Issue 94 AUTUMN 2023

The newsletter of EPCC, the supercomputing centre at the University of Edinburgh

epcc news

Winners!

TeamEPCC brings home the ISC23 Student Cluster Competition trophy

From our Director

Welcome to the Autumn 2023 edition of EPCC News. It's difficult to believe we've now been based in the Bayes Centre for five years.

Despite the complexity of the move in 2018, the Bayes has proven a great building to be based in. We've not lost contact with the many people we continue to work with at King's Buildings, but we've also met many new researchers from across the University since we moved.

EPCC goes from strength to strength. We've grown considerably since we moved to the Bayes and now have 135 staff, 14 PhD students, 60 on-campus MSc students and a further 60 online MSc students.

This is an exciting time for us. Having been confirmed as the hosting site for the UK's first Exascale supercomputer, we're in the middle of complex preparations in Computer Room 4 involving decanting all of its existing equipment into the other rooms and adding the power and cooling capability for what is likely to be a 20MW system.

In May I was incredibly proud of our MSc students who won the ISC'23 Student Cluster Competition in Hamburg, Germany. With 22 teams from across the world this is a major accomplishment. Our team was the outright winner across the competition.

I hope the breadth of our activities comes across in this issue. Many of the projects we're now involved in are benefitting from the past two years of development work on the Edinburgh International Data Facility. This complex infrastructure is now delivering a unique service for advanced data analytics and AI for our region.

As ever, please do let us know if you enjoy this issue and if there is anything you would like to hear about in more detail in a future issue.

Professor Mark Parsons
EPCC Director

Contents

3 **ISC23 Student Cluster win**
TeamEPCC named overall winner

4 **Five years at the Bayes**
A review of events since our move

6 **Better weather forecasting**
Collaboration with UK Met Office

7 **Women in HPC**
Working for greater inclusivity

8 **Reducing the cost of large-scale simulations**
SiMLint project review

9 **Quantum is coming**
EPCC's Quantum Computing Group

10 **New Chancellor's Fellow**
Introducing Joe O'Connor

11 **RSE Personal Research Fellowship**
Our new Fellow describes his work

12 **AI services for business**
Our cutting-edge systems are available for use by industry

13 **Accelerating research**
New EIDF GPU Service now open

14 **Advanced Computing Facility**
Ensuring the ACF remains world-class

16 **Scottish Medical Imaging Service**
Our work with Public Health Scotland

17 **HPC Summer School**
Review of international gathering

18 **Software Sustainability Institute**
Recent activity

20 **MSc programmes**
Saying farewell to the class of 2022–23

22 **Outreach activities**
The Big Bang Fair 2023

23 **Schools research visits**
Nuffield Research Placements

www.epcc.ed.ac.uk
info@epcc.ed.ac.uk

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

TeamEPCC wins ISC23 Student Cluster Competition

Image left shows (left to right): Spyro Nita, Tomas Rubio Cruz, Jaffery Irudayasamy, Ikraduya Edian, Hristo Belchev, Oleksandr Piekhota, Xu Guo.

Above: (front row) Kris Tanev.

In May TeamEPCC took overall first place in the ISC23 Student Cluster Competition in Hamburg, Germany.

The ISC23 Student Cluster Competition brought together 22 teams from across the world, both virtually and in person. TeamEPCC was one of seven to compete in the conference hall, with another fifteen teams competing virtually. Over three days all on-site teams demonstrated their abilities in obtaining the best performance from their computing clusters across a series of benchmarks and applications. Each team designed its own cluster, with TeamEPCC's created with the support of sponsor Hewlett-Packard Enterprise (HPE), and the EPCC systems team.

TeamEPCC performed outstandingly and was the outright winner across the competition. The winning group comprised five students from our MSc programmes in High Performance Computing (HPC), and HPC with Data Science, plus a third-year undergraduate student intern from Edinburgh Napier University's Networking and Cybersecurity programme: Hristo Belchev; Ikraduya Edian; Jaffery Irudayasamy; Oleksandr Piekhota; Tomas Rubio Cruz, and Kris Tanev (student intern).

Members of TeamEPCC give their reactions to their incredible win.

"I think what led us to winning was the collaborative spirit of the team as well as the drive of all its members - I felt like I was part of something much bigger than myself and that every member was giving their best to propel everyone forward."
Hristo Belchev

"To future students/team members: if you are motivated, have a passion for HPC and want to learn beyond the curriculum, joining TeamEPCC is the best thing you can ever do!"
Ikraduya Edian

"Preparing for the competition was exhilarating and challenging. Over six months, we learned, strategized, and honed our skills in high-performance computing (HPC). It was an opportunity to push our limits and apply our knowledge."
Jaffery Irudayasamy

"Being part of EPCC and being able to support our MSc students in HPC is a part of my career that I will remember for the rest of my life."
Kris Tanev

"We have met many interesting people, gotten into, and understood many modern HPC trends, seen many HPC companies, and made many connections."
Oleksandr Piekhota

Xu Guo, Ben Morse, Spyro Nita, EPCC
x.guo@epcc.ed.ac.uk
b.morse@epcc.ed.ac.uk
s.nita@epcc.ed.ac.uk

ISC High Performance

The ISC SCC is an annual event for all the selected student teams to demonstrate their abilities to design their own clusters and achieve the best performance on a selection of benchmarks and scientific applications within a certain power usage limit.

We thank our sponsor, Hewlett-Packard Enterprise (HPE), and the EPCC systems team for their support.

MSc programmes at EPCC
https://www.epcc.ed.ac.uk/education-and-training/masters-programmes

Celebrating five years at the Bayes Centre

Thursday 23rd August 2018 was the last day for EPCC in the James Clerk Maxwell Building (JCMB) at the King's Buildings in Edinburgh, where it had been located since its inception. The following Monday staff reported for work at the new Bayes Centre in the heart of Edinburgh. Mark Sawyer looks back at an eventful five years that have passed since that move.

The transfer from JCMB went very smoothly with just a single day of working at home before we were in place. As befits a fresh start, many staff members (including myself) took advantage of the opportunity for a major de-clutter. Nevertheless, a formidable array of purple crates was waiting for us as we settled in to the new surroundings.

Data Science, DDI, and EIDF

Moving to the Bayes Centre was more than a change of scenery; it marked EPCC's new role in the Data-Driven Innovation (DDI) initiative, a cluster of innovation hubs set up under the Edinburgh and South-East Scotland City Region Deal. The Bayes Centre innovation hub brings together around 600 world-leading applied data science researchers, talented students, and staff from organisations across the public, private and third sectors into one facility.

