

EPCC News

ISSUE 68, AUTUMN 2010



Fireworks over Edinburgh Castle.
Image courtesy 'Edinburgh Inspiring Capital':
www.edinburgh-inspiringcapital.com

Visitor programmes
HPC facilities
Consultancy
Research
Software
Training

20
years
of EPCC

In this issue...

- 2** EPCC at SC'10
- 3** PlanetHPC: The roadmap for HPC in Europe
Faster simulation codes on HECToR
- 4** EPCC visitor programmes
- 6** TEXT: Towards EXaflop applicaTions
PRACE
- 7** HECToR: the journey to multi-core
- 8** SPRINT: speeding up research analysis
- 9** Supercomputing liquid crystal superstructures
- 10** MSc in High Performance Computing
- 11** FPGA researcher wins Industrial Fellowship
All Hands 2010
- 12** Celebrating 20 years of EPCC

Editorial

Iain Bethune



In the spirit of EPCC's 20th birthday celebrations, the autumn 2010 edition of *EPCC News* looks both to the past and the future.

To celebrate 20 years of EPCC we held the successful EPCC@20 event in September, and our Director, Arthur Trew, considers our achievements over the years. One of EPCC's longest-running activities has been our visitor programmes (TRACS and HPC-Europa), and Catherine Inglis reviews over 15 years of visits to Edinburgh. More recently, Judy Hardy highlights some of the research work completed over the summer by graduates of our MSc in High Performance Computing.

Looking forward, Liz Sim, our User Support Project Manager, reports on the installation of the HECToR Phase 2B, the UK's national supercomputing service, and plans for further upgrades.

EPCC will be attending the Supercomputing 2010 conference in New Orleans, and Adrian Jackson shares some of the range of events we are involved in. Finally, we show how EPCC is helping shape the future of HPC in Europe through the PRACE partnership and the PlanetHPC networking activity.

As always, it is a busy time at EPCC as projects are completed and new ones start up. This issue is packed full with reports of what our staff have been up to over the last few months: from successful software development and optimisation to fundamental research in condensed matter physics and field-programmable gate arrays.

I hope you enjoy this issue of *EPCC News*, and join with me in wishing EPCC a happy 20th birthday, and further successes over the next 20 years!



EPCC at SC'10 **Adrian Jackson**

*New Orleans skyline.
Image: Alex Demyan and NewOrleansOnline.com*

EPCC booth: 1623 EUFORIA booth: 4457 PRACE: 4021



With tutorials, posters, BOF sessions, EPCC's own exhibition booth and involvement in a number of other booths as well, and of course the jazz culture of New Orleans, EPCC is looking forward to a busy SC this year.

Together with our friends at Cray, we will be running a tutorial on Co-array Fortran, one of the new breed of Partitioned Global Address Space (PGAS) languages. We are also involved in a tutorial providing an introduction to the shared memory programming standard OpenMP, a model of growing importance as multi-core and many-core processors become more ubiquitous.

We are hosting a couple of Birds-of-a-Feather sessions this year, so if you're interested in high performance workflows or data intensive research, please come along.

The poster sessions will also be interesting. Applications Area Chair, EPCC's Michèle Weiland, says: "As you would expect, the quality of the submissions was very high and we have put together an exciting and varied programme that showcases research into novel applications and technologies from across the globe."

EPCC is heavily involved with the poster sessions, with three posters from the centre: the DEISA Benchmark Suite; enabling applications on the PRACE prototypes and GPU acceleration of scientific applications.

EPCC staff are helping out on the PRACE booth (number 4021) and organising the EUFORIA projects' SC10 presence (booth number 4457). We will also have our own booth in the exhibition hall (number 1623) as usual and we invite you to visit and find out more about our work.

Big Easy here we come!

PlanetHPC: The roadmap for HPC in Europe

Mark Sawyer



In the last issue of EPCC News we reported on the launch of PlanetHPC, a Framework Seven Support Action to produce a roadmap for HPC research in the EU.

Membership of the PlanetHPC network stands at 200, with around 30% of those members coming from industry. These are encouraging numbers, as the aim is to have the views of a broad cross-section of industry and academia represented in the roadmap.

We have carried out a series of interviews with leading people from the HPC industry. These can be found on the project website: www.planethpc.eu. The project has also given presentations at the HiPEAC Innovation Event held in Edinburgh in May and at the ICT2010 conference in Brussels in September. We will be carrying out a series of focussed workshops in the coming months, and publishing the first draft of the roadmap early in 2011.

To ensure that the project covers the widest possible range of industries and applications, and aligns with EC policy on such

issues as sustainability, health and competitiveness, EPCC has been in contact with trade and industry associations such as the Innovative Medicines Initiative, the European Factories of the Future Research Association (EFFRA), and bodies representing the automotive, aerospace and energy sectors.

One of PlanetHPC's aims is to encourage ambitious ideas from end-users, so that European research on HPC can be directed accordingly, and we encourage anyone with ideas for challenging and novel applications of HPC to contact us.

