epcc news

The University of Edinburgh

www.epcc.ed.ac.uk

Issue 54, summer 2005

CONTENTS

New arrivals at the zoo

Editorial

Gavin Pringle

This year is EPCC's 15th anniversary and, to mark this epoch, this issue of EPCC News is focusing on the past, present and future of HPC platforms at EPCC.

The majority of articles are concerned with the most recent additions to EPCC's "zoo" of machines, such as the IBM Blue Gene platform, QCDOC and the new FPGA machine, amongst others.

Indeed, this zoo has become so large, we have built a new facility to house them, the new Advanced Computing Facility (ACF). The ACF was formally opened on the 1st of July by His Royal Highness, the Duke of Edinburgh, and this event is also covered within these pages.

The beginning of June saw EPCC host the IBM users group meeting, namely ScicomP11, and SP-XXL. Joachim Hein helped organise this successful event and reports from the front line.

For the second year running, EPCC had a booth at the International Supercomputing Conference (ISC) in Heidelberg, and we report some of the conference highlights.

Lastly, we present an condensed, paraphrasing of a conversation I had with Prof. Ken Bowler, one of EPCC's

original founders, just one month before he took early retirement. We chatted about his involvement with EPCC, supercomputers he has known, and his thoughts on what the future.

EPCC becomes one of the two UK High End Computing Training Centres

Adam Carter

EPCC, and the Centre for Scientific Computing at the University of Warwick, have been chosen by the UK's Engineering and Physical Sciences Research Council (EPSRC) to be the two UK High End Computing (HEC) Training Centres.

The centre at EPCC will provide postgraduate training at Masters level. Students who come to the centre will be undertaking a PhD in their own scientific discipline at their own UK academic institution but, instead of the usual three-year studentship, they will be on a longer, four-year HEC Studentship programme leading eventually to a PhD. Over the first three years, they will take a number of courses at EPCC followed by an HEC research project. This all leads to an additional qualification: an MSc in High Performance Computing from the University of Edinburgh.

To combine the PhD and MSc programmes successfully, students are formally enrolled as part-time students at the University of Edinburgh. They will usually visit us for two blocks of seven or eight weeks for the taught part of the course, spread out over their first three years of study. This will allow them to meet and share experiences with their fellow MSc students without having to spend too much time away from their home institution where they are doing the PhD research. Students will spend the first six months at their PhD institution and then come to EPCC for the first set of courses after Easter.

The training that we provide covers a wide range of useful and exciting areas in HEC. This includes how to write high-quality computer programs, how to make programs effectively utilise the world's largest supercomputers, how to use Grid technologies to harness the Internet effectively, visualisation of scientific results and the mathematical tools and techniques that underpin computational science and engineering. Graduates of the MSc in HPC will have acquired skills that make them employable in a wide range of careers ranging from postdoctoral academic research beyond their PhD, to jobs in commercial information technology.

The MSc in High Performance Computing is a well-established programme providing an excellent grounding in HPC technologies and their practical application.

Opening of the University of Edinburgh's Advanced Computing Facility

Alison Kennedy

The University of Edinburgh's Advanced Computing Facility (ACF) was formally opened on 1 July 2005 by His Royal Highness, the Duke of Edinburgh, Chancellor of the University of Edinburgh. The facility hosts a range of high performance computers with a combined performance of 20 teraflops and a vast storage area network of 155 terabytes, all connected to national and international networks by a 10 Gbit/s connection.

Around 70 guests attended the opening, including University colleagues, senior IBM management (three of the computers were installed in partnership with IBM), representatives of the Research Councils and of Scottish Enterprise.

Dr Arthur Trew, Director of EPCC, welcomed the guests. He said 'Computer simulation has, over the past 25 years, become recognised as the third scientific methodology, complementing theory and experiment. The emergence of simulation as this vital scientific tool is due, in no small part, to the pioneering work of Edinburgh researchers. It has a unique ability to investigate complex scientific phenomena which are too fast, too slow, too large, too small, or simply too expensive for the traditional approaches.'

The Chancellor's party then went on a short tour of the computing facilities, while Dr Alan Simpson, EPCC's Technical Director, gave a talk entitled 'EPCC: A Tradition of Excellence' on the history of supercomputing at Edinburgh, to the remainder of the guests. This was followed by a video highlighting some of the computational science projects making use of the facilities at Edinburgh. When the Chancellor's party returned from the machine room tour, the Chancellor took the opportunity to speak to some of the researchers and hear more about their work.

The formal opening then took place. Firstly, the Principal, Professor Tim O'Shea, spoke on the importance of the facility to computational science at the University. He also paid tribute to the many colleagues who had played a part in the setting up of the ACF. Next Professor Richard Kenway, vice-Principal for HPC and eScience, and Nick Donofrio, IBM's Senior Vice President of Technology and Manufacturing, spoke about the mutually beneficial, close working relationship between the University and IBM, which had helped to bring about the impressive concentration of high performance computers at the ACF. The Chancellor then unveiled the plaque. This was followed by lunch, a poster session and by further machine rooms tours.

The opening of the ACF underlines the University of Edinburgh's pioneering role in the use of computer simulation to tackle some of the world's most complex scientific puzzles. The ACF will enable scientists to make advances in such crucial areas of research as drug discovery, climate prediction and the design of new materials. It will also help researchers grapple with big scientific questions, such as gaining a better understanding of the structure of the universe, or, at the other end of the scale, exploring the make-up of sub-atomic elementary particles.

Since 1980, the University of Edinburgh has been pre-eminent in exploiting the largest available platforms in High Performance Computing (HPC) to push forward their research in a wide range of disciplines. EPCC, one of the leading HPC centres in Europe, acts as the focus for its interests in this area. The opening of the ACF, supported by an IBM Shared University Research award, is the University's latest strategic investment in HPC and its applications.

The University Principal, Professor Timothy O'Shea, said: 'The University of Edinburgh has a world-wide reputation in science, and also in informatics. These strengths intersect in High Performance omputing and we have enjoyed a strong position in this area through the leadership shown by our physicists in particular, by other elements of the scientific community, and by EPCC. The Advanced Computing Facility is evidence of our very strong and ongoing commitment to maintain that leadership and to support the most advanced possible scientific research.'

