300 K and a high solute concentration of

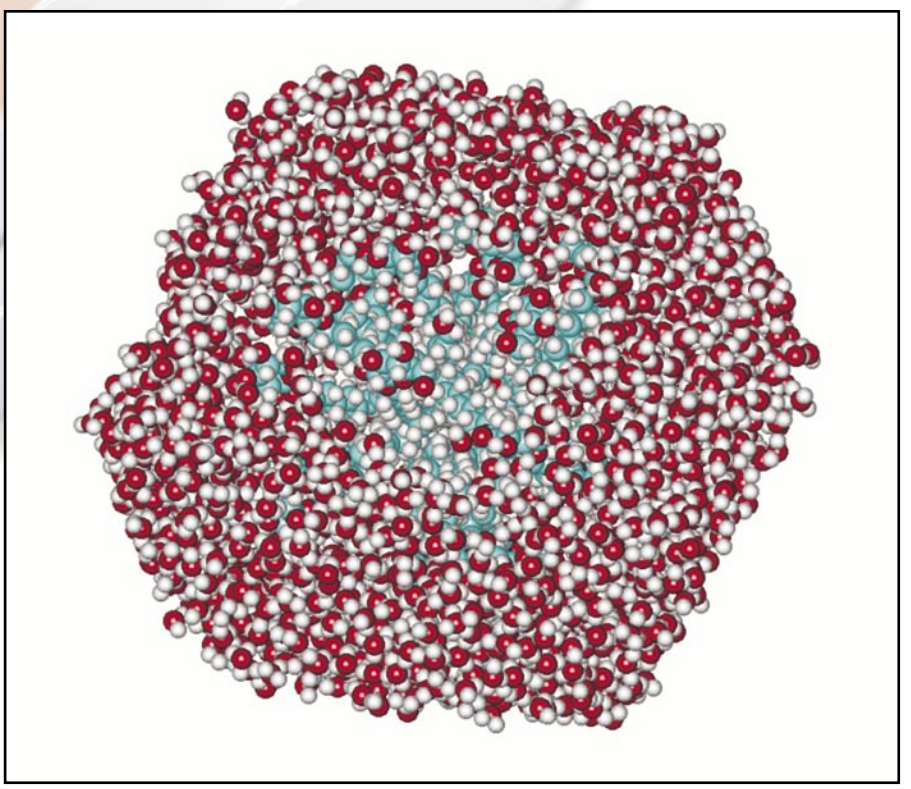

Figure 1 Micelle in a truncated octahedral periodic box of water

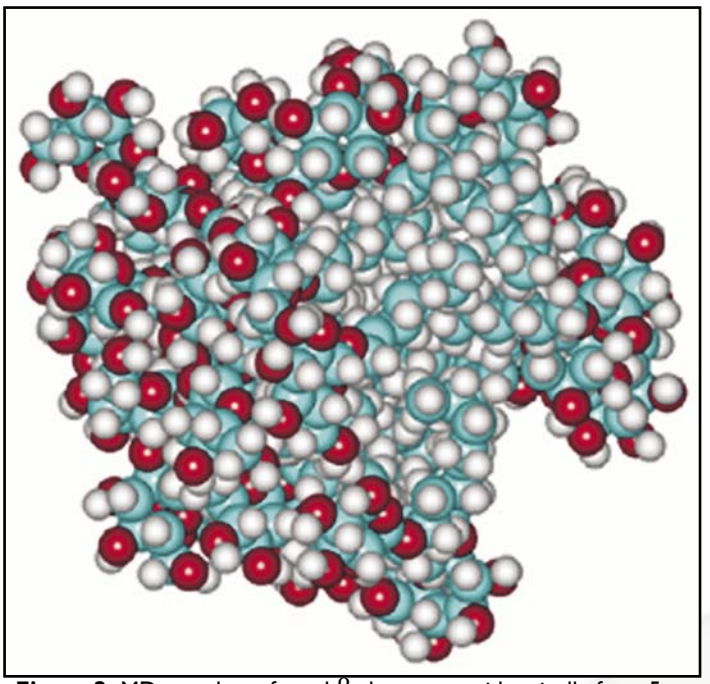

Figure 2 MD snapshot of octyl β-glucopyranoside micelle from 5 ns simulation. Oxygen (red), carbon (cyan) and hydrogen (white) atoms are shown.

0.6 M (well above the critical micelle concentration), performed using the program AMBER 7 on the SGI Altix Newton. The density profile (Figure 3) indicates the broad distribution of the headgroups over 7 to 23 Å from the centre of mass of the micelle. Also apparent is the penetration of the aqueous solvent within the micelle into the headgroup region. This arises in part from the exposure of hydrophobic surface area (Figure 2) to surrounding water. The average headgroup radius is ~14 Å (Figure 3), which is in reasonable agreement with a hydrodynamic radius of 15±1 Å (an upper limit for actual micelle radius), determined from chromatographic experiments for an aggregation number of 27 [1]. Thus, molecular dynamics simulations on supercomputers are enabling access to detailed studies of glycolipid structure and dynamics, deepening our understanding of these therapeutically and commercially important materials.

Acknowledgements

This work is supported in part by grant 09-02-03-9010 SR0004104 (MOSTE).

Contact details

Richard A. Bryce
School of Pharmacy and Pharmaceutical Sciences,
University of Manchester, Manchester, M13 9PL, U.K.
email: R.A.Bryce@man.ac.uk
www.pharmacy.man.ac.uk/rab
www.glycoprojects.kimia.um.edu.my

References

[1] B. Lorber, J. B. Bishop and L. J. DeLucas. *Biochim. Biophys. Acta*, **1990**, *1023*, 254-265.

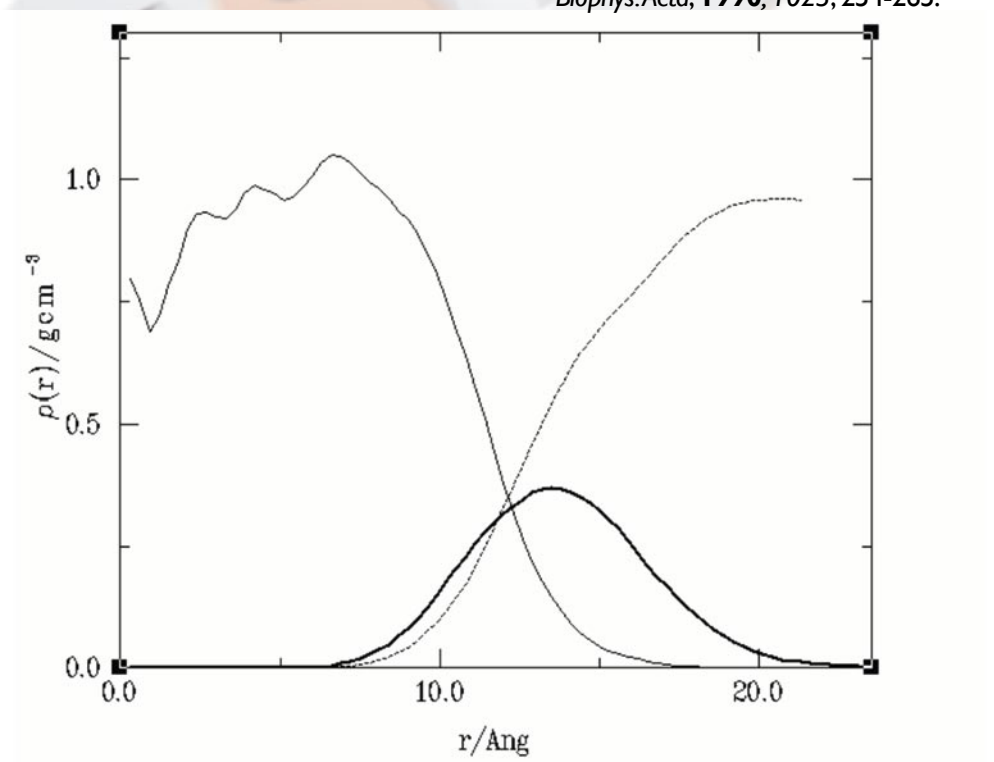

Figure 3 Local density profile of micelle/solvent system: glucose headgroup (bold solid line), lipid tail (solid line) and water (dashed line).