The latest information on the project can be found on the project website, including an interview with Thierry Van der Pyl, Director of Components and Systems Research at the European Commission.

www.planethpc.eu

Speeding up simulation codes on HECToR

Alan Gray

EPCC has been awarded funding to further enhance the performance of two popular scientific applications on HECToR, the UK's national supercomputing service.

CP2K and VASP are used for a wide range of materials science simulations. Both applications are important for improving our understanding of key scientific and environmental processes and aiding development of useful new materials. They are among the most heavily used packages on HECToR. Performance improvements in these codes can therefore have a dramatic impact on the overall effectiveness of the HECToR service.

EPCC has been actively involved in developing CP2K since 2007, enabling the code to more effectively exploit the architecture of large-scale multicore systems such as HECToR.

Communication patterns within CP2K have been optimised, the code has been restructured to improve load balance, and a mixed mode programming model has been introduced. This work has resulted in impressive performance gains for the code.



A new project started in October to build on this work and further enhance performance, focusing on the linear algebra steps which dominate many of the large simulations.

The new VASP project, also building on previous work within EPCC, will improve the Hybrid DFT/HF functional component – a recent addition to VASP which overcomes the limitations of traditional density functional approaches and improves simulation accuracy. A mixed-mode programming model will be introduced into this part of the code to alleviate scaling limitations and allow the code to utilise more of HECToR's computational resources.

Both projects will run throughout 2011 and are funded under the HECToR distributed Computational Science and Engineering (dCSE) support scheme.

Current or potential users of HECToR who are interested in collaborating with EPCC on future dCSE projects should contact Alan Gray: a.gray@epcc.ed.ac.uk

EPCC visitor programmes: creating a lasting impact on research careers

Catherine Inglis

EPCC has been running European visitor programmes for most of its 20-year history, first through TRACS (1993-2003) and then HPC-Europa (2004-2012).

We have supported more than 750 collaborative visits from a total of 28 European countries, with visits hosted by over 100 researchers from 23 UK institutes. Visits often lead to long-term collaborations, with joint publications and further reciprocal visits; some even lead to visitors obtaining a research post in their former host department.

For many visitors, TRACS and HPC-Europa have offered a unique opportunity to gain access to world-class HPC facilities. In 1994, some of our earliest TRACS visitors had access to Europe's fastest supercomputer at that time, and the first UK National Computing Service on a parallel computer – the 256-processor Cray T3D (later upgraded to 512 processors), with a peak performance of 38.4 gigaflops.

As if to prove the saying “plus ça change, plus c'est la même chose”, while the world of technology has changed enormously since 1994, research visitors in 2010 are still drawn here at least in part by the prospect of using one of the most powerful computers in Europe. The HECToR phase 2b machine has a theoretical peak performance of over 360 teraflops – approximately 10,000 times that offered by the T3D in 1994.

Here some former visitors tell us how their visits have been of benefit to them.

Further information

Find out more about EPCC's visitor programme:
www.epcc.ed.ac.uk/visitors

More information about HPC-Europa, including the application closing dates and details about how to apply, can be found at:
www.hpc-europa.eu

Prof Ernst von Lavante Faculty of Engineering, University of Duisburg-Essen TRACS visitor: 1995, 1996 and 1997



In the mid-1990s, the need for a further increase in computational speed for large computational fluid dynamics projects became obvious.

One possible solution was the parallelisation of the corresponding codes. Two of my PhD students hand-parallelised (using assembly language)

a 2-D Navier-Stokes solver, including the post-processing. It was dreadful. It worked but I was thinking: there must be a better way. And there was.

Around the same time, EPCC was equipped with the most powerful and practical parallel computer in Europe – the T3D. I do not remember how I came across that information, but I immediately thought: this is my chance. I had a real monster of a program, a 3-D unsteady Navier-Stokes solver with non-equilibrium chemistry, used for the simulation of hydrogen-air combustion in a supersonic combustion ramjet (scramjet).

I finally had an opportunity to work on a real parallel computer with – and this was the amazing feature – a debugger that actually was worth something! Add to that a useful version of MPI and the very helpful, friendly and efficient staff at EPCC, and you got all the ingredients for success.

But, that was not all – I was also discovering beautiful Scotland: the gentle Lowlands, the wild Highlands and the highly cultural surroundings of Edinburgh.

The entire series of stays in Edinburgh at EPCC was a great success. Most importantly, with the help of the knowledgeable, helpful and competent people of EPCC and the friendly help from the Mechanical Engineering Department, I finished the parallelisation of my code called ACHIEVE with very good results.

The code was used in its fullest form on the Cray T3E in Jülich in Germany on what was at that time the largest computer in Germany. Using block decomposition, it was distributed to 1024 processors with excellent efficiency. The same strategy was applied to my other codes, being executed on everything from PC-cluster to hypercube to massively parallel systems, and it still works like a charm.

Furthermore, I am still in close contact with the nice people in the Mechanical Engineering Department, with whom I am now establishing an ERASMUS exchange program.

Thanks EPCC and thanks Edinburgh.