Computational science on Blue Gene

Lorna Smith

EPCC took delivery of the first IBM eServer Blue Gene system in Europe, in December 2004. This 2048 processor system offers 5.7 Tflop/s performance and provides unprecedented computing performance coupled with very low power consumption, floor space and cost. Designed to be highly scalable with low latency communications, the system offers a significant advancement in supercomputing technology.

The system is a key component of Edinburgh's computational science program and already a diverse group of scientists are exploiting the system.

For example, scientists from the Mathematics Department have investigated the problem of modelling financial assets and liabilities. The application of massively parallel computers to financial modelling is novel and has allowed Edinburgh scientists access to memory on a scale previously unavailable in this research field. Using our Blue Gene system, they have been able to solve a problem over 3 times the size of previous studies.

From the School of Physics, scientists have been looking at how proteins assemble from amino acid chains into biologically-functional structures. This is an important area, as protein mis-folding is attributed to diseases such as cystic fibrosis, CJD and Alzheimers. Also from this School, scientists are focusing on Nuclear fusion. With concerns over the long term strategy for providing energy, this is highly topical. They are simulating the behaviour of materials to understand the processes which govern damage in fission reactors. This work is essential in demonstrating what needs to be done to extend the lifetime of current-day reactors.

On the biological side, scientists are attempting to integrate biological data in an attempt to understand how biological systems function. The ultimate goal is to provide a basis for the true understanding of the complex network of processes involved in Life. Hence scientists at the Edinburgh Centre for Systems Biology are developing models of cellular, organ and organism level processes using new computational models that integrate gene, protein and metabolite data.

More traditional users include scientists from the School of Physics, who have been looking at designing new materials, using Lattice-Boltzmann simulations, and scientists from the School of Chemistry, who have been simulating material behaviour.

EPCC continues to work with these, and other research groups to ensure they can exploit the system effectively. The variety of science and the innovative nature of the machine, makes this an exciting time indeed!

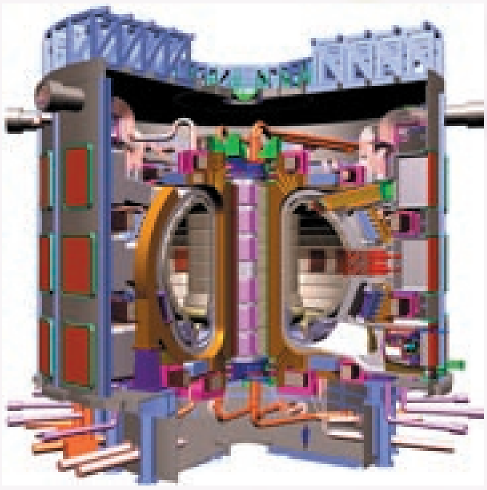

ITER

Working on QCDOC

Craig Morris

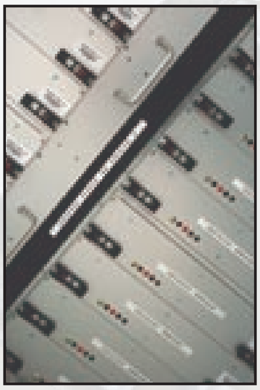

Top: single rack.

Above: Processors running on the motherboard.

Below: Craig standing beside 12 racks

Over the past year, EPCC has been involved with the installation of the first to be commissioned of three QCDOC (QCD-on-Chip) supercomputers in the world. It was installed at the University of Edinburgh's ACF, in November 2004. This purpose-built supercomputer has 12,288 processors and a peak speed of 10 teraflops. There are also a few smaller associated machines, which bring the total number of processors to 14,464. With an investment over 5 years of 6.6M, this supercomputer will be the most powerful computer in the world devoted to lattice Quantum Chromodynamics (QCD). QCD describes quarks and gluons, so these computers were designed to simulate particle physics.

I have been working for EPCC for three years now, and over the past year I have been heavily involved in the final construction, configuration, debugging and maintenance of QCDOC. While working on this project I spent eleven weeks in the US: two weeks at Columbia University (CU) and nine weeks at Brookhaven National Laboratory (BNL).

At CU I completed my training which involved learning about the construction, configuration, debugging and testing of the prototype machines. At BNL, where the main construction began, I was involved with the construction of putting the production machine (twelve racks) together. By the end of the sixth week at BNL I was then involved with the dismantling and packaging of this machine, which had to be shipped to the UK by a fixed deadline. While the UK machine was in the process of being shipped we began working on the second machine which was to stay at BNL. Then, within two very demanding months back at EPCC, reconstructing the machine and with lots of hard work and effort, we were able to get it into production by the end of January 2005.

Most of my time was spent running diagnostics and test programs on motherboards to see if all the processors were contactable and were communicating correctly and that the machine's network infrastructure was functioning correctly as well. To give an idea of the complexity and the amount of processors and networking on this machine, it had 192 motherboards = 12,288 processors, 2304 serial communication cables, 1608 ethernet cables and 78 network switches. The main advantage of the complexity of this machine is that of its scalability: the machine can employ between 2 and 12,288 processors, if required.

While in the US, I had a good time gaining experience and using some of the knowledge and experience that I had acquired while studying at the University of Edinburgh. On the social side I also had a great time travelling round the US while taking some long weekends. I was able to visit some friends in LA and relatives in Florida, where, to my surprise, I was allowed to drive a Chevrolet Suburban 5.8 Litre – you'd be amazed by the amount of gallons you get to the mile…

Personally, to be involved in a project of this size and magnitude and with the high risk of failure and high impact of success, I have found it very rewarding indeed and feel proud to be part of the team that made it work and turned it into a successful production machine. It has also been a really enjoyable experience for me to be involved in the challenges of making this highly public and important project a success. And now I look forward to continue providing a high quality service to our users.