EPCC has a central role in the DDI, building and operating the Edinburgh International Data Facility (EIDF), which underpins the DDI's activities. EIDF comprises a set of high-powered computational and data services which will continue to grow and evolve in response to users' requirements.

The Covid-19 pandemic

In August 2018, as crates were being unpacked in the Bayes and shelves in new offices filled, few would have believed that 18 months later we would be moving out of the building again, this time to our own homes. In March 2020 it became apparent that the only hope of containing Covid-19 was to severely restrict human physical interaction.

EPCC responded to this challenge with remarkable resilience. A committee had been meeting regularly in the weeks leading up to the 'lockdown' to plan for this eventuality. Within a day of the Bayes Centre closing, EPCC had set up virtual meeting rooms to enable our regular project business to continue, with an online coffeeroom to encourage social interaction. Many staff already had experience of collaborating remotely with project partners, which enabled a smooth transition to the new way of working. With the exception of some essential staff working on-site at EPCC's Advanced Computing Facility (ACF), EPCC continued all its projects and services entirely online, also bringing in new staff, running training courses, and even holding a virtual Christmas party.

With the gradual lifting of Covid-19 restrictions, building occupancy at the Bayes Centre is approaching pre-pandemic levels. Hybrid working is the new normal at EPCC and is generally agreed to be a success.

HPC: as important as ever

EPCC's central role in the DDI does not mean that its interest in traditional HPC has declined. Quite the opposite is true. The period of the pandemic coincided with the launch of the ARCHER2 National HPC Service – the successor to the ARCHER service which delivered over six billion core-hours of compute. ARCHER2 has more than ten times the capability of its predecessor and reinforces our position as the UK's leading HPC centre. The service's installation and launch were carried out during the time of greatest restrictions on travel and working conditions – a remarkable achievement for those concerned.

The National HPC Service has since been accompanied by the arrival of smaller-scale systems aimed at exploring new

technologies and fostering early adoption. These include the Fulhame system, based on ARM technology, and a Cerebras CS-2 powered by the high-performance Wafer Scale Engine (WSE) processor combined with an extremely large-memory HPE Superdome Flex Server system for AI applications and massive data handling capability. Its predecessor at EPCC, the CS-1, was the first such system in Europe.

The next five years

The five years since we moved to Bayes Centre have seen many developments at EPCC that go far beyond a change of postcode. We have built on our core strengths to maintain our position as the UK's leading HPC centre and expanded our expertise into data science and data-driven innovation. We now operate Trusted Research Environments (sometimes called Safe Havens) for handling sensitive or controlled data, and the Edinburgh International Data Facility is demonstrating its value to research and business. With our goal of a UK Exascale service within reach, we look forward to exciting times ahead.

Mark Sawyer, EPCC
m.sawyer@epcc.ed.ac.uk

The Bayes Centre
The £45m Bayes Centre building, which became EPCC's new home, is the University of Edinburgh's innovation hub for data science and artificial intelligence. Officially opened in October 2018 by the University Chancellor, HRH The Princess Royal, it provides space for approximately 650 researchers, students, and external partners over five floors. Collaboration is built into its architecture, with open spaces designed to bring together its occupants from academia, industry, and the public sector. Approximately 100 members of EPCC staff occupy the second floor of the Centre, sharing it with our Master's and PhD students.

Images show: (top) the Bayes Centre building, where EPCC moved five years ago; (middle) packing up in JCMB; (lower) arriving at Bayes.

Forecasting: the future

The weather affects us all, and the climate affects not only us, but generations to come. For the highest quality predictions, on time scales from a few hours to a hundred years, simulation codes need to be both scientifically accurate and highly optimised.

Image: Ashley_Hampson via Getty Images

Last year the University of Edinburgh joined the Met Office Academic Partnership (MOAP). This consists of a group of universities, including UCL, Bristol, Exeter, Leeds, Oxford, Reading, and Birmingham working in close collaboration with the Met Office to tackle key challenges in weather and climate science and prediction. In EPCC, Michèle Weiland was appointed Met Office Joint Chair as part of this programme.

Next-generation code

The Met Office is currently developing its next generation of forecasting code, called LFRic after the visionary mathematician Lewis Fry Richardson. A century ago, and long before the invention of electronic computers, he came up with the fundamental concepts that underly modern numerical weather forecasting.

LFRic will replace the Unified Model, which has been the basis of both daily forecasting and climate science at the Met Office for the last 30 years. Weather forecasting is an unusual use case for high performance computing because it has a real-time deadline – it is no use computing a forecast for weather that has already happened!

To obtain a better forecast within the available time, it is necessary to use as much compute power as possible, and on today's supercomputers this means using hundreds of thousands of cores working in parallel. The design of the Unified Model has limitations which mean it cannot achieve the required scalability, but LFRic has been designed explicitly with modern architectures in mind.

Improving performance

EPCC is working with the Met Office to improve the performance and scalability of LFRic before it goes into production. We have ported the code to ARCHER2 and are conducting extensive performance tests find the best configurations and to fully understand where the bottlenecks are and how they might be overcome. The Met Office has recently installed its own latest supercomputer which, like ARCHER2, is an HPC Cray EX system, but with the latest AMD CPUs and an upgraded network.

We have already successfully implemented some optimisations, and more are planned in the coming months. Currently we are just working with the "dynamical core" of the model, called GungHo, which is the part of the model that solves the fluid dynamics equations for the atmosphere. In due course we will be working on the full LFRic model, which includes all the other physical processes that influence the weather, such as clouds, precipitation, convection, thermal radiation, and interactions with land and sea surfaces.

Further into the future, we will work towards running LFRic on upcoming processor designs, as well as providing specialist HPC training for Met Office technical staff, all as part of a long lasting and fruitful collaboration.

Mark Bull, EPCC
m.bull@epcc.ed.ac.uk

Met Office

MOAP https://www.metoffice.gov.uk/research/approach/collaboration/met-office-academic-partnerships

LFRic https://www.metoffice.gov.uk/research/approach/modelling-systems/lfric

GungHo https://www.metoffice.gov.uk/research/foundation/dynamics/next-generation

Lewis Fry Richardson https://en.wikipedia.org/wiki/Lewis_Fry_Richardson

Why are there not more women in HPC?

Some of the WHPC community celebrating Diversity Day at ISC23 in Germany in May.

Women in High Performance Computing (WHPC) is searching for the answer to this question to change the face of HPC.

What does WHPC do?