Prof Andrew Curtis
School of Geosciences, University of Edinburgh
TRACS visitor: January–April 1997

I visited EPCC under the TRACS programme in 1997 to carry out work designing geophysical surveys or experiments with my host Bob Pearce in the Department of Geology and Geophysics.

Having access to the Cray T3D made a difference to my research because survey design is a ‘macro-optimisation’ problem: before collecting experimental data, we design the data acquisition such that the various inference problems that we may have to solve after collecting data are expected to provide as much information as possible. Hence, design problems require far

more compute-power than simply solving any single inference problem.

My TRACS visit was very useful for me personally. At the time I was a postdoctoral Fellow at the University of Utrecht, The Netherlands. The expertise and experience that I gained under TRACS helped me to secure a job in Schlumberger Cambridge Research, UK, and thus to my current post in which I have again returned to Edinburgh as the Total Professor of Mathematical Geoscience!

Dr Massimiliano Porrini
Foundation for Research and Technology Hellas, IESL-FORTH
HPC-Europa visitor: April–July 2010



My participation in HPC-Europa2 was a great scientific experience for me.

I applied for the programme while completing my Marie Curie postdoc fellowship at the Institute of Electronic Structure and Laser, Foundation for Research and Technology Hellas (IESL-FORTH) in Crete, Greece, where, with PI Prof Stavros C. Farantos, I

undertook a study of free energy calculations of Cytochrome c Oxidases interacting with diatomic ligands, using a perturbative approach.

Free energy surfaces (FES) are one of the most important features in studying macromolecules, as from these functions we can extract the majority of the properties of macroscopic states of matter. Performing a complete mapping of the FES of systems is very demanding in terms of computing time, therefore we decided to try to parallelise our code, to subsequently implement it on a supercomputer. Specifically we reckoned that HECToR, the Cray machine at EPCC, could match and satisfy perfectly our requirements.

With the valuable help of my EPCC contact, Dr Andrew Turner, I was able to parallelise the code through a task farm (or master/slave) approach. After testing and running preliminary jobs on the small parallel machine Ness, we carried out our production runs on the Cray machine HECToR, utilising both XT4 and XT6 phases.


In conjunction with my HPC-Europa2 host, Dr Perdita Barran of the School of Chemistry at Edinburgh University, I have undertaken another line of research, concerning structural properties of synthetic peptides in gas phase and solution.

This resulted in me obtaining job contracts to allow me to continue working in Edinburgh with Dr Barran for four more months. The resulting work, conducted over a relatively short period of time, has already yielded a publication in *Chemical Communications* journal, with further papers in preparation.

During the time spent in Edinburgh I learned a lot about all the HPC techniques and strategies, by attending all the advantageous courses about MPI, OpenMP and GPGPU held at EPCC.

Moreover the experience of sharing the flat with other visitors from all over Europe was really enjoyable from a social point of view. The EPCC staff organise events almost every week so the visitors can have fun and mix with each other. To conclude, both social and scientific experiences took place in the very lovely town of Edinburgh.

aa3 enzyme from Paracoccus Denitrificans bacterium and the box representative the 3D Cartesian grid of points. The chains A and B are shown in pale green and pale yellow respectively, the water molecules are red balls and the heme groups are depicted with licorice representation. The 2.5 kcal/mol isosurface of the free energy calculated with molecular oxygen is shown in orange.





PRACE: world-class computing for Europe **Alan Simpson**

*Delegates to the
PRACE-1IP kick-off.*

PRACE (Partnership for Advanced Computing in Europe) is a collaboration of supercomputing centres and funding bodies from some 20 countries throughout Europe. It aims to provide a persistent pan-European Research Infrastructure for HPC.

PRACE forms the top level of the European HPC ecosystem, providing world-class systems to strengthen Europe's scientific and industrial competitiveness. PRACE will maintain a pan-European HPC service of up to 6 top-of-the-line leadership systems (Tier-0), well integrated into the European HPC ecosystem. Each system will provide computing power of several Petaflops in the midterm. A scientific steering committee will advise PRACE and operate alongside a bespoke peer review process through which access to the Tier-0 resources will be granted, based on scientific excellence. This infrastructure is managed as a single European entity.

The first Tier-0 system is the IBM Blue Gene/P in Forschungszentrum Jülich, Germany, which has almost 300,000 processors. The second will be a Bull system located in a new computing centre near Paris: the Très Grand Centre de Calcul. Its first phase will be installed by the end of 2010.

PRACE collaborates closely with DEISA (Distributed European Infrastructure for Supercomputing Applications), which provides access to Tier-1 systems and the two projects are working together on middleware for ecosystem integration.

The PRACE 1st Implementation Phase (PRACE-1IP) started in July 2010 and will run for two years. It involves the partners delivering more than 2000 person-months of effort on activities including: training and dissemination, technical operation, industry and HPC ecosystem relations, applications enabling, procurement support and future technologies.