Seeing the invisible through direct numerical simulation

K. H. Luo
School of Engineering Sciences, University of Southampton, Highfield, Southampton; formerly Department of Engineering, Queen Mary, University of London

Introduction

In the scientific world, seeing is not necessarily believing. However, the ability to reproduce and visualise is vital to achieving full understanding of complex, unsteady and non-linear phenomena. Turbulence and other physical processes occurring in a turbulent medium, for example, are largely invisible to us and remain unsolved mysteries in engineering and physical sciences. Various experimental techniques such as high-speed cameras and laser-induced fluorescence are able to visualize turbulent phenomena to some extent, but they suffer from certain defects: (a) The spatial and temporal resolution is usually not sufficient to resolve small-scale phenomena; (b) Only basic quantities such as velocity and temperature are measurable while derived or high-order quantities (e.g. dissipation rate) are almost impossible to obtain; (c) Correlations between two or more variables are difficult to measure; and finally (d) seeding the flow field alters the phenomenon being investigated.

Computer visualization is a natural tool to complement experimental measurements, which can overcome the latter's inherent disadvantages (a)-(d), provided sufficient computer resources are made available. In the early days when computers were number-crunching below megaflops, only the Reynolds Averaged Navier-Stokes (RANS) calculations were possible. The results are not true visualization of turbulence, because the non-linear processes involved are not reproduced but represented by pre-determined turbulence models. Direct computer visualization of turbulent phenomena became a reality with the advent of gigaflops and now teraflops computers, which allowed Direct Numerical Simulation (DNS). As DNS solves the fundamental governing equations without the interference of empirical models, the simulation process is a reproduction of the physical "reality", and the non-linear turbulent processes are "seen" throughout the simulation. Analysis of any complex and unsteady phenomenon is greatly enhanced through access to the DNS data which contain full-field, time-dependent and accurate information of almost all quantities.

Current Research

Our work on computational fluids and combustion has benefited greatly from access to world class facilities provided by CSAR and HPCx. Generous support has been provided by CSAR staff, in particular, Dr. Neil Stringfellow who helped to port and optimize our codes on Green and Newton. Our research is conducted within two national consortia: the "UK Turbulence Consortium" (EPSRC Grant No. GR/R64964 with Prof. Neil Sandham as P.I.) and the "Consortium on Computational Combustion for Engineering Applications" (EPSRC Grant No. GR/R66197 with myself as P.I.). Each of these consortia consists of active researchers from up to eight universities in the UK, which provides an effective forum for dissemination of results to UK and international audience through regular meetings and international workshops. Current activities in our group fall into three main areas: (1) computational aerodynamics and aeroacoustics; (2) turbulent diffusion flames and partially premixed flames; and (3) mixing and dispersion of solid particles and liquid droplets in turbulent flow. Sample results from these numerical investigations are shown in the accompanying figures.

Direct computation of the invisible sound and its source generated by complex shear flows has always been desired, which presents a great challenge for numerical methods and requires massive computing power. Figure 1 shows results of a DNS of a compressible flow in and around

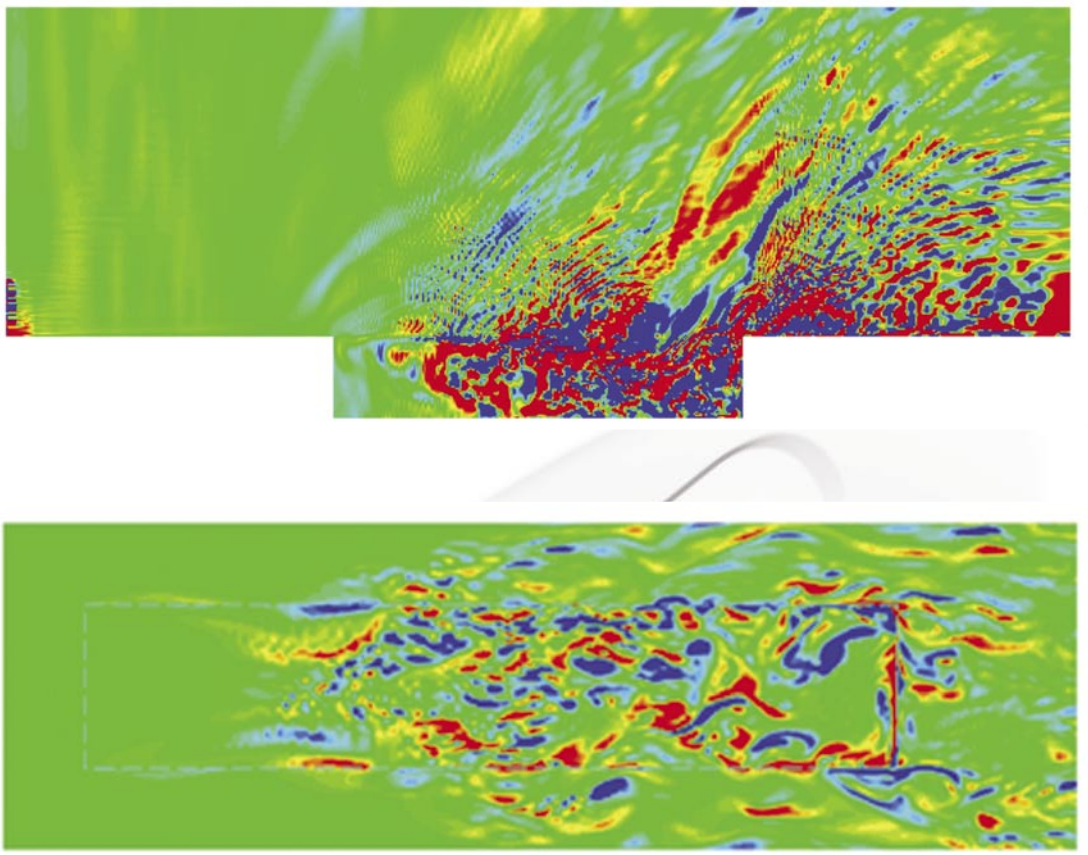

Figure 1: DNS of a compressible open cavity flow and its acoustic near-field at a Mach number 0.85 and a Reynolds number 5000, using 21 million grid points. The top plot shows a side view of the dilatation field and the lower plot a top view of the vorticity field (Lai and Luo, 2004).

an open cavity using a Reynolds number of 5000. The same configuration was used in a series of experimental measurements at DSTL. The geometry was fully three-dimensional, with side walls of the cavity included, which was very rare in numerical simulations. The side view of the dilatation field shows identifiable wave patterns in the near-field. The top view of the vorticity field demonstrates clear asymmetry of the flow in the spanwise direction, indicating strong three-dimensional effects, which were absent from a simulation at a lower Reynolds number of 1000. A Large Eddy Simulation (LES) under identical conditions was also performed. The predicted sound source of the cavity flow, as defined by Lilley, is shown in Figure 2 (overleaf). A comparison of DNS and LES predicted sound pressure level (SPL) with the DSTL measurements is also shown.

