SPRINT: more runners, fewer hurdles

epcc news

In this issue

Clouds: testbeds, future internet and energy efficiency

CRESTA: reaching for exascale

HPC training at EPCC

Powering business

Meet INDY

Our new HPC cluster makes our on-demand service the most flexible and comprehensive in Europe

Also in this issue
Cray XK6: a thousand GPUs and beyond

From the directors

Welcome to the new look EPCC News. We hope you enjoy it.

The changes we've made to this edition link to the new EPCC website, which will be launched soon. We've put a lot of effort into making both our newsletter and website more accessible and useful for our many clients and collaborators.

As many readers will know, EPCC has worked with industry and commerce since it was established in 1990. We believe we've worked in one way or another with over 750 companies in the past 22 years.

We've been expanding our commercial activities recently, with a more comprehensive, on-demand HPC service. This, and the new users brought by the Supercomputing Scotland initiative with Scottish Enterprise, has convinced us of the need for a flexible industry-focused HPC system. You can find more out about this and Supercomputing Scotland's work on pages 4 and 5.

As a largely self-funding organisation, EPCC has built a strong multi-funding approach over the years. We undertake a huge variety of projects across a wide spread of distributed computing domains. In this issue you'll find articles covering our traditional work in HPC for the sciences, our next-generation research into exascale computing, our focus on the challenges brought by big data, our future internet research and software sustainability.

We hope you enjoy this issue of *EPCC News*. Please get in touch if you would like to find out more about any of the topics it covers.

Mark Parsons & Alison Kennedy
EPCC Executive Directors

a.kennedy@epcc.ed.ac.uk
m.parsons@epcc.ed.ac.uk

Contents

EPCC at SC'12

Come and learn about our exascale research, training opportunities and extensive collaborations with industry.

In booth **1919**, we'll display our broad range of activities, including our HPC service provision, such as hosting HECToR (the UK national supercomputing service) and the DiRAC consortium's ground-breaking BlueGene/Q.

We're also appearing alongside PRACE (Europe's supercomputing project), and joining up with our CRESTA partners to talk about European exascale research.

We're participating in a number of SC'12 workshops and Birds of a Feather sessions too.

Hope to see you there!

SC12
Salt Lake City, Utah

Contact us

www.epcc.ed.ac.uk   info@epcc.ed.ac.uk   +44 (0)131 650 5030

EPCC is a supercomputing centre based at The University of Edinburgh, which is a charitable body registered in Scotland with registration number SC005336.

MSc in HPC: welcome to our new students!

Crystal Lei
yuhua.lei@ed.ac.uk

This year we welcome 33 new students from 17 different countries to our MSc in HPC.

To find out more about the MSc, contact Crystal Lei or see: www.epcc.ed.ac.uk/msc

At the recent two-day induction, our new students were introduced to the University, to EPCC, and most importantly to fellow students and teaching staff.

Not only are the students' nationalities diverse, their backgrounds are diverse too: some have just graduated from undergraduate studies, some have been in employment, while others have been doing research work in other fields. They have all come together now to learn about parallel programming and HPC technologies.

Master's Scholarships

For the first time, the School of Physics & Astronomy has offered two Master's scholarships in High Performance Computing for the academic year 2012–13.

The John Fisher Bursary was established in memory of our EPCC colleague who died in 2007. John was a well-liked and highly respected man who made an enormous contribution to EPCC. The recipient of the John Fisher Bursary this year is James Willis, who was one of the top graduates from the School of Physics & Astronomy at the University of Edinburgh.

The other scholarship was awarded to Vanya Yaneva, an outstanding graduate who previously studied in the University's School of Informatics.

PRACE Advanced Training Centre

Catherine Inglis
cafi@epcc.ed.ac.uk

Cray XE6 Performance Workshop
University of Reading, 20-22 Nov

Software Carpentry
University of Edinburgh, 4-5 Dec

PGAS Programming with UPC and Fortran Coarrays
University of Edinburgh, 9-10 January 2013

Upcoming EPCC training
http://www.epcc.ed.ac.uk/training-education/course-programme

Upcoming events in Europe
https://events.prace-ri.eu/

EPCC is one of six PRACE Advanced Training Centres (PATCs) in Europe, with a remit to run around 10 UK courses a year.

Since April, we have held a Cray XE6 Performance Workshop and run courses on Message-Passing Programming with MPI, Shared-Memory Programming with OpenMP, and GPU programming with CUDA and OpenACC. The courses have been well-attended, with people coming from all around the UK and further afield.

Two more courses will run in 2012:

• Cray XE6 Performance Workshop, (University of Reading). Lectures and practical exercises will introduce new users to the more powerful and advanced features of the hardware and software.

• Software Carpentry Boot Camp (University of Edinburgh). Short tutorials and hands-on practical exercises will increase researchers' productivity by teaching basic computing skills such as program design, version control, testing and task automation. There will be online follow-up sessions for 6–8 weeks afterwards, extending the material from the Boot Camp (see p16).

Registration for both is free.

Introducing INDY

Our new multicore HPC cluster for industrial users

EPCC's on-demand HPC service, **Accelerator**, has been extended with the installation of a new multicore HPC cluster.

INDY is a heterogeneous Linux-Windows HPC cluster aimed at industrial users from the scientific and engineering communities who require on-demand access to mid-range, industry-standard HPC.

The addition of INDY positions **Accelerator** as the most flexible and comprehensive on-demand HPC service in the UK. In addition to INDY, **Accelerator** provides access to two of the UK's most powerful supercomputers: HECToR and BlueGene/Q. HECToR is a 90,112 core, 800 teraflop Cray XE6 supercomputer, and the IBM BlueGene/Q is the world's most energy efficient peta-scale machine.

Accelerator users also benefit from access to petabyte-scale data storage, archive and backup facilities, as well as access to specialised GPU compute nodes.