The UK is represented in PRACE-1IP by EPSRC and the technical work is performed by EPCC and STFC. EPCC leads the work package on petascaling of applications: the largest work package within PRACE. We are also collaborating on the bid for the PRACE 2nd Implementation Phase. PRACE is partly funded by the EU's 7th Framework Programme.

www.prace-project.eu

PRACE will be at SC'10 in booth 4021. If you are attending, please drop by and say hello.

TEXT: Towards EXaflop applications **Mark Bull**

In the near future, increases in compute power seem set to come primarily from increasing core counts. This poses a problem for applications: the traditional ways of programming HPC systems seem set to break down at very high core counts.

This is particularly true for MPI. The design of its interface makes it difficult to scale up to millions of cores if we continue, as in the past, to run one MPI task on every core. Several solutions have been proposed. An obvious one is to combine MPI between nodes with OpenMP, exploiting the shared memory parallelism across the cores within a node. But this too has its problems as it is difficult to avoid OpenMP threads sitting idle while MPI communication is taking place.

The TEXT project is exploring an alternative solution by using SMPs instead of OpenMP. SMPs is a directive-based, task-orientated API developed by Barcelona Supercomputing Centre in combination with MPI. The MPI+SMPs model makes it much easier to exploit asynchrony between compute tasks and communication tasks, and can therefore result in higher scalability. TEXT will install the MPI+SMPs environment

at several partners' HPC facilities (including HECToR) and demonstrate how seven real and relevant applications/libraries can be improved using it. The codes fall in the areas of basic linear algebra libraries, geophysics, plasma physics, engineering and molecular dynamics.

EPCC will provide support for external users, and develop synthetic benchmarks which will be used to ensure the quality of MPI+SMPs implementations. Other activities include improving the support for the MPI+SMPs environment, developing debugging and profiling tools, and working on asynchronous algorithms. TEXT will also engage external HPC users by providing access to MPI+SMPs installations, together with in-depth support and training.

Further information

Please contact us if you are interested in porting an application to the MPI+SMPs model: Mark Bull: markb@epcc.ed.ac.uk

Project website: www.project-text.eu

TEXT is an EU-funded project, bringing together partners from Spain, France, Germany, Greece, Switzerland and the UK.



HECToR: the journey to multi-core

Liz Sim

HECToR, the UK national supercomputing service, has been upgraded to include the world's first production Cray XT6 system.

The upgrade to the HECToR service, which is funded by the UK Research Councils, commenced in late May and was completed by early June. HECToR is hosted by EPCC at the Advanced Computing Facility (ACF) outside Edinburgh.

The Cray XT6 – 'Phase 2B'

The Cray XT6, known as Phase 2B to our users, is contained in 20 cabinets and comprises 464 compute blades. Each blade includes 4 compute nodes, and each compute node comprises two 12-core AMD Opteron™ 2.1GHz Magny Cours processors. This amounts to a total of 44,544 cores.

Each 12-core socket is coupled with a Cray SeaStar2™ routing and communications chip. This will be upgraded to the Cray Gemini™ interconnect in late 2010. Each 12-core processor shares 16GB of memory, giving a system total of 59.4TB. The theoretical peak performance of the system is over 300Tflops. The system entered the June 2010 Top500 at position 16.

Users have been online since June, and we are already seeing some fantastic results from the Phase 2B system.

Cray XT4 – Phase 2A

As part of the upgrade, the existing Cray XT4 system (known as Phase 2A) was reduced in capacity. The Cray XT4 has

been reduced in size to 33 cabinets from 60, which brings its capacity down to 3072 quad-core compute nodes from 5664. This amounts to a total of 12,288 cores, each of which acts as a single CPU. The processor is an AMD 2.3 GHz Opteron.

In addition to the Cray XT4, the Cray X2 Vector machine remains in service. The X2 includes 28 vector compute nodes; each node has 4 Cray vector processors, making 112 processors in all. Each processor is capable of 25.6 Gflops, giving a peak performance of 2.87 Tflops. Each 4-processor node shares 32GB of memory.

Next steps

Following on from the Gemini upgrade later this year, our focus will move to the third phase of the HECToR service. Planning is currently underway for 'Phase 3', which is due to run from 2011-2013.

Using HECToR

Examples of HECToR use, including research already completed on the new Phase 2B system, are available here:
www.hector.ac.uk/casestudies

Commercial users with an interest in using HECToR in their research will find details on the EPCC Industry Hub website:
www.epcc.ed.ac.uk/eih/supercomputers