A series of direct numerical simulations have been performed to study the link among mixing, entrainment, vortex dynamics, combustion and the source configurations of transitional and turbulent buoyant diffusion flames. The full Navier-Stokes equations and chemical reactions governed by the Arrhenius law were solved by high-order numerical schemes. Figure 3 (overleaf) shows a diffusion flame from a rectangular source undergoing axis switching. The large number of grid points used, 256 x 256 x 768, is due to the fact that combustion

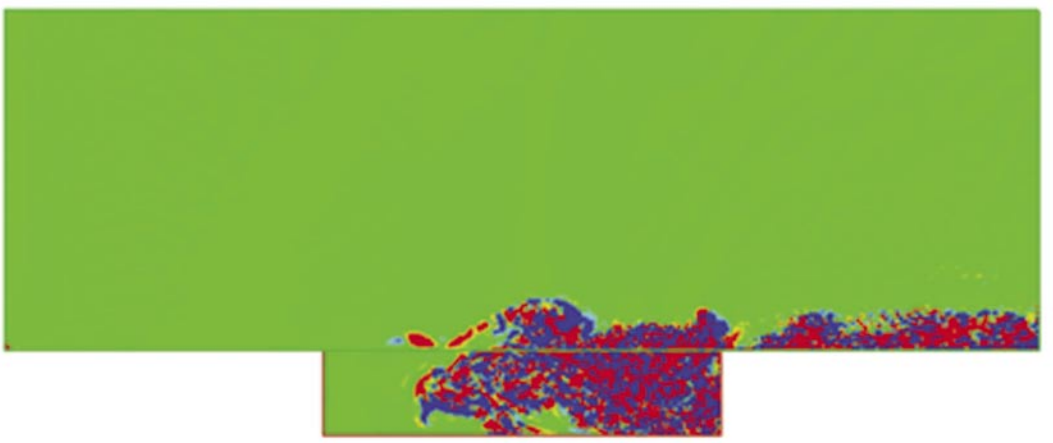

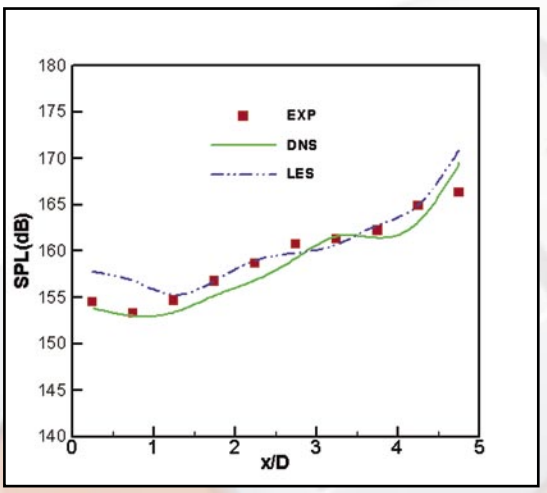

Figure 2: The acoustic source field obtained by LES of a compressible open cavity flow at a Mach number 0.85 and a Reynolds number 5000, using 5 million grid points and a comparison of the sound pressure level (SPL) from numerical predictions with the DSTL experimental data (Lai and Luo, 2004).

takes place at much smaller scales than the smallest scales of flow, the Kolmogorov scales.

Partially premixed flames are known to exist but direct proof by experiments or numerical simulations is not easy, due to the extremely small scales involved and the transient nature of the phenomena. Figure 4 shows a DNS of a lifted methane-air flame. Using the concept of cross-scalar dissipation rate, $\chi_{F,O} = D \nabla Y_F \bullet \nabla Y_O$, which is essentially the product of the gradient of the fuel concentration and the gradient of the oxidizer concentration, the lifted flame is seen to have partial premixing, if the reaction rate is conditioned on $\chi_{F,O}$. By following the evolution of the reaction zone numerically, the dynamics of partially premixed flames can be explained unambiguously.

Summary and Outlook

DNS has enabled unaltered views of fine details of turbulence and turbulence-related phenomena such as turbulent combustion, which are otherwise "invisible" to us. It is the only means by which complex unsteady phenomena in Thermofluids can be recreated with no or minimum empirical interference, providing a valuable tool for obtaining fundamental insight and for model validation. However, the computational cost is prohibitively high for most practical calculations. And despite the availability of today's teraflops hypercomputers,

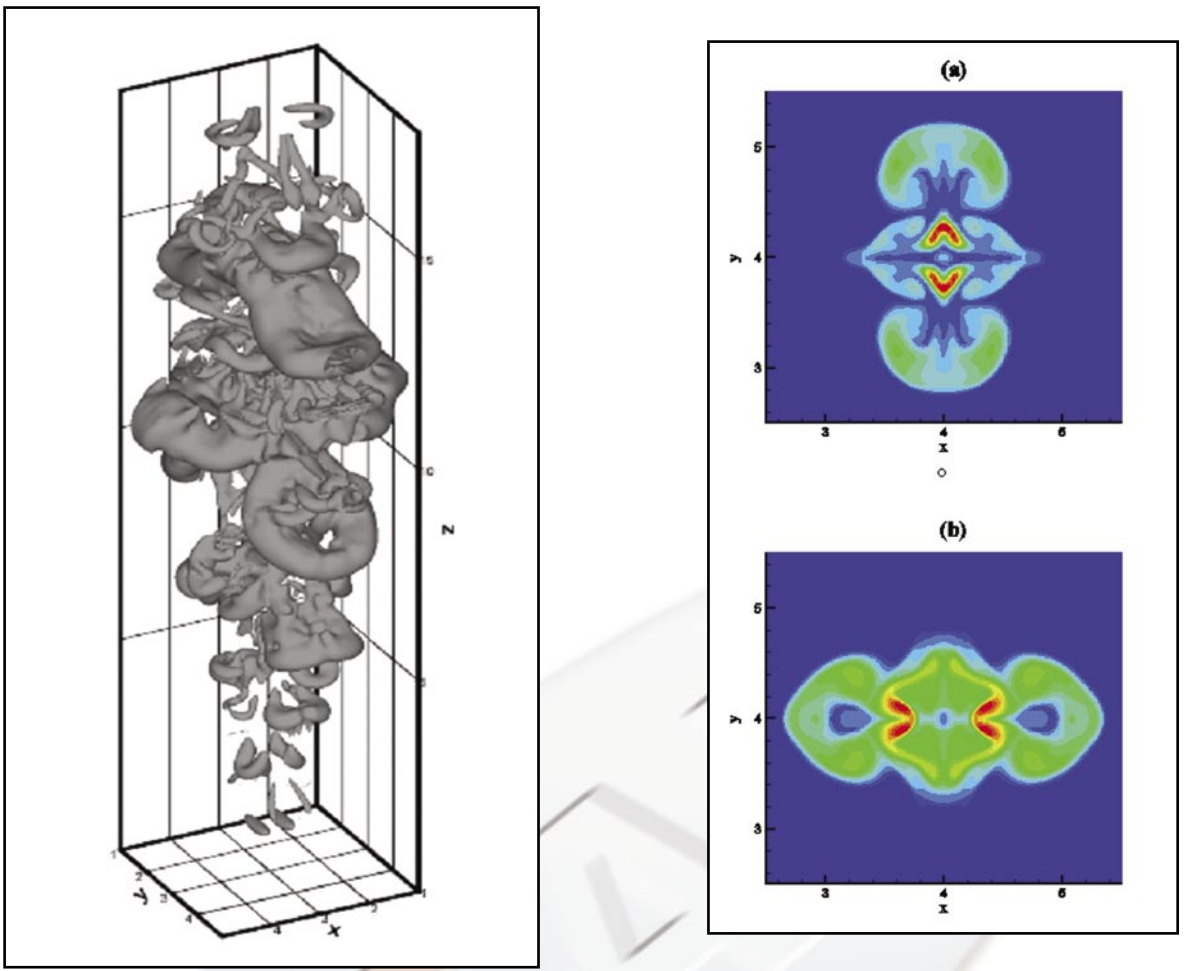

Figure 3: DNS of a buoyant diffusion flame from a rectangular source showing axis switching. Left --- pressure iso-surfaces showing the vortex structures; Right --- combustion product distributions in two vertical planes before and after the axis switching (Luo, 2004).