EPCC is a founding member of Women in High Performance Computing (WHPC), an international organisation created to bring together a network of women in HPC, and their allies, to provide fellowship, education and support, and advocate for a more inclusive community for all.

Through this network, WHPC is raising awareness about the under-representation of women in the HPC workforce and how this impacts our community, research outputs and progress. As well as raising awareness about a lack of diversity, WHPC also provides a platform to showcase the success of women in HPC. This includes promoting the visibility of women role models in HPC, spotlighting contributions and successes in the field and highlighting the positive results of diversity and inclusion initiatives.

The WHPC community is truly worldwide with an international network of Chapters and Affiliates that are officially associated with WHPC. These are community-level groups that serve a specific area or organisation to support and promote the WHPC vision locally. To date, there are WHPC Chapters in Europe, Asia, Australasia, and South and North America. Here at EPCC, we operate a chapter for the University of Edinburgh and our local community. The global community also gathers at international conferences to organise and attend WHPC workshops, networking receptions and conference-wide Diversity Days. These events provide a platform for discussion and networking, and ensure that WHPC has a prominent profile internationally to raise awareness and address issues around diversity and inclusion within the HPC community.

WHPC hopes to energise the community to engage with the cause and support women in the HPC community to succeed in their chosen career.

How can I get involved?

Become a member, follow us on social media, join us at events and join your local chapter or affiliate. Everyone is welcome to join WHPC! If you have an idea for an event that would benefit the community, or would like to work with us, let us know! Can't find a local chapter? Gather a group of like-minded individuals in your community and apply to start your own WHPC community.

WHPC
WOMEN IN HIGH PERFORMANCE COMPUTING

Eleanor Broadway, EPCC
e.broadway@epcc.ed.ac.uk

How is EPCC involved?

As a founding member of the organisation, EPCC continues to be prominent in the WHPC community today. Mark Parsons, Director of EPCC, sits on the Steering Committee and George Beckett, EPCC senior project manager, is the WHPC treasurer on the Executive Board. We also have staff who volunteer to organise WHPC workshops and on-site events at international conferences. For example, in 2023, Weronika Filinger was appointed as Overall Programme co-Chair and I became the Mentoring Chair on the WHPC at ISC'23 Organising Committee. I was also appointed Submissions co-Chair for the WHPC SC'23 Workshop.

Want to support us? We also offer a variety of partnerships, support and sponsorship opportunities for organisations. Find out more on our website: https://womeninhpc.org/

Reducing the cost of large-scale physical simulations

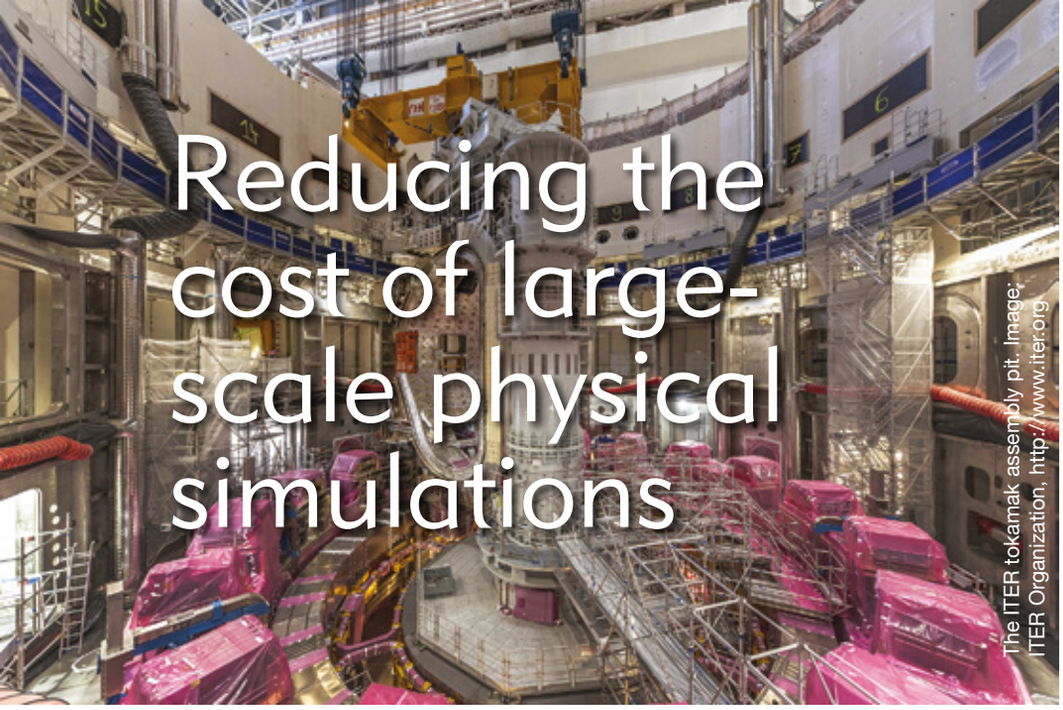

The ITER tokamak assembly pit. Image: ITER Organization, http://www.iter.org

The ExCALIBUR programme is supported by the UKRI Strategic Priorities Fund.

The programme is led by the Met Office and the Engineering and Physical Sciences Research Council (EPSRC) along with the Public Sector Research Establishment, the UK Atomic Energy Authority (UKAEA) and UK Research and Innovation (UKRI) research councils, including the Natural Environment Research Council (NERC), the Medical Research Council (MRC) and the Science and Technologies Facilities Council (STFC).

SiMLInt is being developed at the University of Edinburgh in collaboration between EPCC and The School of Mathematics, as part of the ExCALIBUR programme.

SiMLInt is a multi-disciplinary project that demonstrates how data-driven information can aid large-scale computational simulations of physical systems.

The behaviour of physical systems, such as fusion reactors, is often very complex, and in many cases the only way to model such systems with sufficient precision is through direct numerical simulation. This is computationally a very costly process as it involves evaluating the state of the system at a high number of discreet points across its spatial and temporal domain.

However, the general behaviour of such systems exhibits patterns that can be uncovered using machine learning (ML). It has been shown that such data-driven information can aid numerical computation and speed up the simulation significantly, however the two approaches use a very different set of computing technologies and it is not straightforward to create such a composite workflow.

Machine learning models

Here at EPCC we are developing SiMLInt, an interface that makes it easier for users to employ ML models in their simulations. Based on common technologies and off-the-shelf solutions, SiMLInt demonstrates how to bootstrap an existing physics simulation with machine learning without the need to develop bespoke software or rewrite parts of the workflow using different, and possibly less suitable, languages.