INDY hardware is supplied by Viglen and utilises the AMD Opteron processor. The system consists of 24 back-end nodes and two front-end, login nodes. It is interconnected using a high-performance, very low latency ethernet switch from Gnodal. There are four Opteron, 16 core processors per node giving 64 cores per node and 1536 cores in total. As standard, each back-end node has 256Gbyte of shared RAM, with two large memory back-end nodes configured with 512 Gbyte RAM to support applications with a requirement for a larger shared memory resource. In addition, the system has support for future installation of up to two GPGPUs cards (NVIDIA or AMD) per node.

To support the widest possible range of application codes INDY utilises IBM's industry leading Platform HPC cluster management software providing job level dynamic provisioning of compute nodes into either Windows or Linux depending on a user's specific O/S requirement. Engineering users, for example, using Windows based CAD tools will be able to seamlessly integrate INDY HPC into their current PLM processes, whilst scientific users running Linux based codes on modest internal clusters will be able to facilitate efficiency improvements by executing the same codes directly on INDY.

Potential users of **Accelerator**'s INDY HPC service can benefit from a number of improved use-case scenarios:

- As a cost effective alternative to buying and operating in-house HPC capability.

- To augment an organisation's own HPC cluster during times of peak demand.

- To provide access to extra compute capability as a means of reducing modelling and simulation times, accelerating early stage product development lifecycles.

- As a contingency over internal infrastructure failure.

Accelerator platforms are housed in EPCC's purpose built, world-class, Advanced Computing Facility, providing the highest levels of physical and logical security.

George Graham
g.graham@epcc.ed.ac.uk

Accelerator is one of Europe's most powerful on-demand HPC services.

Offering easy access, a choice of Windows or Linux platform, and the ability to orchestrate resources as required, Accelerator gives direct, desktop access to unprecedented levels of HPC performance.

Accelerator

To access any of **Accelerator's** services, including INDY, contact George Graham:

g.graham@epcc.ed.ac.uk

+44(0) 131 651 3460.

Business innovation through HPC

A diverse and broad range of companies has engaged with the Supercomputing Scotland programme since its inception at the beginning of 2012 and the first HPC Adopter project is due to start in late November.

Ronnie Galloway
r.galloway@epcc.ed.ac.uk

A wide range of Scottish companies is looking to HPC to help meet their business objectives.

www.supercomputingscotland.org

Supercomputing Scotland is a business improvement initiative from EPCC at The University of Edinburgh and Scottish Enterprise which provides Scottish companies with the knowledge to help them decide if using high performance computing (HPC) makes sense for them in reaching their business objectives and, if so, supports the company to proceed with a project part-funded by Scottish Enterprise and delivered by EPCC.

25 companies are currently engaged with EPCC across the three 'levels' of the Supercomputing Scotland programme, with 6 on the point of commencing a project within the next month or so. The rigorous approach of the Level 1-3 process means that a detailed understanding of the companies' needs is achieved and matched to EPCC solutions by the time a project is put forward to Scottish Enterprise for funding.

The diversity of the companies involved to date is evident when considering and comparing the prospective projects:
• Doosan Babcock: turbine modelling
• Xi Engineering: acoustic modelling
• IDEAS Engineering: plant integrity
• WeeWorld: social network/gaming
• Live Visualisation: 3-D rendering of point cloud data
• IES: thermal imaging.

The potential project with Live Visualisation is an interesting development of 3-D technology taken from the company owner's game-designing background into the realm of heritage tourism whereby virtual visits to museums and other places of historical and cultural interest are created.

A preliminary project with the new V&A Gallery in Dundee provides the background to this development. EPCC is currently discussing how to help Live Visualisation to develop and improve their 3-D imaging capabilities to create finer detail of the artifacts and interiors of historical and cultural buildings.

Scaling Soft Matter Physics

A thousand GPUs and beyond

Simulations based around fluid dynamics offer a powerful way to study, predict and ultimately improve the behaviour of "soft matter": everyday materials such as paints, engine lubricants, foodstuffs and cosmetic items; hi-tech items such as liquid crystal displays (LCDs) and even biological fluids within the body.

Over the last 12 years, a collaboration between the University of Edinburgh's Soft Condensed Matter Physics group and EPCC has developed simulations of soft matter systems using the lattice Boltzmann method and the parallel computing code "Ludwig" to accurately capture the physics of systems such as mixtures, suspensions and liquid crystals. Understanding and controlling the phase separation of liquid mixtures, for example, can improve the shelf-life of foodstuffs – one of the many practical applications of the research.

Ludwig performs excellently on "traditional" CPU-based supercomputers, and is able to take advantage of many thousands of compute cores in parallel. But recent work at EPCC has taken this a step further by engineering the code for exploitation of large-scale GPU-accelerated architectures, such as the cutting-edge Cray XK6™ supercomputer. Such machines offer an impressively high ratio of performance to power consumption and can be thought of as a possible template for future exascale systems (for which power consumption will be a key issue).

The main challenge in this work is to maintain good scalability whilst exchanging data between many GPUs. To ensure optimal performance scaling, use was made of NVIDIA® CUDA™ stream technology combined with MPI to minimise communication overheads through overlapping. The new code is observed to scale excellently such that it can exploit a thousand GPUs in parallel on the Cray XK6 – with a view to even larger simulations as bigger machines become available.

"A crucial factor in obtaining accurate results is the size of the physical system that we can simulate," says lead Ludwig author Kevin Stratford (EPCC). "To model complex problems in large systems, we need efficient, scalable application performance over large numbers of nodes."

Further work is underway to enable advanced functionality for the GPU architecture such as the ability to include, in the simulation, particle suspensions together with liquid crystals. This is of particular interest for new generations of LCD display technology.

Alan Gray
a.gray@epcc.ed.ac.uk

Alistair Hart
ahart@cray.com

"Large-scale computer simulation has become a central tool for research in materials physics. Thanks to the Cray XK6 system and similar innovations, the rate of increase in computational capability is breathtaking."