(a) (b)

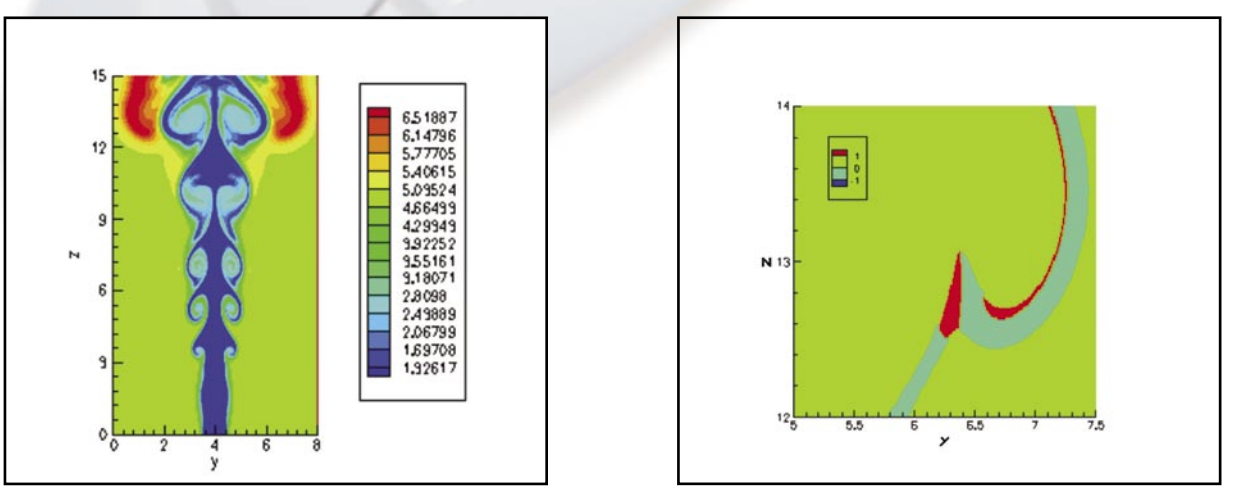

Figure 4: DNS of a lifted flame at a Reynolds number 9500. (a) A representative temperature field showing the reaction zone detached from the source and (b) a close-up view of the conditional reaction rate field showing partially premixed flame (Luo, 2003).

DNS is limited to a small range of physical parameters. The desire of extending such limits is curbed by the fact that the pace of demand from the physical point of view is always exceeding even the very rapid advancement in computer technology. For DNS of turbulence only, the computational cost (memory x CPU requirements) scales with the turbulent Reynolds number to the power of 3. In other words, an order of magnitude increase in the Reynolds number would increase the computational cost by a factor of 1000, which in turn would require the computer industry roughly 10 years to achieve on a constant-cost basis. Nevertheless, the current computing techniques coupled with development of higher-order numerical schemes have enabled DNS to pass the threshold of being able to simulate physical conditions comparable to the lower range of conditions in experiments. Moreover, it is the sole method that can compute 2nd and higher order quantities accurately. In comparison, the RANS approach can only calculate the mean quantities (zero order) and the LES may compute the first derivatives (1st order), with acceptable accuracy. In a word, the desire of the scientific and engineering communities to seek truth and to break the limits will continue to drive the development of DNS and related modelling approaches such as LES.

Acknowledgements

Support from the EPSRC Grant No. GR/R66197 and the EPSRC/DSTL/MOD Grant No. GR/R85303/01 is gratefully acknowledged. Thanks are due to Prof. N. D. Sandham who provided much of the computing resources through the "UK Turbulence Consortium". Work of past and current colleagues, Dr's. E. Avital, Y. M. Chung, Z. W. Hu, X. Jiang , X. Zhou and Prof. J. J. R. Williams as well as expert support from Dr. N. Stringfellow at CSAR is greatly appreciated.

References

[1] X. Zhou, K. H. Luo and J. J. R. Williams (2002). "Vortex dynamics in spatio-temporal development of reacting plumes," Combustion and Flame, 129: 11-29.
[2] Y. M. Chung and K. H. Luo (2002). "Unsteady heat transfer analysis of an impinging jets," Transactions of ASME: J. Heat Transfer, 124 (6): 1039-1048.
[3] Y. M. Chung, K. H. Luo and N. D. Sandham (2002). "Direct numerical simulation of an impinging jet," Intl. J. Heat and Fluid Flow, 23 (5): 592-600.
[4] Z. W. Hu, X. Y. Luo and K. H. Luo (2002). "Numerical simulation of particle dispersion in a spatially-developing mixing layer," Theoret. Comput. Fluid Dynamics, 15(6): 403-420.
[5] X. Jiang and K. H. Luo (2003). "Dynamics and structure of transitional buoyant jet diffusion flames with side-wall effects," Combustion and Flame, 133 (1-2): 29-45.
[6] K. H. Luo (2003). "New opportunities and challenges in fire dynamics modelling," Proc. 4th Intl. Seminar on Fire and Explosion Hazards, 15 pages, Londonderry, N. Ireland, 8-12 September (Invited keynote speech).
[7] K. H. Luo (2003). "A study of partial premixing in flames by direct numerical simulation," European Combustion Meeting 2003, Orleans, France, October 25-28.
[8] X. Jiang, E. Avital and K. H. Luo (2004). "Direct computation and aeroacoustic modelling of a subsonic axisymmetric jet," J. Sound and Vibration , 270 (3): 525-538.
[9] X. Jiang, E. Avital and K. H. Luo (2004). "Sound generation by vortex pairing in subsonic axisymmetric jets," AIAA J., 42 (2): 241-248.
[10] K. H. Luo (2004). "Enhancement of mixing and combustion in buoyant non-circular jets," in ERCOFTAC Series: Direct and Large-Eddy Simulation V, 9, 375-384, R. Friedrich, B. J. Geurts and O. Metais (Eds.), Kluwer Academic Publishers, Dordrecht. ISBN 1-4020-2032-5.
[11] K. H. Luo (2004). "Instabilities, entrainment and mixing in buoyant reacting jets," European J. Mechanics B: Fluids , 23 (3): 443-460.
[12] K. H. Luo (2004). "Axis switching in buoyant diffusion flames," Proc. Combust. Inst., 30: 1-7 (in press).
[13] H. Lai and K. H. Luo (2004). "Direct numerical simulation of three-dimensional compressible flows in an open cavity," in Recent Advances in Fluid Mechanics, F. Zhuang and J. Li (eds.), Tsinghua University Press & Springer, Beijing, pp. 790-793.
[14] K. H. Luo (2004). "Fluid dynamics effects on lifted flames," in Recent Advances in Fluid Mechanics, edited by F. Zhuang and J. Li (eds.), Tsinghua University Press & Springer, Beijing, pp. 381-385.
[15] H. Lai and K. H. Luo (2004). "Large eddy simulation of a compressible flow over a fully three-dimensional open cavity," Flow, Turbulence and Combustion (to be submitted).

Application Performance of Modern Number Crunchers

Gerhard Wellein, Thomas Zeiser, Georg Hager, Regionales Rechenzentrum Erlangen (RRZE), Martensstraße 1, 91058 Erlangen
Peter Lammers, Hochleistungsrechenzentrum Stuttgart, Allmandring 30, 70550 Stuttgart

Is There Still a Need for Classical HPC Systems or Can We Go Commodity Off-the-Shelf?