SiMLInt demo

Our released demo shows how to install and set up SmartSim, which facilitates the in-memory communication between the numerical simulation and the machine learning model, and also the pipeline that simulates turbulence in the magnetised plasma of a tokamak reactor.

Our demonstration numerical simulation is resolved on a coarser level of detail than is optimal, which significantly reduces the computation costs but also causes imprecisions and errors. We therefore introduce a machine learning model that has been trained using higher resolution simulation data and so can supplement the more detailed information, correcting any errors caused by the lower resolution. We are using BOUT++ to resolve the underlying numerical simulation, and Learned Correction and Large Eddy Simulation to supplement the machine learned information.

Our proposed SiMLInt infrastructure is robust, future-proof, and versatile, allowing domain specialists to test different data-driven approaches and different ways to feed the data-driven information to their numerical simulation.

Anna Roubíčková, EPCC
a.roubickova@epcc.ed.ac.uk

ExCALIBUR at EPCC https://excalibur.ac.uk/projects/simlint

BOUT https://boutproject.github.io

github https://github.com/EPCCed/SiMLInt

SmartSim https://github.com/CrayLabs/SmartSim

There are a number of open questions relating to the stability and feasibility of the resulting simulation, which will require collaboration with the wider community to tackle. If you are interested in joining the efforts or want to learn more, get in touch!

Preparing for quantum computing

Image: John D via Getty Images

It has been a very busy year for EPCC's growing Quantum Applications Group! Our research focus is on the interaction of quantum computing and high performance computing (HPC), applications of quantum computing, and classical simulation of quantum computing to support algorithm design and development.

In April 2023, the University of Edinburgh launched its Quantum Software Lab (QSL), which is affiliated with the National Quantum Computing Centre. I am a member of the Lab and co-investigator of its inaugural Quantum Advantage Pathfinder project.

QSL aims to develop a reproducible framework for identifying computational challenges relevant to industry, then designing, implementing, verifying, and benchmarking quantum solutions. It is investigating a number of use-cases, with EPCC involved in three of them.

Quantum Computing Application Cluster

We also engage with the Quantum Computing Application (QCA) Cluster, which brings together academics from across Scotland, bridging the divide between theorists and experimentalists.

In December 2023 I will join a QCA Cluster delegation to Q2B, the largest conference on quantum technologies, in California, to promote Scotland's unique quantum computing ecosystem.

Quantum research

Members of EPCC's Quantum Applications Group have been investigating how to increase the energy efficiency of large scale quantum computing simulations. This is particularly important when simulating 44 qubits on 4,096 nodes of ARCHER2! See our paper "Energy Efficiency of Quantum Statevector Simulation at Scale" (link at bottom of page).

We have also been working with LRZ on quantum task offloading with OpenMP, which will go a long way towards enabling traditional HPC to be coupled to quantum accelerators.

Future plans

Looking forward, we're particularly excited to begin investigating how new projects like Microsoft's Quantum Intermediate Representation can be leveraged to enable portable quantum programming, whether it's for real quantum hardware, or simulated backends.

We expect 2024 to be another busy year for EPCC's Quantum Applications Group!

Oliver Brown, EPCC
o.brown@epcc.ed.ac.uk

I lead EPCC's Quantum Applications Group (QAG), and in April 2023 I was awarded a five-year University of Edinburgh Chancellor's Fellowship. I have been using this new academic role to build EPCC's quantum computing research programme and expand our academic collaborations in the area.

Energy Efficiency of Quantum Statevector Simulation at Scale https://arxiv.org/abs/2308.07402

NQCC https://www.nqcc.ac.uk/quantum-software-lab/

QCA Cluster https://qca-cluster.org/

New Chancellor's Fellow joins EPCC

The University of Edinburgh's five-year Chancellor's Fellowships support interdisciplinary research and innovation. Joe O'Connor has joined EPCC as our second Chancellor's Fellow and he describes his work below.

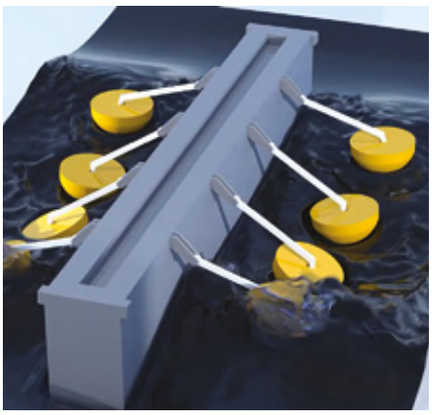

Image shows a smoothed particle hydrodynamics simulation of a wave energy converter.

The main focus of my Fellowship is to accelerate current simulation tools for marine energy systems (eg wave energy converters). Smoothed particle hydrodynamics (SPH) is the current state-of-the-art for simulating such applications, thanks to its meshless formulation which makes it ideally suited to violent free-surface flows (eg breaking waves) and wave-structure interactions. However, SPH is also more costly than traditional mesh-based computational fluid dynamics (CFD) methods. Because of this, SPH has found limited use in typical engineering workflows (eg design optimisation, uncertainty quantification, optimal control). Therefore, there is a need to accelerate SPH.

One approach to acceleration is low-fidelity surrogate models that approximate the high-fidelity SPH but are much cheaper to compute. The recent boom in machine learning has led to a string of developments in data-driven surrogate modelling. However their application to SPH is essentially non-existent. Another approach to acceleration is high performance computing (HPC). For example, low-cost surrogate models for SPH will enable new areas of research in novel parallel computing strategies (eg parallel-in-time). These mixed-fidelity methods have the potential to combine the efficiency of low-fidelity surrogate models with the accuracy and robustness of high-fidelity SPH.

While the focus of my Fellowship is method development and acceleration, the ultimate aim is to apply these methods to real-world problems. Marine energy is an ideal application where this research can deliver significant impact.

A successful Fellowship will enable larger, longer, and more detailed simulations of marine energy applications, as well as unlocking previously intractable engineering workflows. Ultimately, this will improve design capability and enhance progress towards digital twinning technology. My research relates to many of the areas targeted for the transition to exascale and will also contribute towards the UK's Net Zero Strategy.

EPCC is the perfect place for this investigation, with world-leading expertise in HPC, data science, and computational modelling, as well as significant experience of leading national projects relevant to my work. This Chancellor's Fellowship presents a unique opportunity to combine my own skills and experience with world-leading expertise in EPCC to help address an important global challenge problem. And, as one of the developers of the leading open-source SPH code in the community (DualSPHysics), I will be aiming to make the outputs of this research available as open source to maximise its impact.