Prof. Michael Cates, Head of the Soft Matter group

SPRINT: more runners, fewer hurdles

The SPRINT project focuses on the parallelisation of R functions used by biostatisticians and other data analysts.

Kevin Robertson
Kevin.Robertson@ed.ac.uk

Find out more

January 2012 marked the first full release of SPRINT v1.0, providing 7 parallelised functions of generic utility in the analysis of large data matrices.

The SPRINT team was recently awarded a further 2 years of BBSRC funding to promote the accessibility and usability of the framework.

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

1. J. Hill et al., Sprint: A New Parallel Framework for R. Bmc Bioinformatics 9, (Dec 29, 2008).

2. L. Mitchell et al., paper presented at the Proceedings of the second international workshop on Emerging computational methods for the life sciences, San Jose, California, USA, 2011.

3. S. Petrou et al., Optimization of a parallel permutation testing function for the SPRINT R package. Concurr Comp-Pract E 23, 2258 (Dec 10, 2011).

R is a statistical programming language used for data manipulation, visualisation and statistical analyses by diverse organisations such as Thomas Cook and the US Food & Drug Administration. Since 2000, R has become indispensible for the analysis of post-genomic data.

High-throughput biomedical research platforms such as next-generation sequencers are producing large and complex datasets, testing the limitations of serial computational approaches. To overcome these limitations, SPRINT was established to provide a set of easy to use, 'drop-in' parallelised functions allowing the biostatistician straightforward access to high-performance computing.

Established in 2008, SPRINT is a collaborative project between EPCC and the University of Edinburgh's Division of Pathway Medicine. Since inception, SPRINT's functionality has been driven by the requirements of the biostatistics community. This work has produced a full release of the software (available on CRAN or R-forge, it has been downloaded around 700 times a month) and resulted in several publications describing the framework and its specific R functionality[1,2,3].

In the next two years, SPRINT will focus on...

1. **Access**
   By facilitating local installs and providing user training alongside installations on HECToR and the Amazon EC2, the team aims to remove technical barriers that prevent widespread use of SPRINT functionality.

2. **Availability**
   Currently running on Linux, SPRINT developers will migrate the framework to new platforms such as Apple OS X.
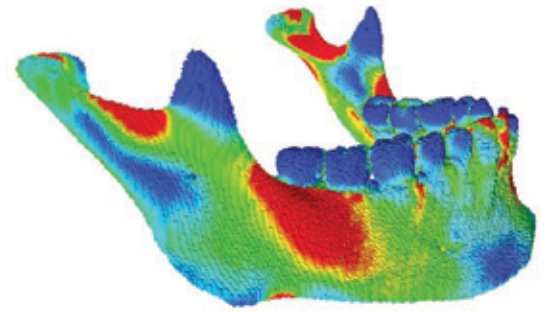
3. **Application**
   The SPRINT team will re-engage with the R community, establish and prioritise the development of new parallelised functions applicable to, for example, explorative machine learning (eg Hamming Distance). As part of this process, the team will engage with new users, beyond their core of biostatisticians, and assess fresh opportunities for application of the software.

Modelling bone structure on HECToR

VOX-FE is a complete, voxel-based finite element bone-modelling suite developed by Prof. Michael Fagan's Medical & Biological Engineering group at the University of Hull[1].

Via the HECToR dCSE scheme[2], the group at Hull has accessed expertise at EPCC to improve the I/O routines in their code to prepare for future research on HECToR.

The VOX-FE suite comprises two parts: a GUI for manipulating bone structures and visualizing the results of applying strain forces to them; and a Finite Element solver, "PARA-BMU", which performs the heavy computation required. An example application would be computing the maximum principal strain in a human mandible during biting.

The dCSE project aimed to improve the I/O performance of PARA-BMU. Examination of the code showed that plain ASCII files were being used for both input and output data. Reading and writing these was done in serial by one process whereas the solver routines operated in parallel. This solution did not scale well to a large number of cores with the result that for systems with a large number of finite elements, the runtime was dominated by I/O and MPI data exchange rather than by computational routines. It also meant that file sizes grew extremely large as the problem size increased, which increased transfer times to and from HECToR.

Prof. Fagan's group intend to routinely study bone models with resolutions of over 100 million elements in the near future.

We replaced both the input and output routines with parallel versions based upon the netCDF library with parallel HDF5 support[3]. netCDF is a convenient interface library which allows easy manipulation of HDF5 files in both serial and parallel. In this case, we chose to make use of the parallel interface for optimal performance and scalability. However, this is an implementation decision and does not affect the underlying file, which can then be read with the serial interface by other applications (for example, the GUI).

In addition, netCDF files can be compressed using the freely available NCO utilities and, importantly, read in parallel directly from the compressed format. This had the triple advantage of reducing file sizes by nearly 200 times in one case, reducing I/O time by a factor of 6 and converting files to a portable, self-describing format which can easily store additional meta-data about the dataset.

In addition, we produced a set of netCDF to ASCII converters to allow output files to be converted back to the original ASCII format and to allow existing input files to be converted to netCDF. These pre- and post-processes can be performed off-line when interaction with the rest of the VOX-FE suite is required.

As a whole, this represents a much more scalable solution than was possible before, especially for increasingly large simulations.

Nick Johnson
nick.johnson@ed.ac.uk

Iain Bethune
ibethune@epcc.ed.ac.uk

Result

It is now feasible to consider HECToR as a platform for running PARA-BMU rather than small-scale local clusters, and paves the way for bone modelling at unprecedented scale and accuracy.

1. http://www2.hull.ac.uk/science/medical__biological_eng/research/vox-fe.aspx

2. http://www.hector.ac.uk/cse/distributedcse/

3. http://www.unidata.ucar.edu/software/netcdf/

International collaboration on research data infrastructure

Rob Baxter
r.baxter@epcc.ed.ac.uk

The tides of digital data are rising exponentially quickly, and major efforts are underway worldwide to find ways to stem them – or at least to save us all from drowning.

www.icordi.eu

Persistent identifiers

Imagine a huge library, shelf upon shelf, row upon row. Without a catalogue that notes the location of every new book, it won't take long before books become lost in the stacks.