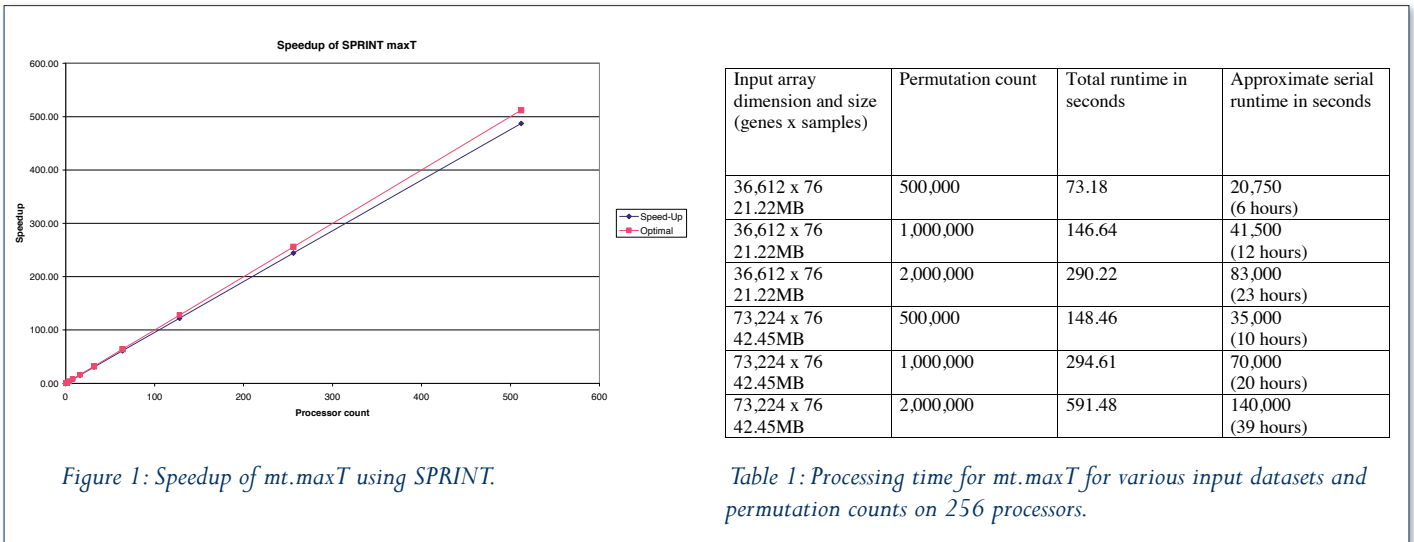
Academic users with an interest in collaborative projects with EPCC on the HECToR service should visit:
www.epcc.ed.ac.uk/research-collaborations/hector-dcse



SPRINT: data analysis in minutes, not days

Terry Sloan

Photo: Mark Sadowski



Input array dimension and size (genes x samples)	Permutation count	Total runtime in seconds	Approximate serial runtime in seconds
36,612 x 76 21.22MB	500,000	73.18	20,750 (6 hours)
36,612 x 76 21.22MB	1,000,000	146.64	41,500 (12 hours)
36,612 x 76 21.22MB	2,000,000	290.22	83,000 (23 hours)
73,224 x 76 42.45MB	500,000	148.46	35,000 (10 hours)
73,224 x 76 42.45MB	1,000,000	294.61	70,000 (20 hours)
73,224 x 76 42.45MB	2,000,000	591.48	140,000 (39 hours)

Table 1: Processing time for *mt.maxT* for various input datasets and permutation counts on 256 processors.

SPRINT is an easy to use, parallel version of the statistical language R. It is now available from the Comprehensive R Archive Network (CRAN), a network of ftp and web servers around the world that store identical, up-to-date, versions of code and documentation for R.

R is very popular with life scientists for gene analysis. *The New York Times* estimates that 250,000 people work with it regularly, not just in biological research but also in companies as diverse as Google, Pfizer, Bank of America, Intercontinental Hotels and Shell.

A common operation in statistical analysis is the permutation test, which calculates the statistical significance of a set of data – how likely it is that similar data would be obtained if the experiment were re-run. In R, the *mt.maxT* function performs this calculation for step-down multiple testing procedures [1], as they are described in Westfall and Young [2].

EPCC’s implementation of the *mt.maxT* function in SPRINT shows near optimal speed-up when using up to 512 cores of HECToR, the UK’s national supercomputing facility (Figure 1). Table 1 compares the serial and SPRINT runtimes for different dataset sizes and permutation counts on 256 processors [3], showing speedup of up to 280 times. This clearly illustrates how SPRINT enables R users to analyse previously intractable datasets in a few seconds or minutes rather than hours or days.

In addition, SPRINT also provides correlation and cluster functionality that provides similar levels of performance to the SPRINT *mt.maxT* implementation.

SPRINT is being developed by EPCC and the Division of Pathway Medicine at the University of Edinburgh through funding from the Wellcome Trust and a distributed computational science and engineering (dCSE) grant for HECToR.

The SPRINT team has recently been awarded a further dCSE grant to add Random Forest and Rank Product to the list of functions that statisticians can use on HECToR.

Further information

SPRINT on CRAN:

<http://cran.r-project.org/web/packages/sprint/index.html>

SPRINT website: <http://www.r-sprint.org>

Footnotes

[1] Y. Ge, S. Dudoit and T. P. Speed, “Resampling-based multiple testing for microarray data hypothesis,” vol. 12, pp. 1–77, June 2003.

[2] Westfall, P.H. and Young, S.S., “Resampling-Based Multiple Testing: Examples and Methods for P-Value Adjustment,” p. 340, 1993. Wiley, New York.