Joe O'Connor, EPCC
j.oconnor@epcc.ed.ac.uk

THE UNIVERSITY *of* EDINBURGH

As a Chancellor's Fellow at EPCC, I will be working towards accelerating marine energy simulations through a combination of machine learning and novel parallel computing. My research will reduce the time and cost associated with performing high-fidelity computational fluid dynamics simulations of wave energy converters.

Research at EPCC

www.epcc.ed.ac.uk/research

Exploring greener weather forecasting

Image: deepblue4you via Getty Images

EPCC's Nick Brown has been awarded a Personal Research Fellowship from the Royal Society of Edinburgh to explore greener supercomputing for weather and climate simulations. Here he explains the focus of his research.

My research will comprise two main strands, with in-situ HPC data analytics the first area of focus. Large-scale simulation codes, such as models from the Met Office, generate vast amounts of raw data that must then be post-processed into information that can be consumed by users.

The state of the art is to leverage existing CPU nodes of the supercomputer for data processing, however this wastes resource that could be otherwise used for compute and, furthermore, requires significant energy. My hypothesis is that Field Programmable Gate Arrays (FPGAs) would be a much better match and I think they can deliver a step change in capability. This is because by combining the very high bandwidth of FPGAs with their inherent energy efficiency, we can put them in the data plane (the part of a network responsible for the movement of data) if they reside at the network level. Consequently, FPGAs could undertake post-processing of vast amounts of raw data in real time as it is streamed out from simulation codes.

My other area of investigation will be the computational acceleration of the Met Office's CASIM microphysics model on AMD Xilinx AI Engines (AIEs). These are provided in arrays of 400 engines, each running at 1.2GHz and capable of undertaking 8 floating point operations per cycle – again at very low energy. I undertook the first study of AI engines for HPC workloads with a Met Office atmospheric advection kernel. This work with CASIM builds on those lessons learned.

While the AIEs deliver significant raw compute power at low energy usage, the key challenge is how they might actually be leveraged in a code! Consequently, using the CASIM microphysics model as a vehicle to explore the potential advantages of this new technology will not only help develop performance best practice, but I hope it will also deliver benefits to this production code as the demands on it, driven by scientific ambition, continue to increase.

I am looking forward to starting my Personal Research Fellowship in 2024. It has in part been made possible by EPCC's involvement in the Met Office Academic Partnership (MOAP) programme. Not only has MOAP opened up access to Met Office simulation codes, enabling me to talk to Met Office staff and understand how innovative research might help tackle some of their challenges, but the MOAP collaborative programme means that I will be able to ensure that the outcomes of my RSE Personal Research Fellowship can be fed back to them.

RSE The Royal Society of Edinburgh
KNOWLEDGE MADE USEFUL

Nick Brown, EPCC
n.brown@epcc.ed.ac.uk

Research at EPCC https://www.epcc.ed.ac.uk/research

Exploring the Versal AI engines for accelerating stencil-based atmospheric advection simulation
https://arxiv.org/pdf/2301.13016.pdf

Decarbonising our supercomputers is crucial to ensure we meet our climate goals and because drawing less energy means these systems will be cheaper and easier to install and operate. The question of how novel, more energy efficient, hardware can be exploited is of paramount interest and will be the focus of my research.

Services for industry: AI-ready architectures on demand

Images show the two Cerebras CS-2 systems at EPCC's Advanced Computing Facility.

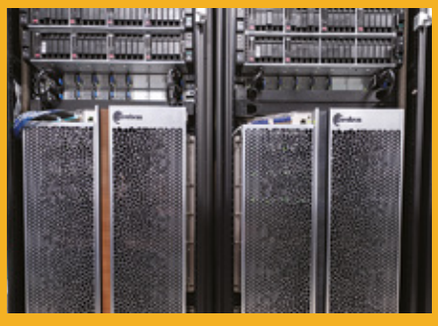

Julien Sindt, EPCC
j.sindt@epcc.ed.ac.uk

Machine learning models have the potential to revolutionalise decision making and problem solving across industry. Here at EPCC, we can provide access to an impressive collection of AI-ready architectures, bringing the benefits of machine learning to businesses of all sizes.

EPCC continually invests in state-of-the-art computing technology. We host and operate experimental systems such as the RISC-V and FPGA testbeds which allow developers to test whether novel architectures can significantly improve the performance of their workflows. In doing so, we also maintain a good understanding of upcoming trends in computing technologies and their potential applications.

The rise in prominence of machine learning has directed EPCC's recent explorations in cutting-edge computing technologies. Training is usually the most compute-intensive part of creating a machine learning model, but the ease with which training can be split across multiple cores has made this task an ideal fit for GPU architectures.

But as machine learning models have become more complex, the number of GPU-hours required to train such models has increased drastically. For example OpenAI's ChatGPT, one of the largest models to be trained, is estimated to require 34 days even using 1,023 top-of-the-line Tesla A100 GPUs. This incredible resource requirement can make machine learning seem out of reach for many smaller companies.

However, novel computing architectures designed specifically to improve the training time of very large models have recently started to become available. EPCC has procured two such AI-ready systems, the Cerebras CS-2 and the Graphcore Pod64, and we have been testing their efficacy for training large machine learning models.

The Cerebras system is incredibly good at training large language models (similar to those used in ChatGPT). We have found that training a BERT large-language model with 100 million parameters on the Cerebras CS-2 is 36 times faster than training it on a Tesla V100 GPU, and expect this speed-up to increase when the CS-2 configuration is upgraded. We expect a similar improvement in performance from the Graphcore Pod64.

These novel technologies make it easier to replicate the most impressive machine learning feats in a more reasonable time and at a more reasonable cost. We look forward to exploring their exciting possibilities with our business partners!

EPCC is one of the few places in the world where it is possible to compare the performance of training your machine learning model on a Tesla V100 GPU cluster, a Tesla A100 GPU cluster, a Graphcore Pod64, or a Cerebras CS-2. Such flexibility means that decisions about the best architecture can be made early in a project, ensuring that our partners get the best computational bang for their buck.

If you would like to try your applications on any of the systems mentioned here, or if you would like to hear about the other exciting systems hosted at EPCC's Advanced Computing Facility, please contact me by emailing commercial@epcc.ed.ac.uk.

Industry services at EPCC https://www.epcc.ed.ac.uk/industry-solutions

Accelerated opportunities for the research community

The Edinburgh International Data Facility (EIDF) continues to develop its resources to support data-driven discovery. One area that has progressed significantly is the formal launch of the EIDF GPU Service.