Persistent identifiers (PIDs) play the same locating role for digital data objects – files, images, spreadsheets, whatever. The same PID, when shown to the "digital librarian", will always retrieve the same digital object, now and forever.

Necessity is the mother of invention, and a number of scientific disciplines have been wrestling with ways to manage their own data better for years. A good deal of effort in Europe in recent times has gone on managing the deluge of bits from the Large Hadron Collider (LHC) at CERN. The International Virtual Observatory Alliance (IVOA) has been coping with larger and larger telescope images on a global scale for over 10 years, and during the same period the European Molecular Biology Laboratory – European Bioinformatics Institute (EMBL-EBI) and others have been wrestling with the explosion of data in post-genomic life sciences.

These data-rich sciences have, out of necessity, developed systems and procedures for managing their own collections but as a consequence the risk of different kinds of data becoming locked in silos with no easy way of connecting them increases year on year. Now connecting LHC data with EMBL data makes little sense, of course – but LHC data with IVOA data? That may be valuable for astrophysicists. And what about EMBL data with biodiversity and habitat data, perhaps including species migration data and human urban dynamics; could such combinations help us understand how H5N1 flu will mutate and spread? Perhaps they could.

This possibility of being able to ask new research questions across discipline boundaries is compelling. The EUDAT project (see EPCC News 71) is an initiative taking steps towards a collaborative data infrastructure to enable just such research in Europe – but research today is global. Connecting European efforts to similar activities in the US, China, Australia and internationally is increasingly important.

ICORDI ('eye-cordy') is an EU-funded collaborative network designed to link European data infrastructure activities with those in the US and beyond. Aligning with parallel funding from the National Science Foundation in the US and the Australian National Data Service, iCORDI and its sister projects across the world aim to drive the agenda of data interoperability and integration through a new policy and standards forum called the Research Data Alliance (RDA). Or maybe "anti-standards" is a better term; rather than invent yet more standards that serve little purpose, RDA seeks to harmonise, to reduce the number of existing standards in key areas critical to the agenda of global data interoperability. Instead of a dozen systems for persistent data identifiers, for instance, let's try to agree on one.

ICORDI will engage strongly with the RDA, but it also runs a programme of "learning by example", a community-led "show-and-tell" activity which will create a series of prototypes of interoperable data systems across the world. Initially four important research areas – astronomy, Earth sciences, chemical safety and publication-data linking – have been chosen to pursue these exemplars and to showcase them in international workshops, warts and all, to help inform the RDA agenda. EPCC leads this community prototyping programme, and over next two years we will be assisting researchers around the world in understanding and overcoming the challenges of creating truly interoperable data systems.

BonFIRE: it's spreading

BonFIRE has grown rapidly since its inception. Today BonFIRE sites have a capacity of over 450 dedicated cores, with 1TB of RAM and 47 TB of disk and almost twice as much capacity is available on request. The EPCC testbed is a prime example of this growth. Funding from the University of Edinburgh's School of Physics has allowed us to upgrade the original infrastructure to 96 cores, with 256GB of RAM and 24TB of storage; an additional 80, low-energy cores are currently being added to the system.

The BonFIRE sites' heterogeneous architectures are a major strength, as they allow experimentation and what-if analyses on various target systems.

For example, IT Innovation at the University of Southampton has executed numerous experiments on BonFIRE's heterogeneous infrastructure and validated its hypothesis that Dwarf benchmarks [1] can be used as a uniform and informative way of characterising compute resources and also to predict the performance of several multi-media and scientific applications.

But how does one make use of this heterogeneity? BonFIRE gives users control with several BonFIRE sites now allowing users to select the specific physical host on which to instantiate their resources. This allows users to make the most of BonFIRE's heterogeneous architecture and, combined with knowledge of the exact specification of each machine on BonFIRE, interpret the results. An example of this level of control on BonFIRE is aptly demonstrated at the Inria BonFIRE site: a large pool of resources is available on request, with exclusive access to the physical host. This isolation increases the repeatability of experiments, giving unprecedented control compared to contentious, multi-tenant, standard clouds.

A major characteristic of BonFIRE is its attention to observability. We recognise the importance of monitoring and collect over 100 metrics for each of the compute resources instantiated on BonFIRE. Motivated by the BonFIRE experimenters, we now provide information that public cloud providers would tend to hide from their users: selected BonFIRE sites provide monitoring of the physical hosts; detailed, time-stamped logs of the user requests as they traverse the BonFIRE stack; and detailed, time-stamped logs of the compute-resource deployment process.

Recently, BonFIRE experimenter CESGA verified its software which calculates radiotherapy doses. One

Kostas Kavoussanakis
kavousan@epcc.ed.ac.uk

BonFIRE is an EC-funded project to build a multi-cloud testbed for advanced experimentation on systems, services and applications.

www.bonfire-project.eu

Open access for experimenters

BonFIRE is entering its last year. A major goal for the very near future is Open Access for experimenters. Through this scheme, researchers will be invited to submit succinct proposals for experimentation, using parts of BonFIRE free of charge for short periods of time.

To register interest contact: bonfire@bonfire-project.eu

Two years in, BonFIRE is fulfilling its promise to be a unique infrastructure for testing and validating Cloud and networked applications.

Ease of use

Early on we provided an Application Programming Interface (based on the Open Cloud Computing Interface specification) to give an entry point for our user community. The graphical user interface of the BonFIRE Portal allows access at a higher level and aims to facilitate adoption.

EPCC is in charge of the BonFIRE Experiment Manager. This allows declarative experiment description using the Open Virtualisation Format standard or the BonFIRE experiment descriptor, which is based on the JavaScript Object Notation.

The Experiment Manager encapsulates the essence of experimentation on BonFIRE: it allows users to easily vary and/or repeat the deployment of their experimental scenarios.