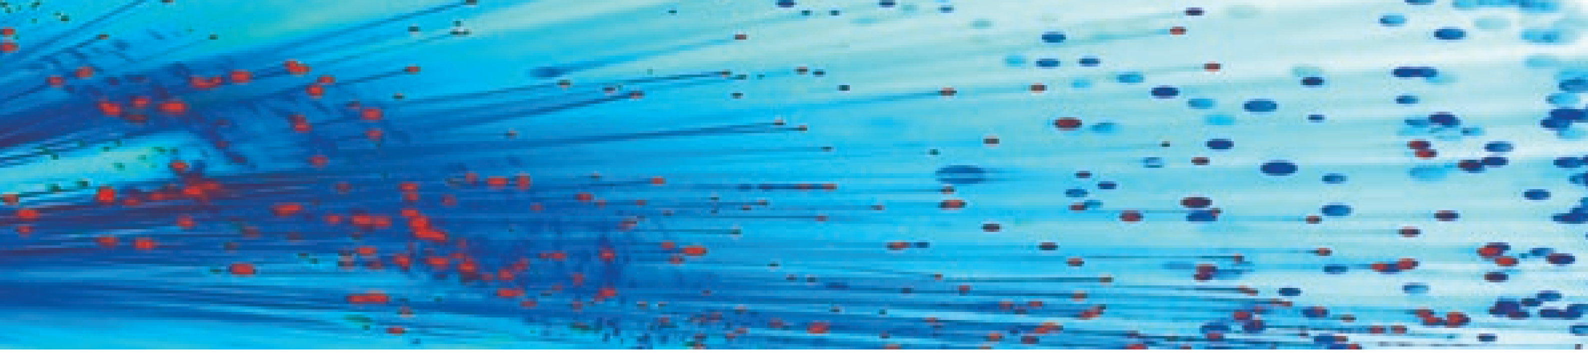

FPGA: a new breed of supercomputer

Rob Baxter, Mark Parsons, EPCC Graham Fairlie, Scottish Enterprise

Scotland is leading an ambitious project to build an exciting next-generation supercomputer with spectacular processing capabilities – and EPCC are at the forefront.

The FPGA High Performance Computing Alliance (FHPCA) will showcase Scotland's expertise in the development of computers built using ultra high-performance chips called Field Programmable Gate Arrays (FPGAs) and could revolutionise the way mission-critical applications are designed in industries such as defence, life-sciences and energy.

Scottish Enterprise has mobilised £3.6m funding and a world-class collaboration of industry, academic and government experts to lead the pioneering three-year project.

Managed by EPCC, Europe's leading high-performance computing centre, the project will draw on effort and technology from across the Alliance partners – Xilinx, the $10 billion NASDAQ-listed global supplier of programmable logic devices based in San Jose, California; Nallatech, based in Glasgow, which has pioneered the development of high-performance FPGA computing solutions for ten years; and Alpha Data, based in Edinburgh, which specialises in designing FPGA systems for customers including Boeing, Lockheed Martin and Motorola.

It will also draw upon the expertise of Dr Tom Kean of Algotronix who undertook pioneering work in this field, and the Institute for System Level Integration who will contribute to the dissemination of skills and knowledge into Scotland's design and application companies. The first commercial computer based on these devices was designed in Scotland in 1991 by Algotronix and this new project builds on Scotland's world-class reputation in the fields of both FPGA technology and high performance computing.

'What's really exciting about this project is that no one's ever tried to build a big general purpose supercomputer around these chips before,' says EPCC's Commercial Director Mark Parsons. 'People are always thinking up interesting designs, but the supercomputer we've come up with is an absolutely unique system. We're trying to join a whole lot of these high-performance chips together so we can tackle very large and exciting problems. It's a real opportunity for Scotland to take the lead in a hugely exciting area.'

FPGAs have potentially many times the processing power of conventional microprocessors because their hardware is reconfigurable – algorithms and complex numerical functions can be programmed directly onto the silicon, unlike conventional microprocessors which rely on software to give them flexibility – software which makes them slower. On FPGAs, certain algorithms have been shown to run over 1,000 times faster than the fastest off-the-shelf microprocessors.

Systems built with FPGAs are also considerably more compact than conventional systems and a new design can be reconfigured by downloading it and then, by using software tools, activate their internal 'switches'.

Until now, FPGAs have been too small in terms of capacity and too expensive to make a supercomputer system feasible. But in recent years the chips have got much cheaper while their capacity has massively increased, meaning they are now becoming widely adopted in mainstream embedded computing.

Continued opposite.

article
>plea

Do we
Walla
Tracy?

EPCC forges unique alliance to advance FPGA-based high performance computing

Maureen Wilkinson

EPCC joined its partners in launching the FPGA High Performance Computing Alliance (FHPCA) at an industry conference held at Our Dynamic Earth in Edinburgh on Wednesday 25th May 2005. In excess of 100 guests attended the launch, including FHPCA Alliance partners and representatives from leading UK and US universities, the Scottish Executive and Scottish Enterprise.

The Right Honourable Jim Wallace MSP, Deputy First Minister of Scotland, Minister for Enterprise and Lifelong Learning, welcomed guests, saying: 'If the Executive is to achieve its key aim of continuing to grow Scotland's economy we need to encourage effective collaboration between industry and academia, to develop new technological products and to explore new markets.

'The High Performance Computing Alliance is an excellent example of this type of successful partnership working - bringing together some of the world's leading computing experts to develop truly ground-breaking technology which will place Scotland at the forefront of microelectronic innovation and design.

'The supercomputer that the Alliance is creating will have far-reaching benefits in areas such as life sciences, climate change and games technology and ultimately generate new knowledge and wealth for Scotland.'

Next, the Principal of the University of Edinburgh, Professor Tim O'Shea, thanked Jim Wallace and spoke on the importance of the Alliance to the University of Edinburgh and paid tribute to those who had played a part in establishing the Alliance. Following lunch, during which guests had the opportunity to visit the stands of several Alliance partners, Dr Mark Parsons, EPCC's Commercial Director, gave a talk explaining how the project would build on Scotland's reputation in the field of FPGA computers.

'This is a fantastic Scottish technology that was developed here and is being exploited by Scottish companies,' he said.

'What's really exciting about this project is that no-one's ever tried to build a big supercomputer with these chips before. We're trying to join a whole lot of these high-performance chips together so we can tackle very large and complex problems. It's a real opportunity for Scotland to take the lead in a hugely exciting area.'