The EIDF GPU Service is made up of HPE Apollo 6500 GPU servers containing NVIDIA A100 GPUs. The service currently has a total of 160 GPUs, of which 112 are available for general EIDF users.

Individual projects can have access to up to 12 GPUs. The service is accessed through a Virtual Machine (VM) set up for each project within the EIDF and is operated via Kubernetes.

Using Kubernetes, users can submit work directly to one or more GPUs or make use of a smaller section of a GPU for any individual job. Work using full GPUs can be run on pods of up to eight GPUs per job. Each GPU allocated to users will, by default, have circa 100GB of memory and eight CPU cores associated with it.

Sub-GPU-scale work uses NVIDIA Multi-Instance GPU (MIG) technology and provides for multiple users, or multiple jobs from an individual user, to run work on a single GPU with complete isolation between different jobs. MIG can partition runs of up to seven jobs per GPU, dependent on the target workload. These jobs will by default have memory and CPU scaled to the amount of GPU used.

Advanced cooling system

The GPU Service is housed in Hewlett Packard Enterprise's (HPE) "Adaptive Rack Cooling System" (ARCS) racks, which means that rather than having to manage cooling of an entire room, the racks are cooled themselves, ensuring that the cooling of the GPU servers is as efficient as possible. This reduces the overall energy required to operate the service, for users. (See ACF article on p14 for more information.)

Early use of the Service delivers

EPCC has been working with the University of Edinburgh's School of Informatics to develop the EIDF GPU Service and this has already resulted in research being undertaken on it.

Two papers have recently been published, based on work performed on the new service:

- Gema et al (2023) 'Parameter-Efficient Fine-Tuning of LLaMA for the Clinical Domain'.
- Kaddour et al (2023) 'No Train No Gain: Revisiting Efficient Training Algorithms for Transformer-based Language Models'.

EPCC is committed to further developing the EIDF GPU Service and we are looking at the next generation of NVIDIA H100 GPUs to enhance this resource in 2023–2024.

Ritchie Somerville, EPCC
r.somerville@epcc.ed.ac.uk

Image: HPE ARCS racks

NVIDIA A100 GPU. Image NVIDIA 2023.

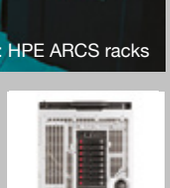

HPE Apollo 6500 Gen10 Plus system. Image HPE 2023.

"The EIDF GPU Service is enabling us to explore new methods for training more explainable, robust, and trustworthy AI systems; to design and experiment with models that can learn to search for the information they need for solving arbitrary knowledge-intensive tasks; and to design statistical models for solving challenging biomedical and clinical problems."

Pasquale Minervini, Lecturer in Natural Language Processing, School of Informatics, University of Edinburgh

EIDF GPU Service https://www.ed.ac.uk/edinburgh-international-data-facility/services/computing/gpu-service

Parameter-Efficient Fine-Tuning of LLaMA for the Clinical Domain: https://arxiv.org/abs/2307.03042

No Train No Gain: Revisiting Efficient Training Algorithms For Transformer-based Language Models: https://arxiv.org/abs/2307.06440

Constant motion at the ACF

The two 6MW plate heat exchangers used to control the temperature of the computer cooling water used in Computer Room 4. Image by Keith Hunter Photography.

EPCC's Advanced Computing Facility (ACF) never stands still. From small one-off changes to major upgrades, we constantly look to improve in all areas. All of which keep us busy and excited about the future – whatever it might be.

Recently we have been working to improve the efficacy of the ACF's Mechanical and Electrical (M&E) infrastructure, which is the backbone that supports the services we operate and maintain.

We have many types of equipment here at the ACF, from directly water-cooled to more standard air-cooled equipment housed in open cabinets. The equipment in these open cabinets is cooled by the "traditional" method of drawing cool air from the front of the cabinet and through the kit to cool it via internal fans. The expelled hot air at the rear of the cabinet is then recirculated as part of the room air, re-cooled by large air-conditioning units and sent back to the cabinets via a sub-floor plenum to continue the cooling circuit. Simple!

This has worked well for many years. However, the efficacy of these systems would be improved if we can contain the cool air locally and recirculate it around the equipment. This is where HPE ARCS (Adaptive Rack Cooling System) cells come in.

These cells generally comprise five cabinets, with the centre cabinet being a sealed cooling/heat-exchanger unit which is directly connected to a chilled water supply and able to cool and circulate conditioned air to two sealed computer cabinets on either side of it. Hence, only conditioning the local air, not the whole room.

Only the air within the cell is cooled, unlike in a more traditional data hall which cools and conditions that air for the whole data room. Each cell also has its own smart controls, which work to condition the air temperature and flow to suit the needs of that particular cell. This local control is obviously far more efficient and targeted than can be achieved when cooling a whole computer room.

The use of ARCS cells is not just a more efficient use of chilled water and electrical energy, but also reduces noise levels in the space, which is an added benefit to those working in the room.

Computer Room 1 upgrade

One of our main projects is a planned upgrade of our Computer Room 1 (cr1). The central aim of this work is to maximise use of available floor space, introduce more UPS-backed power and, more importantly, improve the efficacy of the space from an energy use point of view. A large part of this being the introduction and installation of several ARCS cells which will maximise the use of available floor space and mechanical and electrical infrastructure.

Unfortunately, it's not just a case of rolling the cabinets into the room and filling them with equipment.

Images: (top) HPE ARCS cabinets at EPCC's Advanced Computing Facility; (lower) some of the Dry Air Cooler plant on the ACF roof which supports both the mechanical and the free cooling for Computer Rooms 2 and 3.

Before this, a detailed M&E design must be prepared to suit the installation requirements and ensure the plans are compatible with the site's existing M&E capabilities.

This is where we brought in the Project Management service of the University of Edinburgh's Estates department to evaluate the requirements of the project and drive it forward. The Projects service appointed an M&E design consultancy with a well-proven track record in data centre design (and who were also the main designers of our Computer Room 4) to provide a full feasibility study and design of the requirements.

This design phase was then followed by a lengthy period to install the Electrical and Mechanical infrastructure designed by the consultancy to support the cells.

In addition to these M&E works in Computer Room 1, there has been a need for extensive room reconfiguration works which included the removal of redundant Computer Room Air Conditioner (CRAC) units, re-siting of the gaseous suppression system, and other ancillary works to provide the physical space for the ARCS cell installation. All of which is time consuming and hard work for all concerned, especially when we need to ensure there is no disruption to the local live production services in the room!