1. Asanovic, K., Bodik, R., Demmel, J., Keaveny, T., Keutzer, K., Kubiatowicz, J., Morgan, N., Patterson, D., Sen, K., Wawrzynek, J., Wessel, D., Yelick, K.: A view of the parallel computing landscape. Communications of the ACM 52(10) (Oct 2009) 5667

of their conclusions is that, given sufficient monitoring data, its application can adapt to the loss of a computing facility and deliver the results on time by increasing capacity on a different site.

BonFIRE has always aimed at providing more than public clouds, and its advanced features are mostly centred on extending experimentation towards the network.

BonFIRE's Virtual Wall facility allows network experimentation based on emulated, but highly configurable, networks. The TurboCloud BonFIRE experiment used this facility to validate the video streaming solution of Irish SMEs RedZinc and Cloudium Systems. Through experimentation only possible on this facility, these SMEs now know the exact capabilities of their solution and can market it accordingly.

BonFIRE attracted further funding from the European Commission to interconnect with the GÉANT AutoBAHN bandwidth-on-demand service. EPCC, through the UK's JANET National Research and Education Network, and PSNC will offer this functionality. Access to AutoBAHN through BonFIRE is imminent and will allow our users to experiment with controlled bandwidth. Similarly, federation with the FEDERICA facility will allow

BonFIRE experimenters the opportunity to extend their experimentation to include aspects of routing network traffic.

So, BonFIRE is spreading: more capacity, more features, more users, but that's not all. Three new projects all involving EPCC – Fed4FIRE, BonFIRE and CityFlow – started in October at the Future Internet Research and Experimentation (FIRE) EC unit.

Fed4FIRE will federate FIRE facilities, including BonFIRE, and EPCC is leading the work towards integrating the services and applications communities within this federation. ECO2Clouds will utilise the BonFIRE facility to introduce environmental aspects to computation on the Cloud; here EPCC is in charge of the architecture which will combine eco-metrics with application performance and support an eco-aware, multi-cloud scheduler. CityFlow (coordinated by EPCC) builds on the results of BonFIRE's TurboCloud experiment to explore the behaviour of traffic-sensitive applications like Web TV and on-demand video in a simulated city of 1 million broadband connections.

You can read more about these new projects overleaf.

Energy-efficient Cloud computing

How much energy is required to deploy a Cloud infrastructure and how can the computational load on the Cloud be allocated in order to maximise both energy and performance efficiency? These are questions that are becoming increasingly important in a world where on-demand computing is now part of everyday life. For instance, many on-line applications use Clouds as a backend for their computation.

ECO_2Clouds (Experimental Awareness of CO_2 in Federated Cloud Sourcing) will aim to address the subject of energy efficient Cloud computing over the next two years. The project will define a set of metrics encompassing energy efficiency and carbon emissions. We will then directly measure these metrics on private Cloud test beds provided by the BonFIRE project, establishing the impact of factors such as workload distribution and virtualisation on the efficiency of the Cloud-based applications, both in terms of energy consumption and runtime performance. Based on this knowledge, the project will then develop a scheduler that will place workloads on the Cloud with the aim to achieve optimal performance within agreed service level parameters, while keeping the energy usage and environmental impact as low as possible.

Michèle Weiland
m.weiland@epcc.ed.ac.uk

EPCC leads the work on requirements capture and the overall design of the architecture.

FIRE federation

Future Internet Research and Experimentation (FIRE) testbeds aim to foster larger-scale, cheaper experimentation than is possible through project-specific testbeds. Experimentation is varied, with users from both industry and academia. The testbeds cover wired and wireless networking, systems and services as well as Smart Cities (cities with sensors).

The BonFIRE project has shown the benefits of federating heterogeneous facilities. Federating horizontally, ie similar testbeds, achieves scale. Federating vertically, ie linking fundamentally different facilities, enables research that encompasses cross-cutting concerns. Both ways enable new, creative research.

Fed4FIRE takes this further, by institutionalising federation on FIRE according to user demand. Fed4FIRE will develop an open, extensible framework that fosters a dynamic ecosystem of facility providers and experimenters to accelerate Future Internet research.

Fed4FIRE is investing in experimentation tools for federation, monitoring and trustworthiness. Such tools will foster efficient experimentation by facilitating the use of heterogeneous platforms while promoting the value of individual testbeds and the benefits of federation. Monitoring data will fuel research and increase the impact of FIRE experimentation. Finally, the trustworthiness work on BonFIRE goes beyond the complex issues of federated identity management and access control, to address quality of service and experience for the experimenters.

Two parallel activities will motivate the development of the tools through community requirements for federation, and deploy them on the respective infrastructure and services facilities. A professional user support infrastructure will ensure that our experimenters make the most of Fed4FIRE's ground-breaking capabilities.

Through Fed4FIRE, EPCC's BonFIRE site will grow, supporting even more, varied, impactful research on the Future Internet.

Kostas Kavoussanakis
kavousan@epcc.ed.ac.uk

Drawing on expertise gained on Grid and Cloud projects, EPCC is leading the workpackage on services and applications.

Fed4FIRE's first Open Call for experimentation and new facilities will be in May 2013.

What does the future Internet look like?

Rob Baxter
rmb@epcc.ed.ac.uk

The topic of "net neutrality" has sparked polarized debate in recent years. Should traffic on the Internet be "differentiated" to give some types of traffic higher priority?

1. FCC's Net Neutrality Order (10-201) December 23, 2010: http://www.fcc.gov/ Daily_Releases/Daily_Business/2010/db1223/ FCC-10-201A1.pdf

2. CityFlow is funded by the European Union under grant agreement 317576.

In many ways the argument is one of both wanting our cake and wishing to eat it. Few of us like the idea of network operators discriminating against certain kinds of Internet traffic – but none of us want any glitches in our on-demand video services. Currently the public Internet is undifferentiated whereas many private networks and VPNs use differentiation to give quality of service support to critical applications (eg voice, video, priority data traffic). The US recently, and with much controversy, gave the green light to reasonable network management practices in the context of proposed net neutrality principles of transparency (traffic management should be clear and above board), non-blocking (lawful content will not be blocked) and no unreasonable discrimination.