In preparation of upcoming procurements in Germany a comprehensive performance evaluation for a selection of applications and kernels has been performed. The main focus of our study was on codes from computational fluid dynamics, which are known to be memory-intensive. In the light of our results we want to comment on the discussion about the need for tailored HPC systems such as SGI Altix or vector computers.

Introduction

Looking at the TOP500 list [1] of recent years, more and more classical HPC systems have been replaced by commodity off-the-shelf (COTS) clusters which do not mainly focus on HPC requirements, but dominate the HPC market due to their (often) moderate price-performance ratio. However, it has also been acknowledged recently that the gap between *sustained* and *peak* performance for scientific applications on COTS platforms is growing continuously [2]. Tailored HPC systems, such as SGI Altix or vector computers, have been designed to meet the numerical requirements of scientific, memory-intensive applications. High *sustained* single processor floating point performance, high memory bandwidth, a balanced interconnection network and a mature software environment (compilers and libraries) are the characteristics for the latter systems. In this report we comment on performance and scalability of tailored HPC systems versus COTS clusters using two applications from computational fluid dynamics (CFD).

Benchmark systems

We have chosen a GBit/Xeon cluster, a Myrinet/Opteron cluster, a SGI Altix 3700 system and a NEC SX6+ based system. Details of the configurations can be found in Table 1.

An important difference between the Intel x86 and AMD x86_64 design is the memory subsystem. While Intel still promotes (also for its new EMT64 architecture) bus based architectures where two or four processors share one path to main memory, AMD uses a separate memory interface for each CPU providing scalable bandwidth within a shared memory node. For this reason, the AMD design is favourable for memory bound applications.

System	GBit/Xeon	Myrinet/Opteron	SGI Altix 3700	NEC SX6+
Basic building block	2-way SMP node with 1 memory path	2-way SMP node with 2 memory paths	4-way SMP node with 2 memory paths	8-way SMP node
CPU	Intel Xeon 2.66 GHz	AMD Opteron 2.0 GHz	Intel Itanium2 1.3 GHz, 3 MB L3	565 MHz
Peak performance per CPU	5.3 GFlop/s	4.0 GFlop/s	5.2 Gflop/s	9 GFlop/s
Memory bandwith per building block	4.3 GByte/s	2x 5.4 GByte/s	2x 6.4 GByte/s	8x 36.0 GByte/s
Interconnect	Cisco 4503 GBit Ethernet switch	Myrinet2000	SGI NUMALink3 2x 1.6 GB/s bidirectional	NEC IXS crossbar 8 GB/s bidirectional
Operating system	Debian Linux 3.0	SuSE SLES 8 Linux	Redhat AS2.1 with SGI Propack 2.4	SUPER-UX
Compiler	Intel ifc 7.1	PGI 5.0	Intel efc 7.1	Native NEC SX

Table 1: Details of platforms and compilers used in the benchmarks.

Application scenarios

Two representative application codes have been chosen, which have been developed at the Institute of Fluid Mechanics (LSTM-Erlangen, Prof. Dr. Durst) and at the University of Erlangen-Nuremberg. Both programs are currently in intense use ranging from single processor runs on Intel Xeon through moderately parallel jobs on SGI Altix (RRZE) up to high end simulations on 512 processors (64 nodes) of the Hitachi SR8000 TFlop/s system at LRZ Munich. The codes have been ported and optim ized by the HPC group of RRZE. Concerning the computational requirements of CFD, usually two scenarios can show up:

Speed-Up: A problem of *fixed size* should be solved as quickly as possible. Time-to-solution is a critical point for applications e.g. from engineering.

Scale-Up: The problem size is scaled with the number of CPUs/compute nodes used. In this context, basic turbulence research can serve as an example.

Architecture	1 CPU	2 CPUs
Intel Xeon	1.9	3.0
AMD Opteron	2.8	5.7
Intel Itanium2	5.0	7.2
NEC SX6+	38	74

Table 2: Speed-up performance within 2-way nodes given in million lattice site updates per second (MLup/s). On the Itanium2 Altix system the two CPUs chosen for this measurement share one path to the memory.

Scalability of lattice Boltzmann simulations

Owing to the high scientific potential for large scale applications, we have chosen the lattice Boltzmann method (LBM) [3] as a first test case. LBM is a recent method from CFD which is characterised by algorithmic simplicity owing to the explicit nature of the algorithm and equidistant Cartesian grids.

In Figure 1 the parallel performance of both scenarios is presented. A handy performance unit for LBM is million lattice site updates per second (MLups), where 5 MLups ~ 1 GFlop/s holds for our calculations.

Of course the *scale-up* problem is well suited for cluster configurations, where parallel efficiencies of more than 80% on 64 processors can be achieved: A low ratio of communication vs. computation was chosen which remains constant in our application if the problem size is scaled linearly with processor count. Due to higher

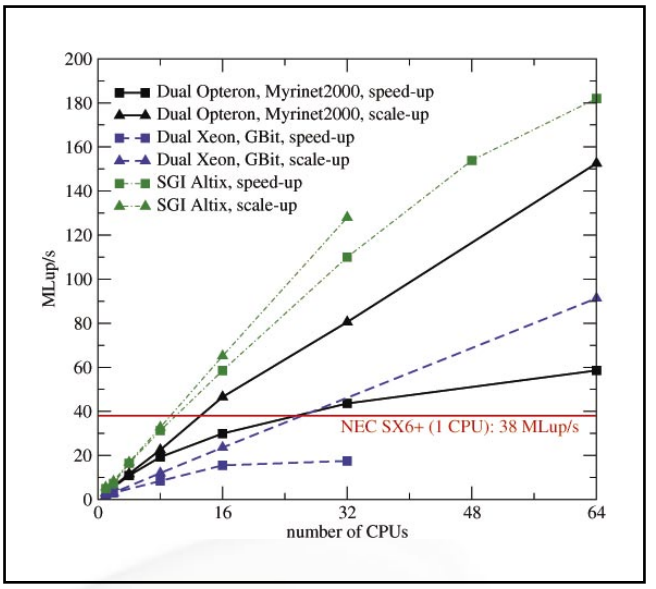

Figure 1: Scalability of LBM for modern clusters and HPC systems. The domain size is 256×129×128 for speed-up and 128^3 per processor for scale-up tests.

single processor performance (see Table 2) and almost perfect scalability the SGI Altix system is significantly ahead of the COTS clusters.

The *speed-up* case is a more appropriate test for the balance of parallel computer architectures. With increasing processor count, the computational domain per processor decreases while the communication per processor remains almost constant. Figure 1 clearly demonstrates that the GBit cluster does not scale at all beyond 16 processors, while the Myrinet interconnect does significantly better. Nonetheless, a Myrinet cluster with 24 Opterons is required to achieve the same performance as a single NEC SX6+ vector processor. The SGI Altix scales very well even for the speed-up case sustaining a parallel efficiency of roughly 80% on 64 processors (if using the two processor performance as the base and thus ignoring the bandwidth problem within the 2-way nodes as shown in Table 2). In comparison to the cluster interconnects, the NUMALink3 is in a class of its own, providing roughly 5 times the MPI bandwidth and only 20% of the MPI latency of Myrinet2000. Although the SGI Altix is much better balanced than the COTS clusters, it suffers from limited scalability on node level compared to AMD compute nodes (see Table 2).

Large Eddy Simulations using LESOCC[1]

LESOCC is a finite volume code with the strongly-implicit solver (SIP-solver) according to Stone [3] as the core routine.