These M&E and reconfiguration works have now been completed as planned and the installation, commissioning and racking of the new ARCS cabinets will commence in late 2023/early 2024.

Calum Muir, EPCC
c.muir@epcc.ed.ac.uk

Advanced Computing Facility
https://www.epcc.ed.ac.uk/hpc-services/advanced-computing-facility

Developing the Scottish Medical Imaging (SMI) Service

Image: Portra via Getty Images

The PICTURES project curates a large, linkable, research-ready dataset of "real-world" medical images from the Scottish population.

Millions of CT, MRI, X-ray, and other image types are taken every year across Scotland as part of routine medical care and stored in a national database. The Scottish Medical Imaging Service (SMI) has permission to hold a copy of this database, covering data from 2010-17 totalling around 1.7 petabytes.

In addition to the images themselves, each file in the database holds a rich set of metadata relating to the patient, imaging device, and often a text-based report describing the radiologist's findings. Using the Scottish Community Health Index (CHI), images can also be linked to other health and administrative data such as GP records.

The dataset has the potential to support a broad range of research including early-prediction tools, population models, and longitudinal studies on patient outcomes. This has been the goal of the SMI project since its inception in 2017 and, through the outstanding efforts of the PICTURES team, we were able to fully launch the service in April 2022.

The challenges associated with managing such messy, real-world data cannot be overstated. For example, even determining the body part which has been imaged has proved difficult! To curate the dataset and enable clinical research, we have applied novel cleaning techniques including natural language processing (NLP) to de-identify radiology reports, character detection models to remove text "burned into" images, machine learning to classify image types, and build a searchable catalogue of metadata covering the whole dataset. Through this work, we have developed and honed our data "wrangling" skills, which are now being applied to other projects at EPCC.

Since opening the service, we have been supporting a handful of exemplar studies. One example is using brain images and linked reports to build a risk model for future health conditions such as dementia and stroke. Recruiting enough patients for this, and monitoring their ongoing medical care, would have been impossible to do at scale without using the dataset curated by SMI.

Looking to the future, we will soon add X-rays to our service catalogue, and have made progress in de-identifying mammography images. Governance arrangements are being finalised to allow the database to be brought – and kept – up to date. Early next year we will migrate to EPCC's new Trusted Research Environment (TRE), which will provide HPC and GPU resources to our researchers.

We are also pursuing funding to further develop the service and add value to the community, such as validating AI algorithms on representative population data before they are deployed in the healthcare setting.

Ruairidh MacLeod, EPCC
r.macleod@epcc.ed.ac.uk

PICTURES

PICTURES is a Medical Research Council funded collaboration between the University of Edinburgh, University of Dundee, Abertay University and NHS Scotland.

"EPCC staff have worked tirelessly to rise to the many technical challenges and governance constraints in collaboration with colleagues in PICTURES and ourselves in Public Health Scotland."

Jackie Caldwell, SMI Strategic Lead in Public Health Scotland

PICTURES public website:

https://www.imageonamission.ac.uk

The International HPC Summer School

The week-long International HPC Summer School 2023 (IHPCSS'23) was held in Georgia, USA, in July and brought together over 80 students and 30 instructors. For the first time in the event's 13 years, EPCC participated as an independent partner rather than European project member, sponsoring the attendance of five students and three staff members.

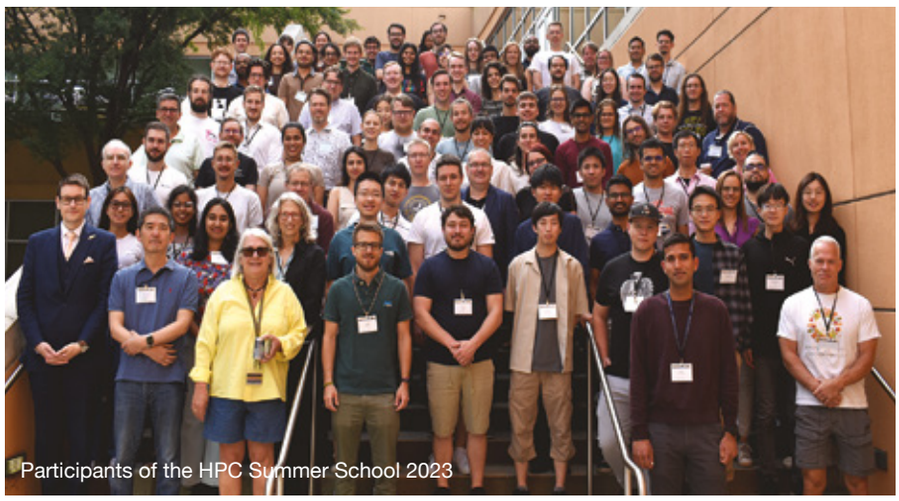

Participants of the HPC Summer School 2023

Technical content

The Summer School is a mixture of lectures, hands-on sessions, science talks, and extensive mentoring programme focused on career development, interpersonal skills, and networking. It mostly targets PhD students, but also accepts MSc students and early-career postdocs who are either already using HPC or have immediate need for it in their research. Throughout the week students are exposed to concepts and techniques used in HPC, artificial intelligence, machine learning, and data science. They hear from computational scientists across several domains about how HPC enables research breakthroughs, and then learn about the tools and practices that will enable them to do the same in their own research.

Mentoring programme

One of the most valuable parts of the IHPCSS is the mentoring programme, which permeates the whole event. Students are encouraged to think about career and personal development goals during themed sessions, group and one-to-one mentoring.

Communicating research

Students have the opportunity to present their research during the poster session. As HPC may be the only common thread between attendees, they are forced to find a new way to explain their work, and I'm pleased to say the UK students succeeded brilliantly.

Programming challenge

An optional component of the Summer School was the programming challenge run by EPCC Teaching Fellow, Ludovic Capelli. The challenge utilised the OpenMP and MPI parallel programming skills that were taught at the event to optimise a parallel PageRank code. Its goal was to complete as many iterations of the algorithm as possible within ten seconds, with the starting point fewer than 100 iterations. The winning code ran an average of 330,000 iterations, an incredible achievement! The winning team comprised EPCC MSc student Katie Worton, EPCC PhD student Shrey Bhardwaj, and Isaac Alonso from The Instituto de Astrofísica de Canarias.

IHPCSS'23 was another successful and fun event. Planning for IHPCSS'24 has already started, with student applications opening towards the end of 2023. Perhaps we'll see you there?