In this context, new standards and products in the field of reasonable Internet traffic management are emerging with the goal of achieving the effect of guaranteeing a bandwidth end-to-end for a user without interference from unwanted traffic, while remaining true to the net neutrality principles. But how will these new approaches work? How can you run large-scale experiments on the way this future Internet might behave without breaking the current one?

The European-funded CityFlow project is one answer to this question. By building on the results of TurboCloud (see BonFIRE article on previous page), CityFlow will explore the behaviour of traffic-sensitive applications like Web TV and on-demand video in a simulated city of 1 million broadband connections. A core part of CityFlow is the VPS Controller, a novel traffic management technology from partners RedZinc Services ltd of Ireland. The project will deploy the VPS Controller on the large-scale network simulation platform known as the Virtual Wall hosted at the Interdisciplinary Institute for Broadband Technology (IBBT) in Belgium, and subject it to a number of data-heavy scaling tests developed by EPCC. Portuguese partners ONE Source lda round off the team with expert knowledge of the underlying network protocols. The whole project is managed by EPCC and kicked off on the 1st of October this year, running for 18 months.

So if you want to know what the future Internet looks like, ask us in 18 months' time!

CRESTA

EPCC's technical challenges

CRESTA, EPCC's flagship EC-funded exascale research project, has reached the end of its first year. Here we describe two of the technical areas that EPCC is involved with.

Simulation and modelling technologies

Simulation and modelling are important tools in the development of exascale systems. This is because there are few other mechanisms for evaluating designs for exascale hardware and software. EPCC has recently been investigating CRESTA's modelling requirements and assessing the most appropriate technologies.

In the early stages of the design process, models need to be quite simple and abstract. This allows us to develop and evolve our designs quickly and efficiently. If we attempt to use overly complex models in these early design stages we will waste time and resources performing overly detailed simulations of design choices that will be abandoned before the final system is built.

Current thinking about exascale hardware design is that it will be highly constrained by system power consumption. To keep power consumption within acceptable levels, exascale systems will need to utilise very high degrees of parallelism and the performance of their communication systems may also have to be limited. This implies that we should use software models that explicitly capture the available parallelism and the communication requirements of an algorithm. One way of capturing this information is to consider modelling the parallelizable sections of the algorithm as directed acyclic graphs.

Application behaviour can be simulated at a high level by simulating its communication pattern. This allows the application's behaviour to be extrapolated to different (possibly theoretical) hardware platforms, so we can explore the behaviour of applications on exascale hardware well before such hardware becomes available. Various approaches to application simulation exist. Particularly interesting is the use of simple skeleton applications to drive the simulation. These are lightweight, simple codes intended to capture the essential behaviour of larger, much more complex applications. They provide a mechanism for exploring the behaviour of new designs without the cost of first developing the design into a fully-functional application. As this approach aims to capture the communication pattern rather than the details of the computational sections, it provides a mechanism to develop a directed acyclic graph model into a form that can be simulated.

Our research has identified a number of suitable simulation platforms. These can give useful insights into the limitations the network imposes on application performance. Many of these

Stephen Booth
spb@epcc.ed.ac.uk

David Henty
dsh@epcc.ed.ac.uk

Lorna Smith
l.smith@epcc.ed.ac.uk

CRESTA focuses on six applications with exascale potential, using them as co-design vehicles to develop an integrated suite of technologies to support the execution of applications at the exascale.

www.cresta-project.eu

EPCC coordinates this large, FP7 project and provides technical expertise on four of the five technical work packages.

platforms are explicitly targeting the development of exascale systems. Though these simulation tools are useful, they are also fairly complex and can be difficult to use.

Compiler technologies and autotuning

One of CRESTA's early deliverables involved a survey of compiler technologies relevant to exascale. Two particular issues were identified: the growing requirement for compilers to support CPU accelerators; and the possible advantages of auto-tuning to produce better performing code on today's increasingly complex and heterogeneous processors. Motivated by this, EPCC has been studying key kernels from one of the CRESTA co-design applications, Nek5000. We have investigated to what extent user-level source modifications affect performance at different levels of optimisation across a range of compilers. We have also studied new GPU-enabled version of the kernels, written using the new OpenACC standard for accelerator directives, to explore current compiler capabilities for heterogeneous architectures. Lastly, we attempted to optimise the performance of OpenACC using auto-tuning technology.

We deliberately took a very naïve approach, using simple kernels to investigate the accelerator

capabilities of today's compilers and to see how much performance improvement could be gained from compiler auto-tuning. The results were very positive overall, with significant speedups achieved on an NVIDA GPU compared to the single CPU versions of the Nek5000 benchmarks. However, the performance achieved was still well below the peak of the GPU, and probably even inferior to a simple multicore CPU version.

As OpenACC is such a new technology the auto-tuning done here is quite basic, looking only at the space of high-level loop scheduling options with the actual source code remaining fixed. Many more GPU optimisations can be investigated for codes written at a lower level using CUDA. For example, in [1] the full CHILL framework is extended to include CUDA and the auto-tuned performance of matrix-matrix code approaches that of the CUBLAS library. A similar study (although not performed using CHILL) was carried out in [2] for 3D Fast Fourier Transforms, where the auto-tuned code out-performed the vendor-supplied CUFFT library. However, because OpenACC is much more portable and easier to use than CUDA, we believe that auto-tuning the higher-level OpenACC parameters will become increasingly important in the future as HPC application developers move away from CUDA.

1. G. Rudy, et al., "A programming language interface to describe transformations and code generation" in Proceedings of the 23rd international conference on Languages and compilers for parallel computing (LCPC'10).

2. A. Nukada and S. Matsuoka, "Auto-tuning 3-D FFT library for CUDA GPUs" in Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis (SC '09).