[3] Savvas Petrou, Terence M. Sloan, Muriel Mewissen, Thorsten Forster, Michal Piotrowski, Bartosz Dobrzelecki, Peter Ghazal, Arthur Trew, Jon Hill, “Optimization of a parallel permutation testing function for the SPRINT R package”, *Proceedings of the 19th ACM International Symposium on High Performance Distributed Computing, Chicago, Illinois, USA, WORKSHOP SESSION: Emerging Computational Methods for Life Sciences (ECMLS) Pages: 516-521, 2010, ISBN: 978-1-60558-942-8, http://doi.acm.org/10.1145/1851476.1851551*

Supercomputing liquid crystal superstructures

Kevin Stratford

EPCC has had a long-standing collaboration with the Soft Matter and Statistical Physics group led by Prof. Mike Cates FRS at the University of Edinburgh. Over a period of more than 10 years, this collaboration has worked on simulation of complex fluids by combining expert knowledge in both theory and experiment with state-of-the-art numerical methods and high performance computing.


Our work has been responsible for some ground-breaking discoveries in the areas of fluid mixtures, gels and emulsions, and the associated rheology (or flow behaviour) of these complex fluids. Recently, we have turned our attention to a number of new areas of research, one of which is liquid crystals.

“Liquid crystal” sounds like an oxymoron. On the one hand, a sample of this intriguing substance can flow and be poured from one vessel to another like a liquid; on the other, it has the long-range ordering more usually associated with a solid crystal. In liquid crystals, the important ordering is on a much larger length scale than is seen in a conventional crystal, and the interesting optical properties occur in the visible part of the spectrum (as opposed to the x-ray region for atomic crystals).

This is the origin of the technologically important properties of liquid crystals, and is made use of in everyday devices such as liquid crystal displays.

A liquid crystal consists of a liquid in which rod-like molecules arrange themselves in regular structures. Depending on temperature, and their exact constituent make-up, a wide variety of structures or phases is observed. The all-important feature which characterises these phases is that of the defects, or “disclination lines” in the ordering structure. This defect structure itself can show a crystalline make-up. It is the superstructure of disclinations which is responsible for the optical properties, so it is important to understand how it behaves. As an example, in a recent paper in “Proceedings of the National Academy of Sciences of the USA” [1] we used supercomputers at Edinburgh in the UK and at Argonne National Laboratory in the USA to study the “blue phases,” so-called for their striking blue appearance.

Simulations were used to look at the growth a small nucleus with a particular blue phase structure into a larger crystal. Here, one might expect that the resultant disclination structure would reflect simply the initial nucleus. Surprisingly, one sees instead a rapid proliferation of disclinations forming an amorphous structure (see illustration).



A complex topological structure of defect lines or “disclinations” in a liquid crystal produced by simulation.

Supercomputers prove invaluable in allowing accurate simulations of these systems. There are a number of important features of any such calculation. First, the fluid motion must be represented, requiring the solution of the governing Navier-Stokes equations on a discrete grid. Second, the ordering of the rod-like molecules must be represented: this is done in a coarse-grained fashion where the degree of liquid crystalline ordering is represented by an ‘order parameter.’ This order parameter adjusts locally to minimise the energy of the system, and interacts with the fluid flow. For example, a shear flow will tend to rotate rod like molecules; likewise, the presence of the molecules affects the flow. Furthermore, the superstructures addressed here require that a large system size is used to prevent artificial finite size effects giving rise to unreliable results. All this means that the liquid crystal calculation requires significant computing resources.

Understanding the complex mechanisms which govern the structures seen will improve the chances of exploiting the technological potential of these fascinating materials. In some phases, the superstructure of disclination lines can become very complex, and can even be difficult to pin down experimentally. This hazy picture has given rise to some exotic naming, such as the “blue fog”, a phase which is thought to be amorphous. Even more exotically, it has been suggested that the complex topological defects seen in liquid crystals might provide a model system for those predicted by string theory in the early universe.

Whatever the case, we believe computer simulation will give rise to further exciting discoveries.

[1] O. Henrich, K. Stratford, D. Marenduzzo, and M.E. Cates, *Ordering dynamics of blue phases entails kinetic stabilization of amorphous networks*, *PNAS*, 107, 13212–13215 (2010); doi:10.1073/pnas.1004269107



Over the last ten years we have welcomed more than 150 MSc students from 26 countries.

MSc in High Performance Computing: another successful year Judy Hardy

At this time of year we say goodbye to the MSc class of 2010, who have completed their dissertations and will graduate in November. And we welcome the class of 2011, who are now getting to grips with HPC concepts, message passing programming, shared memory programming, software development and the other topics in the MSc taught courses. This year's MSc class includes 34 students from 13 countries – a record number on both counts!

The dissertation project is one of the highlights of the MSc, giving students the opportunity to work on a four-month independent research project in an HPC-related area.

A wide range of HPC resources are available for student projects, including HECToR, the UK's national service; Ness, a parallel machine used to support the MSc and other local projects; a parallel FPGA machine (incorporating 64 FPGAs and 32 Xeon processors); a top-of-the-range multicore consumer laptop; and a multi-processor GPGPU machine for evaluating the use of graphics processors for general HPC.