Owing to the complex communication pattern required by the numerical scheme, there is a severe restriction on scalability. Thus we have chosen a typical workload with 12 MPI processes. In Table 3 we show the relative performance of the COTS clusters and the SGI Altix relative to 12 NEC SX6+ vector processors. The vector system is still a factor 5 ahead of SGI Altix, while the COTS clusters only achieve half of Altix performance. Interestingly, performance can also be significantly improved on the Opteron cluster if only one processor per node is used. However, the effect is not as large as on the Xeon. A detailed analysis of the communication pattern on the Altix has identified that a large fraction of runtime is spent in MPI communication with low performance. Whether this is an effect of bad load balance or inefficient communication patterns is currently under investigation. Figure 2 gives a typical example of computations done with LESOCC.

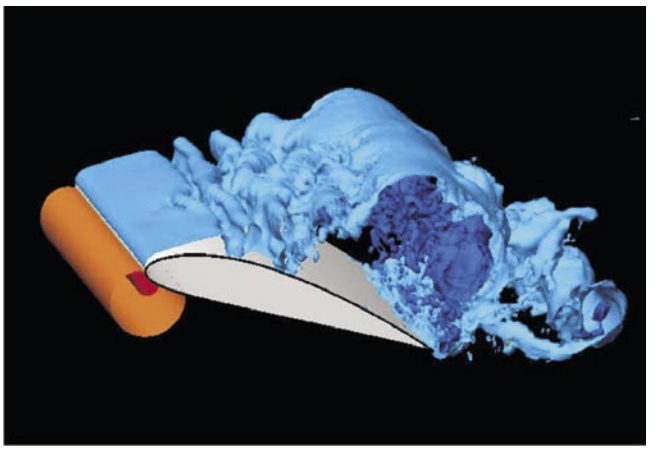

Figure 2: Isosurfaces of constant pressure of the flow around an inclined airfoil using LESOCC (Re=20000). Computations have been performed on the Hitachi SR8000-F1 at LRZ Munich. (Courtesy of N. Jovicic, LSTM - Erlangen)

Conclusions

COTS present a highly competitive HPC architecture if the overall system balance is only of minor importance. However if network capability affects the performance of parallel applications, the SGI Altix is the system of choice. With a performance equivalent of 5-8 Itanium2 processors per NEC SX6+ vector processor, SGI Altix has further reduced the gap between cache-based and vector processors.

Concerning our experiences with the SGI Altix, we have found that performance can still vary between identical runs (depending on the load and/or the buffer cache size) and that only minor changes in compiler versions can still have significant performance impact. Thus, we expect a continuous slight performance increase on SGI Altix with improved system parameters and enhanced compiler technology.

Acknowledgements

Part of this work is financially supported by the Competence Network for Technical, Scientific High Performance Computing in Bavaria KONWIHR. We would like to thank CSAR in Manchester for providing access to their computing facilities and helpful support.

References

[1] Top 500 list, available at http://www.top500.org (November 2003).
[2] L. Oliker, et al., Evaluation of cache-based superscalar and cacheless vector architectures for scientific computations, in: Proceedings of SC2003, CD-ROM, 2003.
[3] S. Chen, G. D. Doolen, Lattice Boltzmann method for fluid flows, Annu. Rev. Fluid Mech. 30 (1998) 329–364.
[4] H.L. Stone, Iterative solution of implicit approximations of multidimensional partial differential equations, SIAM J. Numerical Analysis, 5 (5), 1968

(Courtesy of N. Jovicic, LSTM-Erlangen)

(Footnotes)
[1] Large Eddy Simulations on Curve Linear Coordinates (LESOCC)

System	12 MPI processes using	Perform. relative to NEC
Xeon/GBit Cluster	6 dual-way nodes	0.06
Xeon/GBit-Cluster	12 dual-way nodes	0.10
Opteron/Myrinet Cluster	6 dual-way nodes	0.15
Opteron/Myrinet Cluster	12 dual-way nodes	0.18
SGI Altix	6 dual-way nodes	0.22
NEC SX6+	2 eight-way nodes	1.00

Table 3: Performance numbers of LESOCC relative to the NEC SX6+ measurements. If 12 dual-way nodes are used in the COTS clusters, only 1 CPU per node is used for computation while the second CPU is idle.

Three-Dimensional Simulation of Rotor-Stator Cavity Flow

T. J. Craft, S. E. Gant, H. Iacovides and B. E. Launder (School of Mechanical, Aerospace and Civil Engineering, University of Manchester)
E. Serre (MSNM-GP, Université d'Aix-Marseille)

There have been numerous numerical simulations and experimental studies of flow between rotating and stationary discs with a stationary shroud and no through-flow (a "rotor-stator cavity"). The flow has significant industrial applications, such as internal gas-turbine flows and computer hard disks, and the geometry is relatively simple. Most of the numerical simulations to date have treated the flow as axisymmetric. However, recent experimental flow visualization studies have shown that large coherent eddy structures exist within the cavity at certain operating conditions.

Some of these structures have since been simulated, albeit at low Reynolds number, using Direct Numerical Simulation (DNS) by Serre and co-workers at Marseille University [1].

The purpose of the current work is to investigate whether less costly Unsteady Reynolds-Averaged Navier-Stokes (URANS) and Large-Eddy Simulation (LES) models are also able to predict these time-dependent flow features.

Two sets of calculations have been performed using the parallel MPI version of the in-house CFD code, STREAM [2]. The first used a relatively coarse grid (61 x 50 x 38 nodes in the radial x circumferential x axial directions) and an embedded-grid wall-function which solves numerically boundary-layer-type equations [3]. The ratio of the disc-spacing (H) to radius (R) of the flow domain tested was H / R = 0.126 and the Reynolds number, based on the maximum rotor velocity and disc spacing was Re = Ω RH / ν = 112,000. These conditions

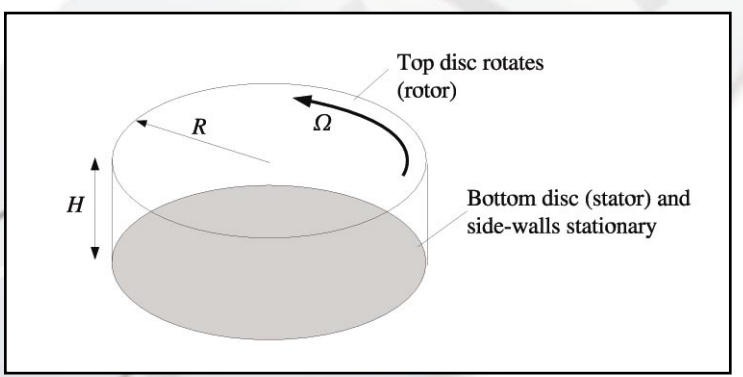

Figure 1: Computational Model

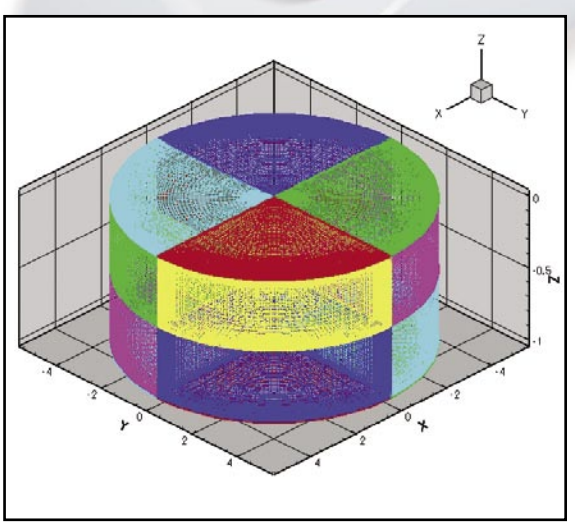

Figure 2: 16 Block Computational Grid (150 x 100 x 141)

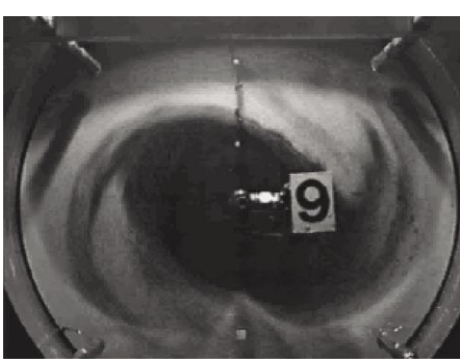

Figure 3: Experimental ink-in-water experiments for H /R = 0.126 case from Czarny et al.

corresponded to a particular experiment undertaken by Czarny et al. [4] in which a distinct two-lobed coherent eddy structure was observed (see photo in Figure 3).