Image © Chris Cannam

Mike Jackson presenting at a bootcamp at Newcastle University earlier this year.

Software Carpentry goes from strength to strength

As part of its goal to enable better research through better software, the Software Sustainability Institute will help run three Software Carpentry boot camps for UK researchers this autumn.

EPCC is the lead institution of the SSI.

The Software Carpentry project teaches scientists how to quickly build the high-quality software they need, and so maximise the impact of their research. Short intensive workshops (boot camps) are followed by 4-8 weeks of self-paced online instruction.

Boot camps do not only teach specific products but also practices to help scientists rapidly develop high-quality software. Instead of just stating that "these are good because software developers say they are", each practice is justified with reference to individual and group psychology and empirical studies into how software development works.

The Software Sustainability Institute (SSI) attended the first UK boot camp, led by Greg Wilson, at University College London in April 2012. A fortnight later, in conjunction with the Digital Institute at Newcastle University and

SoundSoftware, SSI delivered the first boot camp to be run entirely by UK tutors, independent of Greg Wilson's team. And now it's helping run three boot camps this autumn:

1. For institutions in the north-east of England, organised by The University of Newcastle.

2. For bioinformaticians based at, and organised by, The University of Oxford.

3. A boot camp open to all organised by EPCC's PRACE Advanced Training Centre (see page 3).

In the new year SSI will help run two boot camps in Germany, to seed Software Carpentry in Europe. It is also actively working towards a sustainable future for Software Carpentry in the UK. This is part of a strategy to ensure researchers have access to training in software development and we look forward to reporting on progress.

Mike Jackson
michaelj@epcc.ed.ac.uk

Links

SSI website
http://software.ac.uk/

SSI training
http://software.ac.uk/what-do-we-do/training

Upcoming boot camps
http://software.ac.uk/boot-camps

Software Carpentry
http://software-carpentry.org/

PRACE
http://www.prace-project.eu/

SSI Agents

A success for cross-disciplinary intelligence

Last year, faced with the problem of gathering intelligence from across research disciplines, the Software Sustainability Institute created the Agents network. The first stage of the network will end in December. What did we get for our investment?

Simon Hettrick
s.hettrick@software.ac.uk

We began the search for our Agents in July 2011. These would be researchers who had an excellent understanding of their field and experience of software in research.

The response was fantastic, with over 130 people applying for the 10 positions we had on offer. This surplus made it relatively easy to select 10 Agents who were outstanding researchers and communicators, and who represented a range of disciplines from glaciology to sports science. In return for their expertise and advice, each Agent received a £3000 travel budget.

Over the last year, we have received hundreds of communications from the Agents covering everything from news about the latest research and interesting new software, contact details for influential people and even requests for collaboration. This information keeps us at the cutting edge of the research software community, and provides plenty of content for our news service and blog, which attracts more researchers to the Institute.

The Agents went way beyond their original intelligence-gathering remit. They were an invaluable sounding board for our ideas, helped out at workshops and conferences, and they acted as a seed contact within their communities – a spark to start off wider dissemination.

Many of the Agents were already active in the research community, which meant that we received secondary benefits like invitations to speak at conferences, and blog post and articles written by the Agents in their own time.

The Agents network has been a fantastic success by any metric: intelligence gathered, participation in events, feedback on strategy or content generated for the research community. With cross-disciplinary intelligence gathering so important to research organisations, the Agents network showed that a small investment in something that researchers need – travel money – can pay dividends.

The Agents network has been expanded this year and its name has changed to the Fellowship programme.

Fellowship programme
http://bit.ly/ssi-fellows

Agents network
http://bit.ly/SSIAgents

Profiles

Two of our colleagues talk about their work here at EPCC

I obtained my degree in Computer Science in China before coming to Edinburgh in 2005 to study the MSc in High Performance Computing at EPCC. After graduation I joined EPCC as an Applications Consultant.

I really love the warm environment here, people are very friendly and open to new ideas. It's a great pleasure to work with my colleagues. HPC is used in many areas of academia and industry, so we have plenty of opportunities to work on interesting projects.

My job focuses on applications enabling and performance optimisation and I like to work on the most challenging tasks in the morning. Nothing is better than solving problems that have been stacked for a while.

EPCC is engaged in many research projects, and we work with both internal colleagues and external partners. The project I'm mainly working on – PRACE (Partnership for Advanced Computing in Europe) – requires me to travel occasionally. As of October 2012, there are 25 member countries in PRACE and many supercomputing sites are involved in the implementation phase projects. We meet our external colleagues about every 6 months, which allows us to make faster progress on collaborative tasks, review the work done and discuss the next stages.

Sometimes I need to prepare and gain new knowledge before starting a task, which can be challenging, especially when there is a pressing deadline. But it's also a great chance to learn something new and gain new experiences.

Training is an important aspect of EPCC's activities. I enjoy teaching on the MSc and supervising projects – and we can always learn a lot from our students too.

Xu Guo
Applications Consultant
xguo@epcc.ed.ac.uk
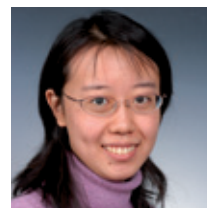
It's exciting that we can face new challenges every day and apply the very latest technologies in our daily life.

I work in the User Support and Liaison (USL) team and the majority of my time is spent helping people exploit the various computing resources hosted by EPCC by optimising their parallel codes and troubleshooting their problems.

One of my big interests is teaching and there is ample scope to teach at EPCC: I run an undergraduate course in scientific computing that has around 100 students annually and which I work closely with the School Teaching Office and postgraduate teaching assistants to deliver. This job also gives me the opportunity to collaborate on funding proposals with researchers; develop open source code; and do some web development on the side.

My background is scientific: I have a degree, M.Phil. (both from Bangor University) and Ph.D. (from the University of Liverpool) in chemistry. I have been a scientific programmer since my M.Phil. and am mostly self-taught, originally in Fortran 90 but have since branched out into a variety of languages and also into parallel programming.