Below we highlight some dissertations from 2009/10. Previous years' reports are available on the MSc website: www.epcc.ed.ac.uk/msc

Dissertation highlights 2009/10

Mixed Mode Programming on HECToR **ANASTASIOS STATHOPOULOS**

The majority of modern HPC systems are clusters of SMP nodes. The mixed mode programming paradigm may be more suited to the architectures of such systems and deliver better performance than the traditional programming model (message passing). This project investigated the performance of the mixed mode (MPI + OpenMP) programming model on HECToR. A number of benchmark codes were run on both the Cray XT4 and Cray XT6 systems. It was found that in some cases the mixed mode versions achieved better performance than the equivalent pure MPI version by improving the load balance and, on the XT6 system, reducing congestion on the network.

Porting the z-finder algorithm to GPU **CHRIS JONES**

A high-throughput particle physics code, the z-finder algorithm, runs in the ATLAS detector in the Large Hadron Collider (LHC). The algorithm was ported to GPU using C for CUDA and the performance tested on both Tesla and Fermi architectures. No significant speedup was obtained using low luminosity samples, however with high luminosity samples, speedups of typically 20x were obtained using CUDA streams. These results suggest that it is worth investing further effort investigating porting areas of ATLAS trigger codes onto GPU.

Analyzing and Optimizing Global Array Toolkit for the Cray Gemini Interconnect **VAIRAVAN MURUGAPPAN**

The Global Array Toolkit provides an efficient and portable "shared memory" programming interface for distributed memory systems. The main aim of this project was to analyze and optimise the Global Array Toolkit for Cray XT Systems with the Seastar and Gemini Interconnects. The project included benchmarking the Toolkit together with a detailed analysis of different optimization strategies.

Optimisation and Parallelisation of a Combustion Large-Eddy Simulation Code for High Performance Computation **MICHAEL PETTIT**

Computational fluid dynamics (CFD) is driven by the need for a greater understanding of complex, turbulent flow processes, which require significant computational effort to describe accurately. A CFD code utilising the Large Eddy Simulation approach has been optimised using an improved domain decomposition and parallel communication model. A thorough performance study has revealed a maximum reduction in execution time of 60%. The strong scaling characteristics of the code have also shown a significant improvement compared to the original code, with a parallel efficiency of at least 0.6 maintained on up to 32 processors for one of the test cases.

FPGA researcher wins Industrial Fellowship

The Royal Commission for the Exhibition of 1851 has awarded an Industrial Fellowship to Thomas Perry for his research into ultra high-performance chips – called Field Programmable Gate Arrays (FPGAs). The Fellowship will last three years and includes an enhanced stipend.

Thomas is an EngD student with the Institute for System Level Integration (iSLI) and is part-supervised by EPCC's Mark Parsons. Thomas' research is part of an initiative known as FHPCA, an alliance of high-tech organisations including Scottish Enterprise. The Alliance is using FPGAs to create new levels of performance for the technical computing market. It hopes to revolutionise the development of mission-critical applications and help deal with the massive processing requirements needed in areas such as drug design, defence, seismology, medical imaging and mobile telecoms.

The Royal Commission for the Exhibition of 1851 Industrial Fellowships aim to encourage profitable innovation and creativity in British industry. Only four Industrial Fellowships are awarded each year to students enrolled for EngD degrees.



Dr Khaled Benkrid (SEE), Nigel Williams (1851), Richard Walke (Xilinx), Thomas Perry and Mark Parsons (EPCC).

Further information

Thomas' research project:

www.sli-institute.ac.uk/research/engd-research-xilinx_tp.php

iSLI: www.sli-institute.ac.uk

FHPCA: www.fhpca.org

The Royal Commission for the Exhibition of 1851 Industrial Fellowships: www.royalcommission1851.org.uk/ind_fellow.html

All Hands, Cardiff Neil Chue Hong

All Hands brings together the leading groups working in e-Science from across the UK and beyond. This year's theme was "novel research, new science and enduring impact" and appropriately the opening keynote, "Narrowing the E-Science Rhetoric-Reality Gap", was given by Dan Atkins, W.K. Kellogg Professor of Community Informatics, who chaired the 2009 RCUK Review of e-Science. This presented a world-leading picture: the UK has created a "jewel" – a pioneering, vital activity of enormous strategic importance to the pursuit of scientific knowledge and the support of allied learning. However he also highlighted the challenges to come, not only from financial changes but from the lack of formal training in STEM (Science, Technology, Engineering & Mathematics) within the Whitehall departments and a rigidity in policy created by the "competency trap" where the current best in a field assume they are on top because they are doing the right things, the right way and are therefore reluctant to adopt new things.