Nine different turbulence models were tested:

- "No-model" approach (i.e. only molecular viscosity, in conjunction with an upwind-biased convection scheme)
- Linear k - ε model of Launder & Sharma
- Linear production model of Guimet & Laurence
- Organised Eddy Simulation (OES) model
- Cubic non-linear *k* - ε model of Craft et al.
- Speziale, Sarkar & Gatski (SSG) differential stress model
- Two-Component Limit (TCL) differential stress
- Filter-based RANS model of Johansen et al.
- Smagorinsky LES

An ongoing second set of calculations using a finer 2.1 million-node grid (150 x 100 x 141) is resolving the near-wall flow (i.e. not using wall functions). The domain geometry is H / R = 0.195 and Reynolds number, Re = 166,000. Three-lobed structures were observed in the experiments at these operating conditions. Calculations are being performed using 16 processors on "Newton" with a reasonable speed-up of 14.3. At present, results have been obtained using a "no-model" approach and computations using a k - ε model are in progress.

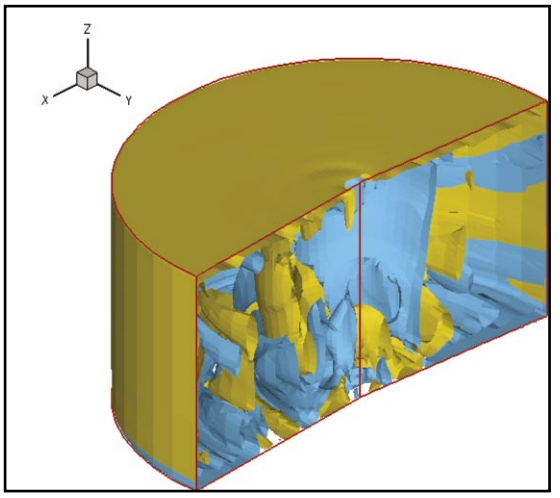

Figure 4: Instantaneous radial velocity isocontours: yellow = positive, blue = negative

In the thinner cavity (H / R = 0.126) the "no-model" approach produced a disturbed, non-axisymmetric, flow field. A plot of the instantaneous radial velocity isocontours (see Figure 4) shows the positive (outward) radial velocity in the Ekman layer near the upper rotor surface and predominantly negative (inward) radial velocity in the layer near the lower stator surface. Within the bulk of the cavity itself there are numerous, disordered regions of positive and negative radial velocity. Spiral structures, which have been observed in DNS results for the annular rotor-stator cavity, cannot easily be discerned.

None of the URANS and LES models tested in this geometry produced coherent flow structures. The calculations were started from the disturbed "no-model" result and over time, the fluctuating velocity field decayed to practically steady, axisymmetric flow.

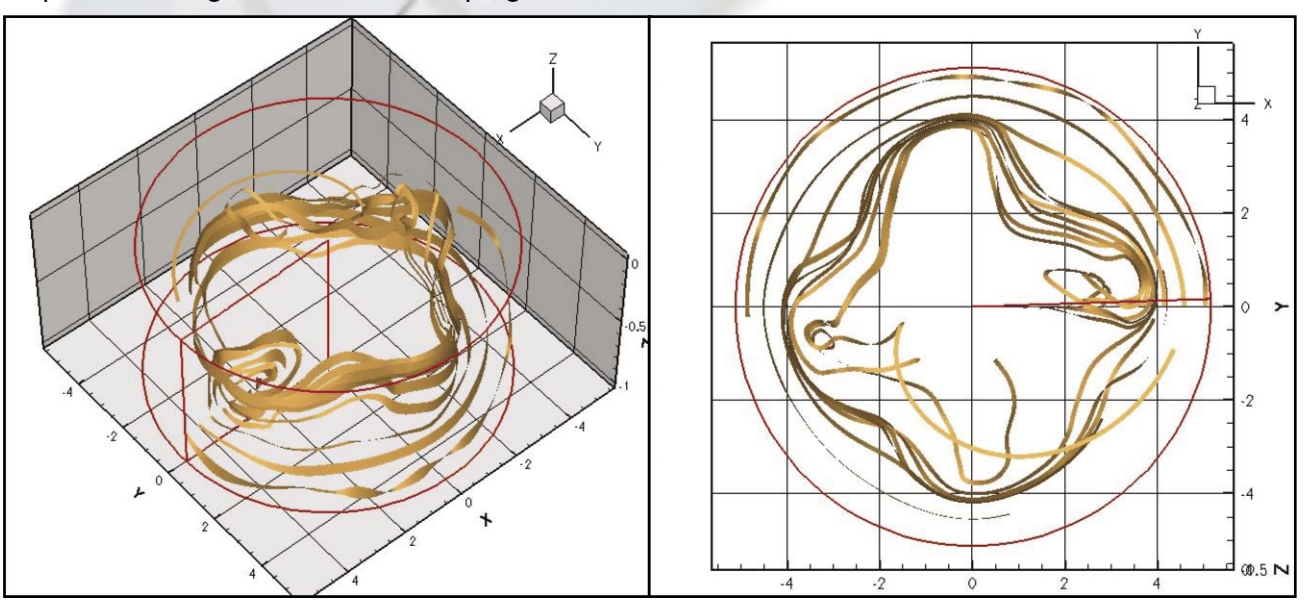

Figure 5: "No-model" 3-D and 2-D steamtraces after 15 revolutions

A very small degree of unsteadiness was observed with those models that gave lower eddy-viscosity (namely the OES and filter-based models).

In the deeper cavity (H / R = 0.195), the "no-model" approach gave an unsteady three-dimensional flow field. After around 15 revolutions of the rotor a distinct four-lobed pattern was observed in the pressure and streamtrace plots (see Figure 5). A further 15 revolutions later, the four-lobed pattern merged into two lobes (see Figure 6). The magnitude of these disturbances were far greater than those which appeared previously in the thinner cavity. Recent results have indicated that as the calculation progresses still further, the two-lobes merge into a single-lobed axisymmetric flow pattern. Preliminary results using a standard low-Reynolds-number k - ε model show spiral structures emerging in the velocity contours near the stator surface, similar to those observed in the DNS rotor-stator simulations.

These preliminary results show that large coherent turbulent structures can be resolved in rotor-stator flows. At present, it is difficult to conclude whether the different results in the shallow and deeper cavities are a consequence of flow in the deeper cavity being naturally more unstable or a result of the numerical methods employed being different (grid resolution, turbulence model and wall treatment). In the experiments, they noted that coherent structures appeared more freely in the deeper cavity.

It is also not yet clear what causes the magnitude of the resolved unsteady motions to diminish in the thicker cavity. The calculation was started from an initially zero flow field with the rotor spinning at full-speed and therefore resembled computationally an impulsively-started rotor-stator flow. Whether the structures are only a transient phenomena linked to the initial conditions or whether they are being damped by some numerical diffusion effects is being investigated.