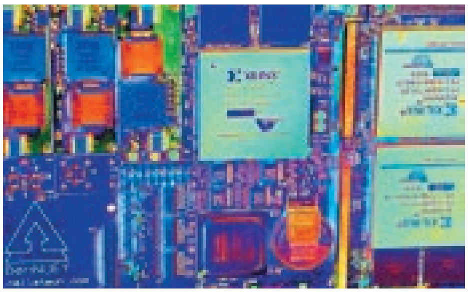

The partners aim to use off-the-shelf components and custom software to produce a computer with 1 teraflop of processing power, equivalent to 200 times the speed of the highest performing Pentium processor.

Rather than targeting one particular application or market, this supercomputer is designed as a general-purpose system targeting many potential applications and markets. The partners initially aim to develop two or three demonstration applications to illustrate the scope and scale of the supercomputer.

These are likely to target established markets that demand huge processing power, such as modelling the behaviour of complex systems; designers requiring ever-greater processing capacity; imaging systems in industries including defence and health; drug design and bio-analysis, particularly in relation to DNA and the human genome, oil exploration; and large volume data processing in industries such as financial services.

Once complete the supercomputer will be available for third-party access, allowing researchers, manufacturers, developers and government organisations worldwide to test and develop high-performance FPGA applications and devices.

If you have any questions about the FHPCA Alliance, please contact m.parsons@epcc.ed.ac.uk.

FHPCA Alliance Members:
Algotronix – www.algrotronix.com
Alpha Data – www.alphadata.co.uk
EPCC – www.epcc.ed.ac.uk
The Institute for System Level Integration – www.sli-institute.ac.uk
Nallatech – www.nallatech.com
Xilinx – www.xilinx.com

QCDOC and Blue Gene Workshop

Next Generation of HPC Architecture National e-Science Centre, Edinburgh, 4th–6th October, 2005.

Lorna Smith

The Blue Gene and QCDOC architectures offer a significant breakthrough in supercomputing technology. With the growing popularity of these systems, a workshop investigating the full potential of these machines across a wide range of applications is highly topical.

The first IBM e-Server Blue Gene system in Europe was delivered to EPCC, the University of Edinburgh, in December 2004. The University has also recently installed a QCDOC machine, providing one of the world's most powerful systems dedicated to the numerical investigation of Quantum Chromodynamics. As the only centre in the world to host both a Blue Gene and QCDOC system, the University is in an unparallelled position to hold such a workshop.

The program consists of a one-day tutorial giving the attendees the opportunity to gain hands-on experience on the Blue Gene system. This is followed by a two-day workshop, with presentations from users, developers and programmers utilising both systems.

We have an excellent range of key-note speakers, from US organisations such as Lawrence Livermore National Laboratory and Argonne National Laboratory, and from the key Blue Gene sites in Europe such as ASTRON, Jüelich and, of course, Edinburgh.

The workshop will be sponsored by IBM, and hosted by EPCC, University of Edinburgh at the National e-Science Centre, Edinburgh, UK.

Keynote speakers:
Steve Louis, Lawrence Livermore National Lab
Kjeld van der Schaaf, ASTRON
Henry Markram, EPFL
Thomas Lippert, NIC, Juelich
Andrew Siegel, Argonne National Laboratory
Eric Kronstadt, IBM
Arthur Trew, Director, EPCC

...Stop Press...

Blue Gene System Software Workshop, Edinburgh, UK

We will be holding a System Software workshop in conjunction with the QCDOC and Blue Gene Applications Workshop.

The workshop will have agenda's supporting the System Software and Management aspects of Blue Gene.

Register now at:
http://www.epcc.ed.ac.uk/BGworkshop/

Self-assembly and hydration in biological systems: lessons from small molecules

Jason Crain

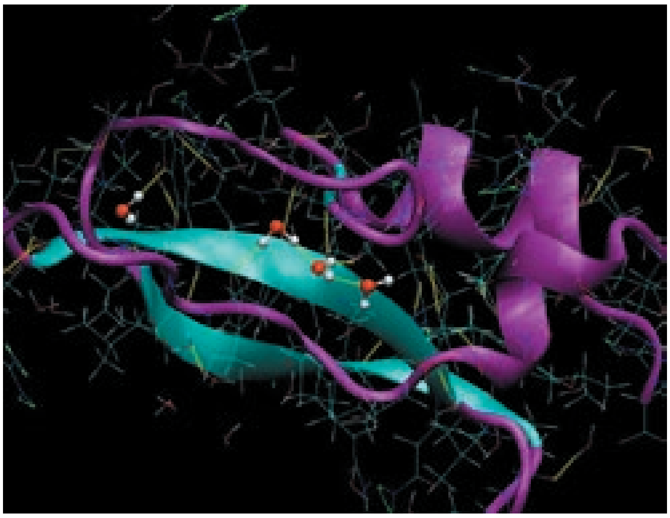

Understanding the principles and processes by which proteins in aqueous solutions assemble from amino acid chains into biologically functional three-dimensional structures is recognised to be one of the major challenges spanning the physical and life sciences. In addition to fundamental interest in the factors surrounding this process, protein mis-folding events are also attributed to diseases such as cystic fibrosis, Creutzfeldt-Jakob disorder and Alzheimers. Therefore, the long-term impact of advances in this area also reaches into medicine and may have future ramifications for protein and peptide design. The basic molecular-scale issues that underlie folding and other related 'self-organisation' phenomena are very general and concern the following themes:

1) Hydration structure: How are the water molecules organised around the polar, non-polar and peptide groups of complex biomolecules? How is this organization influenced by externally-controllable conditions such as temperature and pressure?

2) Molecular flexibility: How flexible are the bonds of complex molecules that define overall macromolecular shape? And how is this flexibility affected by hydration structure and other environmental factors such as pressure and temperature?

3) Ionic (Hofmeister) Effects: How does the presence of simple inorganic ions influence hydration structure and flexibility?

The size, complexity and diversity of real proteins and biomolecules still renders a thorough molecular-level investigation of these issues impractical for systems of direct biological relevance. A promising emerging strategy in the study of processes with inherently high levels of structural complexity, multiple degrees of freedom and competing interactions makes use of so-called model systems which represent extreme approximations of 'real' macromolecular materials. Chosen carefully, and with appropriate combinations of computational and experimental input, coordinated investigations on model systems can be very powerful leading to insight into much more complex systems. In the specific case of the protein folding problem and self-organisation phenomena in biomolecules, there are potentially many views as to what level of abstraction relative to a real biomolecule is 'useful' for a particular purpose. Ultimately, there must be some compromise found between the degree of 'realism' of the system and the level of detail at which the system is described. In those cases where the fundamental questions are general rather than specific to particular biomolecules, extreme simplifications of real biological systems may be made. The benefit of this approach is that very sophisticated new computational (such as polarisable models and path integral molecular dynamics) and experimental tools (isotope-labelled neutron diffraction and empirical potential structure refinement) can be deployed together for the first time to build confidence in new models.