The All Hands Meeting is seen as the primary point for knowledge exchange in the subject and EPCC was well represented at this year's event. EPCC Commercial Director Mark Parsons had the honour of introducing keynote speaker Prof. Alex Szalay: lead guitarist of prog rock band Panta Rhei and a lead researcher into data-intensive computing, including the Sloan Digital Sky Survey. Prof. Szalay's work with the late Jim Gray on the GrayWulf architecture is the inspiration for EPCC's own version of a specialist data analysis supercomputer which can optimise Amdahl's number using commodity

hardware components, and will support EPCC's data intensive research strategy.

A number of talks featured work with our project partners on ADMIRE (data mining), NeISS (social simulation), and Teragrid/DEISA (supercomputing applications). EPCC's George Beckett and Rob Baxter presented work done by EPCC on, respectively, "Building the International Lattice Data Grid" and "Tracking Community Intelligence with Trac".

The Software Sustainability Institute, centred at EPCC, launched with a workshop featuring some of the software projects and research communities that have looked at the problem of creating software that lasts. This was complemented by a booth where the SSI pilot collaborative projects in the areas of fusion power, climate policy, journal indexing, geolinked data and crystal structure were on show.

Further information

Links to some of the EPCC presentations can be found here:

www.epcc.ed.ac.uk/news/epcc-at-all-hands-meeting-2010

ADMIRE: www.epcc.ed.ac.uk/projects/admire

DEISA: www.epcc.ed.ac.uk/projects/deisa

DiGS: www.epcc.ed.ac.uk/software-products/digs

NeISS: www.neiss.org.uk

OGSA-DAI: www.ogsadai.org.uk

Software Sustainability Institute: www.software.ac.uk

Software Sustainability workshop: www.allhands.org.uk/session/w3

Worth celebrating: 20 years of EPCC

Arthur Trew

In September 1990, the University of Edinburgh formally inaugurated Edinburgh Parallel Computing Centre as a technology transfer centre with a mission to promote parallel computing across industry and academia.

Although we have subsequently broadened that remit to promote novel computing, and shortened our name to EPCC, we have stuck by the founding vision and used an entrepreneurial approach to grow from a handful of staff to over 70. In the process we have become one of the world's leading high performance computing centres and, probably, the best example of converting methods and techniques honed in research projects into innovative solutions for business.

EPCC was born through the ideas of three people – David Wallace, Stuart Pawley and Ken Bowler – in the, then, Department of Physics. Indeed, EPCC had an unusually long gestation period of 9 years, with the first simulations run on parallel computers in London and the only method of communication by post. Never say that things have not improved, or that network performance has never been worse – it has! Computers at Edinburgh soon followed, however, and the first two ICL Distributed Array Processors (DAPs) enabled researchers in Physics to publish 186 computational science papers in 6 years. This is a truly remarkable feat, especially when you remember that these computers had less performance and memory than a typical mobile phone today.

EPCC developed through links with Roland Ibbett in Computer Science and Richard Field in the University's Computing Service and gained a business-orientated culture through its first Chairman, Jeff Collins, who had been loaned by Lothian Regional Council. Our experience clearly demonstrates the importance of starting from the right place with the right ethos because no-one 20 years ago would have predicted the range of activities that we would have undertaken successfully, or the degree to which our working practices are gradually being adopted across many sectors of University life – a trend which we see continuing as economic pressures rise.

So, what have we done, and what do we do? First, we operate effectively as a not-for-profit business, generating 95% of our turnover from external grants and contracts. Working with industry and commerce is still critical to this, both because it gives us an important differentiating role and because it brings in over 50% of our income.



We run highly-successful training, now largely focused around our MSc in High Performance Computing, which doubled its student intake this year, and our European Visitor programme which since 1993 has brought some 750 researchers to Edinburgh each year to undertake HPC training and a research project with a local collaborator. In parallel with the industrial projects we undertake a wide range of academic research projects, working both on the production of new programming techniques, and on the application of HPC to particular scientific applications. While it is the fate for the former to be forgotten as their outputs become part of the accepted infrastructure – “the plumbing” – and the latter is still most prevalent in the physical sciences, I have no doubts that both are critical for EPCC's future. The increasing reliance on multi-core microprocessors to maintain Moore's Law and the understanding in other areas that computation has something to offer them gives, I believe, strong drivers for a successful future for EPCC.

Last, but not least, we also provide HPC facilities for UK academics. Over the last 20 years we have won three of the four national competitions to host these facilities bringing prestige, funding, a significant upgrade to the University's infrastructure, and close relationships with leading technology providers, notably IBM and Cray. However, while these are probably our most visible product, they are not our most important. EPCC's success depends today, as it has always done, on our staff – their intelligence, drive and commitment. It is a pleasure and a privilege for me to have worked with so many over so many years. I would like to thank them all for their efforts to make EPCC the success that it is today and with such dedication I have no doubts that we will be able to toast another glorious 20 years in 2030.



EPCC is a European centre of expertise in developing high performance, novel computing solutions; managing advanced systems and providing HPC training. Our clients and partners include local and global industry, government institutions and academia.

EPCC's combination of advanced computing resources and expertise is unique and unmatched by any European university.

www.epcc.ed.ac.uk

info@epcc.ed.ac.uk