I enjoy the variety of the work here, the different people I work with and the ability to direct my job into areas that interest me. My job also gives me the chance to learn new things on a daily basis.

I frequently meet with researchers, which gives me the opportunity to visit different institutions around the UK. EPCC has a large involvement in pan-European projects so there are often meetings on the Continent to attend too.

I love the fact that my role at EPCC is very outward facing. It is extremely rewarding to spend time with the researchers who use our computing facilities, working to solve their problems and helping them to get the most out of the resources.

Andy Turner
Project Manager
a.turner@epcc.ed.ac.uk

One of the great things about my job is that there is no typical day.
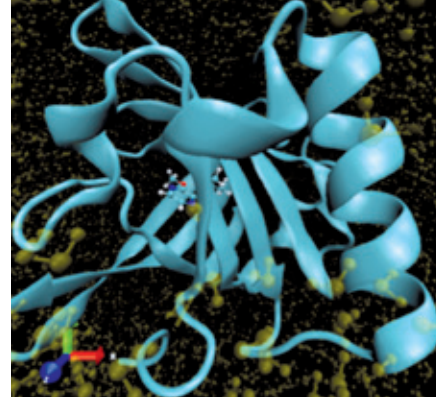
Left: participants enjoying the EPCC demo at Melrose's Bang Goes The Borders event.

Right: Mouse Major Urinary Protein shown as a ribbon structure in a sea of water molecules. Embedded inside the protein is the pheromone molecule.

Supercomputing for everyone

Mario Antonioletti
m.antonioletti@epcc.ed.ac.uk

Iain Bethune
i.bethune@epcc.ed.ac.uk

We demonstrated the power and utility of supercomputers at the British Science Festival in Aberdeen and Bang Goes The Borders in Melrose.

At both events the audience showed a high level of interest in finding out what supercomputing is all about.

Our goal was to encourage the general public to interact directly with a supercomputer. We already had the supercomputer – but we needed an application.

Visual Molecular Dynamics[1] (VMD) is an application that allows systems to be visualized while being simulated by an underlying code. We chose NAMD[2] to run the simulation as it is a supported application on HECToR and is well integrated with VMD. Not only does VMD let you visualize a system as it runs – including zooming, rotating and panning - but it also allows external forces to be applied to an atom or molecule. These are fed back into the simulation in real time.

Now we required an interesting simulation to tell a compelling story.

Professor Charles Laughton (University of Nottingham) has studied the thermodynamics of ligand binding using the mouse major urinary protein[3]. The simulation itself is simply a protein embedded in a water molecule. Inside the protein is a pheromone molecule – the protein acts as a slow release for any of the more volatile pheromone particles that may be trapped inside.

Using NAMD to model the system and VMD to visualise it, we created a demo where the audience could apply external forces to try to pull the pheromone out of the protein.

As well as the simulation, we had a Cray XT4 blade, courtesy of Cray. This was a big crowd pleaser!

We also wanted to convey how a task could be performed faster by following a parallel algorithm. We printed cards numbered 1–60, shuffled a subset of them and asked a willing volunteer to sort them. Invariably this took too long, so we stopped this early having demonstrated the point. We then got several members of the audience to act as computer cores. They were each given a subset of cards which they sorted by swapping highest and lowest numbers, until all the highest cards were at one end of the line and all the lowest were at the other.

Having played the part of a supercomputer, the audience moved on to trying to pull the pheromone out of the mouse protein, with the demo running on 256 cores remotely on HECToR while being displayed locally using VMD at Aberdeen. So not only did our audience get to play the part of a supercomputer, they also got to play with the real thing!

We also took our demo to Melrose's "Bang Goes The Borders" festival, which offers children a chance to experience science from laser physics to zombies – and now supercomputers too.

1. http://www.ks.uiuc.edu/Research/vmd

2. http://www.ks.uiuc.edu/Research/namd

3. J. Roy and C. Laughton, Biophysical Journal 99, 218-226, Jul 2010

Call for participation: Exascale Applications and Software Conference

Edinburgh, 9–11 April 2013

This conference will bring together all the groups with a stake in solving the software challenges of the exascale: from application developers, through numerical library experts, programming model developers and integrators, to tools designers.

The scale of today's leading HPC systems has put a strain on many simulation codes – both scientific and commercial.

However, many of the scientific challenges behind these codes are driving the need for the next generation of exascale HPC systems. For simulation codes that are already struggling to scale up to petaflop levels, major investment is required to enable these codes to run at the exascale.

Application optimisation and algorithmic modifications only represent part of this challenge. Systems of the scale envisaged present enormous challenges in terms of reliability, programmability, power consumption and usability. Programming models, libraries, languages, compilers and tools all need adaption and improvement. Applications must interact with many of these software aspects to be able to exploit exascale systems efficiently.

Authors are invited to submit novel research and experience in all areas associated with developing applications for exascale and the associated tools, software programming models and libraries.

Abstracts are requested on relevant topics including:

· Enabling and optimising applications for exascale in any scientific area

· Developing and enhancing numerical algorithms for exascale systems

· Aiding the exploitation of massively parallel systems through tools, eg performance analysis, debugging, development environments

· Programming models and libraries for exascale

· Evaluating best practice in HPC concerning large-scale facilities and application execution.

Abstracts should be no longer than 500 words and must be submitted by 10 December 2012.

Submission instructions: www.easc2013.org.uk

The conference is being organised by EPCC and will be held at the University of Edinburgh.

Irina Nazarova
i.nazarova@epcc.ed.ac.uk

Confirmed speakers

- Satoshi Matsuoka
  Tokyo Institute of Technology

- Vladimir Voevodin
  Research Computing Center, Moscow State University

- Bill Tang
  Princeton Plasma Physics Laboratory

- George Mozdzynski
  European Centre for Medium-Range Weather Forecasts

- Peter Coveney
  Centre for Computational Science, University College London

- Jack Dongarra
  Electrical Engineering & Computer Science Dept, University of Tennessee