With this strategy in mind, a major collaboration between the University of Edinburgh and IBM Research has been established to combine massively parallel supercomputing with high-resolution neutron diffraction experiments at atomic resolution. To support the project IBM has recently announced a major Shared University Research (SUR) award to Edinburgh including leading-edge POWER-5 microprocessor technology (16-way 1.9 GHz P5 570 + 128 GB RAM) and high performance graphics workstations enhanced by world-class 9 million pixel LCD monitors for visualization of complex structures. The pSeries system will be housed at the ACF, where support for the machine will be provided by EPCC. The project also makes extensive use of the Blue Gene systems in both the ACF and in the TJ Watson Center.

Interview with Professor Ken Bowler

Prof. Ken Bowler was one of EPCC's original founders, the so-called Gang of Four, along with Richard Kenway, Stuart Pawley and David Wallace. Gavin Pringle spoke to Prof. Ken Bowler prior to his retiral. Here is a paraphrased version of that meeting.

What was the first parallel machine you were involved with?
That was the ICL DAP, in 1980, housed at Queen Mary's College (QMC) in London. Stuart Pawley was really the pioneer, using it for molecular dynamics simulations.

It was that time we realised that our problems were suitable for that kind of architecture. We got together with Stuart and did an initial study with the QMC machine, which was an interesting exercise in itself, because we had to use an acoustic modem to link up to the machine in QMC, and we got our results back through the post!

These were interesting results, and we spent some time talking to Stuart Reddaway, from ICL, who was an architect of the DAP. We formed a good relationship with ICL which resulted in us getting one machine here in '81, and a second AMT DAP in '88. The machines were hosted by Edinburgh University Computing Services (EUCS). The then deputy director of the EUCS, Andrew McKendrick (who sadly died last year) and Mike Brown, the DAP's system manager, helped us a lot in those early days.

ICL DAP.

The Meiko Computing Surface.

So the DAP, for us, was very successful. It enabled us to compete with people elsewhere who had time on big, expensive vector machines like the Cray and the Cyber 205, whereas the DAP was relatively modest, but, as it has a parallel architecture, we were able to compete effectively.

The DAP was an adaptation of a memory module of the 2900 Series ICL mainframe computer, which were great big monster machines, where it had processing elements wired into these memory modules. In a sense, it was a back end processor, where the code ran on the ICL mainframe host machine, and the DAP subroutines were executed on this rather bizarre piece of iron-mongery.

After the ICL DAP, you moved to a Meiko platform.
KB: This also predates EPCC. It was 1986 when we got our first Meiko machine; it was a very small Meiko, called the Computing Surface, and it had the old T414 integer transputers as the processing elements. But, the original machine wasn't a serious compute engine for QCD problems, as it didn't have the computational oomph. But it was an interesting learning experience for us, dealing with a MIMD architecture, compared with the old DAP SIMD. We did look at a few particle physics calculations, but they were really toy calculations, rather than QCD in anger. I can still remember coming in at Christmas and New Year and putting in Exabyte tapes in the Meiko machine to save our data. It was very much a string-and-sealing-wax kind of operation.

The follow up to the Meiko machine was a big transputer machine called the Edinburgh Concurrent Supercomputer and that was about that time that EPCC was established. A range of different applications were ported to that machine.

It was a very big transputer-based computing surface, with 400 T800 processors; it was certainly a serious machine. That was an interesting exercise because it was essentially the start of UKQCD collaboration. A group of us from across the UK working in lattice-gauge theory/lattice-QCD decided that the only way forward was by collaboration, and put together a joint proposal for a big machine to do lattice QCD.

We got awarded a grant from the Nuclear Physics Board of the Science Research Council, which funded the purchase of the Meiko i860 box which, by the standards of the time, was quite a powerful machine. It had a peak performance of 5 gigaflops, which was a big deal in those days!

That was the first instance of a UK-wide national service, but it was only for one consortium?
It was for one consortium, but then it was extended. We got more memory and a second smaller machine, by collaborating with the Car-Parinello Consortium, led by Dominic Tildesley from Southampton. That collaboration really launched UKQCD on to the world stage and we got a lot of good physics out of that machine from 1990 to 1995, or so.

So Edinburgh Parallel Computing Centre emerged as a support centre for these consortia. And then we got the Thinking Machines' CM-200 box.
KB: Yes, the CM-200 was used by some people of UKQCD for some of our calculations, but we didn't have exclusive use of it, nor exclusive ownership, so it was a relatively small contribution to the overall computing needs that we had.

People may have thought that the flashing light on the front of the box had a relationship to each processor, and that the light patterns mirrored the current calculation; however, it seems this was just a random pattern just to make it look interesting!
I heard a similar story as well. It was piece of marketing. It looked really nice. They also had this wonderful data store that looked like a cocktail bar: the Data Vault!

Then, after the CM-200, came the Cray T3D, which was funded as a national computing resource and various consortia bid for time on that, including UKQCD. We made very good use of the T3D, as we had a pretty significant allocation of time on it; however, it was a change of model for us. The Meiko had been very much our own baby: an in-house machine, whilst the T3D was run as a shared resource with a professional service. But following on from that, we converted back to the old model, applying to PPARC for a large-scale QCD machine. There was procurement exercise carried out and the outcome of that was that was the Cray T3E.

Again, it was a shared facility amongst three PPARC consortia: Virgo, ourselves and the MHD Consortium. So, the T3E was our mainstay for two or three years or so.