References

[1] E. Serre, O. Czarny, H. Iacovides, P. Bontoux and B. E. Launder Precessing vortex structures in flow within rotor-stator disc cavities: DNS and visualization studies. In I. P. Castro, P. E. Hancock and T. G. Thomas (editors) Advances in Turbulence IX: Proceedings of the Ninth European Turbulence Conference, Barcelona, 2002

[2] F. S. Lien and M. A. Leschziner. A general non-orthogonal finite-volume algorithm for turbulent flow at all speeds incorporating second-moment turbulent-transport closure. Comp. Meth. Appl. Mech. Eng., 114:123-167, 1994

[3] T. J. Craft, S. E. Gant, H. Iacovides and B. E. Launder. A new wall function strategy for complex turbulent flows. Numerical Heat Transfer, Part B: Fundamentals, 45(4):301-318, 2004

[4] O. Czarny, H. Iacovides and B. E. Launder. Precessing vortex structures in turbulent flow with rotor-stator disc cavities. Flow, Turbulence and Combustion, 69:51-61, 2002.

Contact Details

Simon Gant, School of Mechanical, Aerospace and Civil Engineering, University of Manchester, Manchester M60 1QD simon.gant@umist.ac.uk

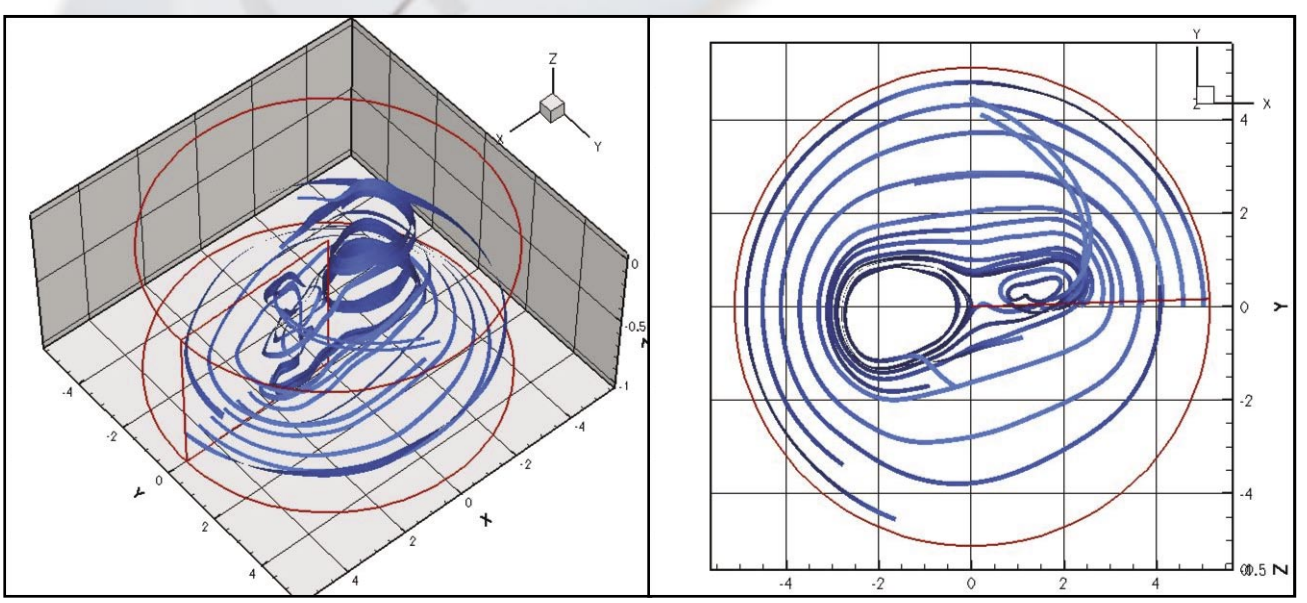

Figure 6: "No-model" 3-D and 2-D streamtraces after 30 revolutions

Manchester Computing Wins Bid to Host the Access Grid Support Centre

Michael Daw
Access Grid Support Centre Manager

Regular readers of CSAR Focus will be familiar by now with Access Grid. Manchester built the first node in the UK and has been prominent in SC Global and in the worldwide Access Grid development community.

Growth of UK Access Grid

The number of Access Grid nodes in the UK is increasing fast. From an initial rollout of 12 nodes, there are now over 20, with more being commissioned and installed every month as new institutions and disciplines realise the potential of this technology. The University of Manchester itself is committed to building another six nodes across campus as part of the investment in e-Science North West, the aim being to encourage all disciplines to play their full part in the potential of e-Science.

Initially, support for Access Grid was informal – the number of sites in the UK was small enough for there to be a recognisable community where everyone helped each other. However, with the number of sites rising quickly, there has been a need for a more formal support solution.

Manchester Wins Bid for AGSC

And so the Access Grid Support Centre (AGSC) was born, following a successful bid by the University of Manchester. The bid was supported by every one of the National and Regional e-Science Centres, which shows how our expertise has been acknowledged and recognised by the rest of the community. UKERNA oversee and manage service provision, as they do for the related JANET Videoconferencing Service (JVCS). The attraction of the original bid was enhanced by our proposal to make use of the CSAR Frontline helpdesk as first line of support for the AGSC.

On-Demand Support and Training

The AGSC provides on-demand support for users and potential users, addressing topics such as procurement advice, help on AGSC services and general troubleshooting. A major role is to give in-depth training through a series of courses, which covers issues from introductions to Access Grid through to debugging and fault-finding techniques.

Quality Assurance Tests

The AGSC also run Quality Assurance (QA) tests to aim for improvements in the perceived quality for users. QA tests check audio, video and network quality at remote nodes, as well as areas specific to Access Grid (as opposed to traditional videoconference) facilities, such as use of the Multi User Domain (MUD) – text chat software used for sideband conversations – and shared presentation software.

AGSC Services

In addition to this, the AGSC operates a number of services for registered users to improve the Access Grid experience. A Virtual Venue Server supports UK-only venues in addition to the standard venues. A persistent Multicast-Unicast Bridge helps those sites which do not yet have multicast enabled, or which are experiencing difficulties in this area. For the first time for many UK sites, the AGSC offers IG Pix, which allows the presentation of Microsoft PowerPoint over a web browser without the need to distribute slides beforehand. IG Pix is easy to use and receiving sites have no need for specialised client software. Also available is a facility to record meetings and an experimental (i.e. not production quality) Access Grid/H.323 Bridge. These services run on software written by InSORS Communications Inc.

Supporting Collaboration

It is anticipated that in the coming years, collaborative tools will come to be developed that enhance, or that can be used alongside, the Access Grid. The AGSC is likely to be a major player in ensuring that these tools are rolled out so that users of Access Grid facilities can extract maximum benefit from them. The aim of the AGSC is to help the UK to fully realise the potential of Access Grid as a tool for highly effective remote collaboration.

For more information, see http://www.agsc.ja.net

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/about/bulletin

CSAR Focus is published twice a year and is also available to view on the Web - http://www.csar.cfs.ac.uk/about/csarfocus. 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.

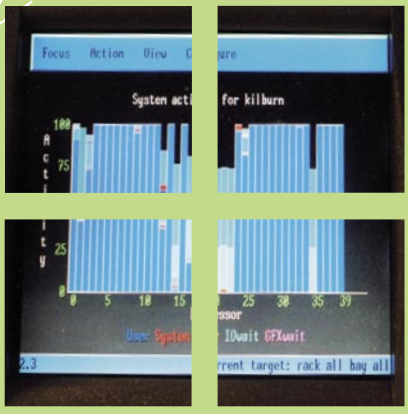

ISBN 1470-5893