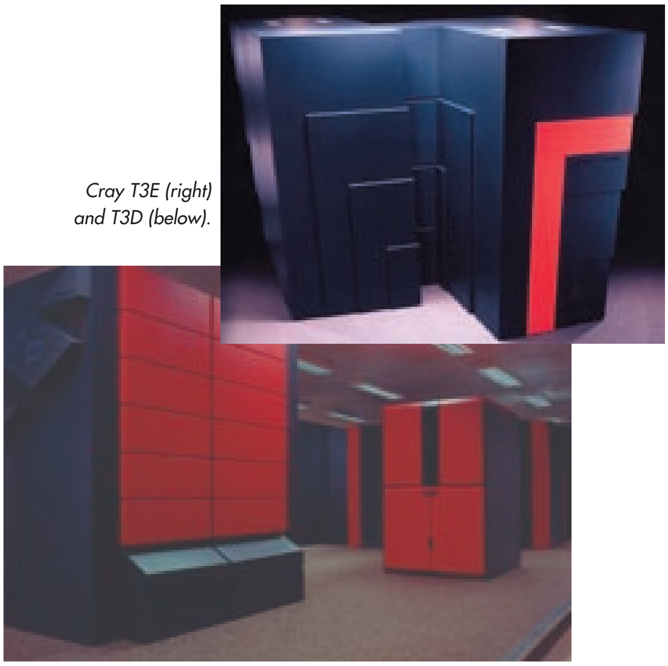

Cray T3E (right) and T3D (below).

There was something of a hiatus for us after the demise of the T3E. The demands on QCD were going up dramatically and we had to be more ambitions and think on a bigger scale, and that was the thinking that lead to the QCDOC project.

We put together a bid under the JIF heading (Joint Infrastructure Fund) without too much hope of success, as JIF didn't typically fund supercomputers. The bid was originally based on a machine called QCDSP, designed and built in Columbia University (CU), by Norman Christ's group. That had been a successful machine and a big version of it had been built and installed in Brookhaven. It was based on Texas Instruments Digital Signal Processing chip with some custom interconnect.

But in the meantime, Norman's ideas became more ambitions and he started talking with IBM. There was now an upgraded project based around QCDOC which involved IBM and we were funded, somewhat to our surprise. After the open procurement, we went into partnership with CU and IBM.

There's around 10 teraflops out at the ACF and there are two similar machines out at Brookhaven: one for the Brookhaven/

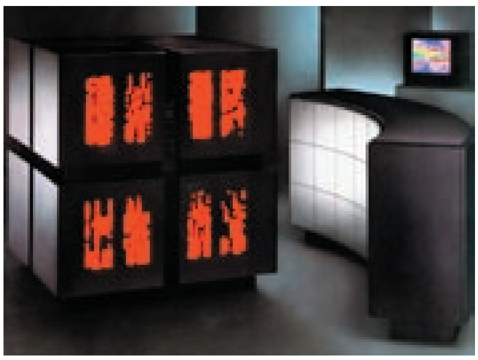

Thinking Machine's CM-200 and the Data Vault.

CU collaboration, and one funded by a US DoE for the US lattice community.

This sounds like great news for British QCD!
It is, it is. It puts us at the leading edge world-wide. Actually, we're starting to get results, but most of the calculations we've done so far are exploratory. We're looking in a relatively big parameter space, to find exactly where we should be running our big simulations. Within the last couple of weeks we've settled on a parameter set and we've just started on some large volume simulations.

I understand that QCD is only possible thanks to supercomputers?
Absolutely, yes. Indeed, it's fair to say that you always need more computing resource for lattice QCD, since it's an open-ended problem. What we are trying to do is simulate a portion of space and time with increasing accuracy, using Monte-Carlo methods, so the results we do have a statistical error associated with sampling over the ensemble. There's also the systematic errors of both simulating a finite volume, rather than the infinite volume of space and time, and the discretisation of that space.

However, the big break through for us is that the new generation of machines enables us to abandon the so-called Quenched Approximation, which we had to make just because of computational cost. In a quantum field theory, what you call the vacuum isn't actually empty. Until very recently, we had to make an approximation where you had to neglect that effect, but QCDOC gives us now the ability to do what's called a Full QCD Simulation, taking into account this vacuum polarisation effect. That's the big difference.

The downside is that it costs an enormous amount more to do the computations; it's a factor of 1000 to 10,000 times more compute time than before. This is important, however, as our experimental colleagues need results from us to extract fundamental parameters for the Standard Model of Particle Physics. They can't get them directly: there's an intermediate step that connects the results for decay rates with the fundamental parameters of QCD which requires us to give the experimentalists what we called QCD Matrix Elements. We are under some pressure to really squeeze down our errors so we can give them computations to an accuracy of a few percent.

How long do you expect QCDOC to be switched on for?
Tough question. We would hope to have a sensible working life in the order of 5 years, and certainly a minimum of 3, but you know as well as I do how fast the technology changes. So people have already begun to think about the next generation. There's already a competition between custom-built machines, like QCDOC, and more general purpose clusters. At the moment, it's probably true that a machine with the architecture like QCDOC has the edge. I mean, it's highly optimised for our problem, in terms of bandwidth to memory in particular, which is the bottleneck in QCD calculations. The QCDOC chip has got a relatively large on chip memory, and a very fancy controller for pre-fetching data.

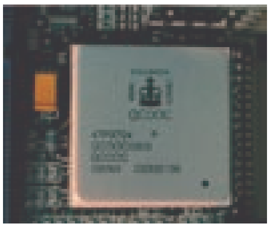

There is competition from larger clusters, but price/performance wise QCDOC has the edge at the moment. And certainly, clusters have a problem in terms of power consumption.

Actually, our machine is quite modest in terms of clock speed, as we had to clock down a wee bit to make sure that the machine runs reliably and reproducibly. You want to be confident that the machine will be stable for runs which take around 6 months.

What is interesting is how much more care we are taking, compared to someone who is using a big commercial machine: they can just take it for granted that the machine runs correctly.

We've got various methods of checking reproducibility, such as sophisticated checksum routines built into our codes, which are monitoring communications between processors. The real whiz on that is Peter Boyle, who is one of the chief architects on QCDOC. He went out there as a UKQCD person to collaborate with CU and ended up playing a very major role in the design and running of the project.

Where does QCDGrid fit into all of this?
It's helping us as collaborate. QCDGrid is the distributed data repository for our lattice data. We have a number of collaborating sites, several UK groups and a node at CU. As QCDOC generates these ensembles of gauge fields, at least two instances of the field data are put on QCDGrid and, with increasingly sophisticated tools and user interfaces, people can look at the datasets that they want, pull the data down onto their machine and processes it.

Is QCD your only area of interaction with EPCC?
In the early days, we ran training courses in OCCAM, when the Meiko machines were here, for anyone to come and learn parallel programming.

I've also been involved with was the TRACS visitor program (the forerunner to HPC-Europa). I was involved in both the scientific assessment panel and in hosting visitors. It was good and I met a lot of interesting people. I believe the visitor program did EPCC a lot of good. There was a two-way flow of people.

My final involvement is through the MSc, where I've been chair of the board of examiners since the start.

And what about the future for EPCC?
I guess collecting machines and hosting them isn't really EPCC's prime function: that is really technology transfer. I would have thought that the demand for that can only grow, so prospects for the future for EPCC are still very bright.

I think it's great that they've been involved in supporting HPC, and I hope they continue to do so. They've got a tradition of attracting very high quality young people, with HPC experience. I think there's a very fruitful interplay between computational science and the sort of thing EPCC is interested in doing, and with any luck that will continue.

Then there's also the big unknown in what direction the Grid will take. How much that will destroy the conventional machine with its local group of users. Obviously, if the Grid lives up to its promise, that kind of model will disappear, as people will just tap into the resource they need, although I'll believe it when I see it.

What plans do you have for your retirement?
I used to play the clarinet, so I'm hopefully going to pick that up again and maybe learn the keyboard, which is very different. I hope to do a bit more hill-walking, and keep my hand in physics, as far as I'm able.

But then, like any true academic, you won't be retiring completely, will you?
Absolutely, no. What I'm giving up are the admin jobs I had and the teaching responsibilities, so I hope I'll still be part of the research excitement.


Ken with QCDOC.

EPCC at Sun HPC

Gavin Pringle

EPCC attended the Sun HPC meeting in Heidelberg on 20-21 June, 2005, where Kevin Stratford presented a talk entitled 'Cosmology using MPI on the Grid'.

This talk described the work that Kevin is carrying out for a joint research acitivty of DEISA, namely Cosmolocal Applications. Specifically, Kevin presented results on running a modified version of GADGET2, the famous cosmological simulation from the Virgo Consortium, over two Sun SMP platforms using PACX-MPI.

The practice of running a single MPI program, using N sets of processors on N separate supercomputers, is known as meta-computing. It is considered that meta-computing is only suitable for loosely-coupled applications, where each sub-computation is relatively independent of the other sub-computations and can be executed within each individual supercomputer.

However, we have found that this cosmological application, which is obviously not loosely-coupled, can be run efficiently in a meta-computing environment. Indeed, the initial results from this experiement, as presented at Sun HPC, are both surprising and extremely encouraging. Hopefully, this work will allow Virgo to run far greater simulations of the Universe than ever before, employing multiple supercomputers and their combined availble memory.

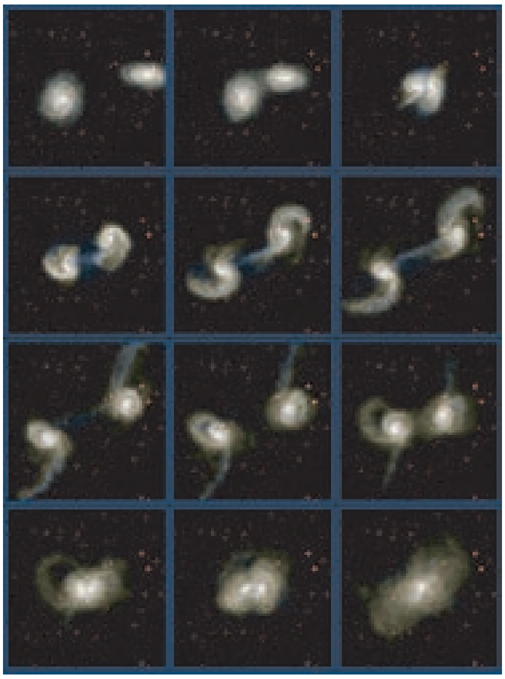

For more information on Sun HPC 2005: www.sun.com/products-n-solutions/edu/events/archive/hpc/heidelberg05.html
DEISA: www.deisa.org
Virgo Consortium: www.virgo.dur.ac.uk
GADGET: www.mpa-garching.mpg.de/galform/gadget
PACX-MPI: www.hlrs.de/organization/pds/projects/pacx-mpi

Global cooling in the supercomputing ecosystem?

Kevin Stratford

Heat was definitely on the agenda at the 20th anniversary meeting of the International Supercomputer Conference held in Heidelberg in June. As delegates wilted in the baking mid-summer temperatures, a number of speakers were also concerned with heat: modern supercomputers just produce too much of the pesky stuff. Large machines are increasingly a strain not only on building space, but also on the electricity supply needed to power large numbers of highly clocked chips and the air conditioning required to remove the resultant waste heat in order to prevent an expensive meltdown.

So it was interesting to note that the leader in the new Top500 [1] list of supercomputers, presented by Erich Strohmaier of the Lawrence Berkeley National Laboratory in California, provided a promise of more temperate climes ahead: the 65,536 node Blue Gene system at Lawrence Livermore National Laboratory (also in California) has clocked 136.8 teraflops on LINPACK, the standard benchmark used to produce the list. The Blue Gene system is interesting because its designers have specifically taken pain to address the problem of space and heat efficiency [See Lorna Smith's article on EPCC's Blue Gene on page X]. Each chip, or node, has two processors with relatively modest clock rate and hence power consumption, allowing many of them to be packed closely together. This should allow quite powerful beasts with hundreds of thousands of processors to be housed in existing machine rooms.

The concomitant issue of scaling real scientific problems to such large numbers of processors appeared in the following talk of Horst Simon, again from Lawrence Berkeley National Laboratory. To mark the 20th anniversary of the conference, Simon gave a fascinating personal overview of the top breakthroughs in supercomputing in the preceding 20 years, and what he considered the most important challenges to be faced in the next 20. As is traditional at such events, these were announced in reserve order. Number three on the list of challenges ahead was exactly the problem of huge machines and scaling real problems to many thousands of processors. Simon likened this shift in the emphasis of supercomputing to a change in an ecosystem, where new dominant species will feed off a changed application landscape. The top of the list of challenges, which continued the ecological metaphor, he considered to be 'living in a world of no growth', the change that will inevitably occur when the growth of processor power (famously described by Moore's law --- that processor speed will double roughly every 18 months) will finally level off. Simon also noted the worrying lack of current fundamental research and new thinking to enable such limits to be circumvented.

(For those wondering what Simon rated as the top three breakthroughs of the last 20 years, they were, in reserve order: 3) the introduction of Scientific Visualisation to really bring home the impact of results of large numerical simulations; 2) the availability of the Message Passing Interface standard; and 1) John Gustafson's work [2] on scaled speed-up. Simon reminded the audience of the heated debate that took place in the early days of supercomputing about the ultimate limits of scalability suggested by Amdahl's law and the work Gustafson did in overcoming this 'mental block'.)

So, less power-hungry machines may lead to computer facilities in future to be able to operate without costly air conditioning. This correspondent, at least, hopes the same is not true for the Dresden conference centre, where next year's ISC meeting will be held, preferably in cooler conditions.

[1] http://www.top500.org/
[2] J.L. Gustafson, G.R. Montry, and R.L. Benner, SIAM J. Sci. Stat. Comput., Vol 9, 609 (1988)

ScicomP11 hosted by EPCC

Joachim Hein

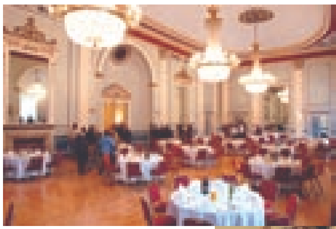

As part of the events marking EPCC@15, EPCC hosted the 11th meeting of ScicomP in Edinburgh from the 31st of May until the 3rd of June 2005. The meeting was co-hosted with the summer meeting of SP-XXL.

ScicomP is an international organisation focusing on the performance and scalability of scientific applications on IBM HPC platforms. SP-XXL's interests are similar to ScicomP; however, their focus is more systems oriented whereas ScicomP is more application driven. With EPCC being one of the key players behind the HPCx system using IBM's Power4 technology and operating the first European IBM eServer Blue Gene, the interests of both groups and EPCC match very well. We were able to welcome over a hundred participants in Edinburgh.

As a member of the local organising committee for ScicomP11 and having attended the ScicomP sessions, I will report on the event from a ScicomP perspective. The meeting started with the traditional tutorial, focusing on IBM's Power technology. Charles Grassl (IBM) described the features of the latest Power5 processor, highlighting differences to the previous Power4 processor. Roch Archambault (IBM Compiler Development) explained how to use the compiler to make efficient use of the processor. With Power5 being an option for the forthcoming HPCx Phase3, the tutorial proved to be of great interest.

The second day started with the keynote speeches by Richard Kenway (Assistant Principal of the University of Edinburgh and Chairman of EPCC) and Rebecca Austen (Director of IBM Deep Computing Marketing). The afternoon saw HPC users from around the world describe their experiences with IBM systems. Some speakers discussed code improvements while others reported on the latest tools to aid performance analysis. It was interesting to see how other people solve similar problems to the ones I face on a daily basis.

Most of the third day centred on Blue Gene technology. With a 65536 processor Blue Gene system being the fastest computer and an increasing number of Blue Gene systems being installed at various computer centres in the world, these presentations were very timely. Manish Gupta and Gyan Bhanot (both IBM) reported on the present and future technology, including application performance. Lorna Smith (EPCC), Tom Spelce (LLNL) and Richard Loft (NCAR) reported on experiences their centres made with their Blue Gene systems. On the one hand, the presentations showed that there are a number of teething problems left, whilst on the other hand, they showed a number of early successes, demonstrating the exciting potential Blue Gene systems have to offer. In my view, this is a clear case of 'watch this space'.

Paul Madden (University of Edinburgh) reported on the scientific software development within the UK Car-Parinello Consortium. The presentation focused on the constraints imposed by the academic environment and how to ensure quality standards are maintained for the application code.

The last day of the program offered a mixture of IBM presentation and talks contributed by our participants. These presentations focused on the technology layers underlying one's application. In particular, the presentations on MPI by Dick Treumann and the compilers by Roch Archambault (both IBM) contained many hints on tweaks which might improve application performance on our systems.

As part of the social program IBM invited everyone into the Surgeons' Hall on Wednesday night. The main social event was on Thursday evening, starting with a walking tour through Edinburgh's Old and New Towns, which was followed by a whisky tasting reception and a dinner in the Georgian splendour of Edinburgh's Assembly Rooms.

As is typical for ScicomP meetings, there were many opportunities to discuss matters with IBM staff members and experts in the field from other computer centres. To me these exchanges are as valuable as the actual presentations. In summary these were four exciting days. On behalf of the organising committee I would like to thank everyone who participated in the meeting and made it such a success. The next meeting of the series will be hosted by NCAR in Bolder, Colorado, starting on the 17th of July 2006. You should seriously consider going.

[1] For further information on ScicomP: http://www.spscicomp.org
[2] For further information on SP-XXL http://www.spxxl.org/
[3] Slides of the presentations given at the ScicomP meeting are available http://www.spscicomp.org/ScicomP11/

SuperComputing, 2005

Adrian Jackson

As in previous years, EPCC will be exhibiting at this year's conference for High Performance Computing, Networking and Storage; SuperComputing 2005. SC|05 is being held in the Washington State Convention and Trade Center, in Seattle, from the 12th to the 18th of November.

Under this year's theme of the 'Gateway to Discovery', SC|05 promises to be an exciting mix of exhibitions, lectures, tutorials, and technical events from the world of High Performance Computing.

As well as having a booth in the exhibition space, we will also be giving a half day tutorial on 'Application Performance on the Blue Gene Architecture', presented by Dr Lorna Smith, Dr Mark Bull, and Dr Joachim Hein.

If you are attending SC|05, you can find us at booth number 136: hope to see you there!

http://sc05.supercomputing.org

Image courtesy of Seattle's Convention & Visitors Bureau

Third HPCx Annual Seminar

The 3rd HPCx Annual Seminar will take place on Monday 5 December, 2005.

It will be held in the Merrison Lecture Theatre, in the CCLRC Daresbury Laboratory, Warrington, UK.

Further details and registration information can be found at: http://www.hpcx.ac.uk/about/events/annual2005