Weronika Filinger, EPCC
w.filinger@epcc.ed.ac.uk

Comments from participants

"I had a wonderful and enriching experience at IHPCSS'23. I met experts in HPC and researchers whose work I had read about. The coding challenge was very useful in bringing together all the techniques learnt from OpenMP and MPI sessions." ***Shrey Bhardwaj*** *PhD student at EPCC*

"Having a broad group of scientists from various disciplines and geographies provided many avenues for discussions. I recommend IHPCSS to anyone looking to increase their exposure to HPC." ***Harry McHugh*** *Online MSc in HPC with Data Science*

"The summer school was an excellent balance of technical development and networking. It was an amazing opportunity to make connections and work with others to deepen my technical knowledge." ***Katie Worton*** *Online MSc in HPC with Data Science*

IHPCSS'23
https://ss23.ihpcss.org/agenda/

Updates from the Software Sustainability Institute

From the Hidden REF to the publication of our Midterm Report, the SSI continues its mission to cultivate better software for better research by organising key events, policy initiatives, and training courses related to software in research throughout the UK. Here's what we've been up to recently.

SSI Midterm Report

We've launched a new publication that draws together some of the achievements of the Software Sustainability Institute since its foundation in 2012. The document comprises 12 case studies that demonstrate the critical role that the SSI has played in improving research culture, revolutionising access to software training and working with our collaborators to develop policies that better recognise and support the vital role of software in research.

https://www.software.ac.uk/blog/2023-08-16-ssi-midterm-review-presenting-ssis-achievements-and-impact-last-13-years

Research Software Camps

The FAIR (Findable, Accessible, Interoperable and Reproducible) principles are a concept that originated in data management to promote good practices and have recently been applied to research software. Some of the principles have also been adapted and extended, due to the nature of research software as a digital object, and are revolutionising the way we approach research software, enhancing transparency, reproducibility, and the sharing of knowledge.

"Brilliant content, engaging presenters, really useful topics explained in an accessible way without overwhelming jargon, and a well thought out way to present it with... interactive elements and plenty of resources for further reading." Research Software Camp attendee

The SSI's latest Research Software Camp (RSC) on FAIR Software took place from the 19–30 June 2023. We hosted talks and workshops such as a panel discussion focusing on key aspects of FAIR research software principles.

After attending an RSC, we offer a mentorship programme to provide one-to-one support with the help of volunteer mentors to further develop research software skills. The Mentorship Programme ran for eight weeks and saw a total of 20 mentees graduate!

https://www.software.ac.uk/blog/2023-07-27-highlights-research-software-camp-fair-software

"This [mentorship] programme was a valuable learning experience, and has encouraged me to continue to learn more about the field of Research Software Engineering, especially within the area of Psychology."
Research Software Camp mentee

Collaborations Workshop discussions

From building and sustaining healthy communities to tracking the environmental impact of research computing, the Collaborations Workshop 2023 (CW23) blog posts address the most up-to-date discussions in the research software sphere. The SSI's annual Collaborations Workshop brings together researchers, software developers, managers, funders and more to explore important ideas in software and research and to plant the seeds of interdisciplinary collaborations. These speed blog posts are written during CW23 and published on the SSI blog:

https://www.software.ac.uk/tags/cw23-speed-blog-posts

UNIVERSE-HPC

The UNIVERSE-HPC project will create a comprehensive collection of training materials to train the next generation of research software engineers (RSEs), with the aim of increasing both the skill and diversity of people working in the sector.

In August, the UNIVERSE-HPC project successfully piloted a new online training course at the University of Southampton for 20 researchers, aimed at those looking to expand their intermediate software development skills:

http://www.universe-hpc.ac.uk

The Software Sustainability Institute's Collaborations Workshop 2023 brought together researchers, developers, innovators, managers, funders, publishers, policy makers, leaders and educators to explore the theme of Sustainable Career Development for those in the research software community.

Hidden REF

The Hidden REF campaign works to recognise all research outputs and every role that makes research possible. The Hidden REF Festival, which took place in September, provides a focus for its activities. Simon Hettrick, SSI Deputy Director, and Kirsty Pringle, SSI Project Manager, have both been involved in the campaign's committee since the start. The committee recently published an article in Times Higher Education which highlights why the movement is pushing for change and how it's already made a difference.

https://software.ac.uk/news/new-hidden-ref-article-times-higher-education

Hidden REF website
https://hidden-ref.org/

SOFTWARE SUSTAINABILITY INSTITUTE

About us

The Software Sustainability Institute helps people build better software and more sustainable research software to enable world-class research.

Selina Aragon, Software Sustainability Institute
s.aragon@epcc.ed.ac.uk

"The vast majority of current research would not be possible without a group of people who provide a diverse range of vital skills – both those that are REFable and those that are hidden by the current evaluation regime."

The Hidden REF committee

Software Sustainability Institute website:
https://www.software.ac.uk

Farewell to the MSc class of 2023

In August we held a celebration for our MSc students and staff in recognition of the end of the 2022-23 MSc programme at EPCC.

It was a pleasure to see so many students enjoying the end-of-year reception, which marked the submission of their dissertations after more than three months' hard work. The wide range of projects undertaken across high performance computing and data science, some including industrial placements, reflects the breadth of staff expertise and computing facilities available to our students.

This is the first full MSc year since the Covid pandemic that has not been subject to some kind of restriction, and we have welcomed the return to normal interactions with students. We're very proud of our high staff to student ratio for dissertations which allows supervisors to get to know students individually over the eight months of their project, from initial topic proposal just after Christmas to submitting the finished work in summer.

One of our goals when we moved to the Bayes Centre five years ago was to make full use of the building for the benefit of our students. The second floor, which is occupied by EPCC, includes an embedded study space designed to encourage interaction and enable students to freely consult EPCC staff. Also, as the photographs here show, the ground floor atrium provides an excellent space for gatherings.

Support for students

This was the first year of a new model of student support across the University of Edinburgh. Here at EPCC we are fortunate to have excellent colleagues supporting us in Jemma Auns and James Richards. Jemma especially in her role as Student Adviser embodies our ideals in building quality relationships with our students, using every interaction as an opportunity to gain a better picture of what is going on with a student.

We have also expanded our academic team with the addition of Ludovic Capelli, our first dedicated Teaching Fellow. Ludovic is himself an EPCC MSc graduate, which enhances the insight he brings to his role of providing technical and academic support. As the new Academic Cohort Lead, Ludovic will focus on ensuring our students are given high quality, targeted support throughout the academic year.

ISC 2023 Student Cluster Competition

We congratulate the MSc students of TeamEPCC, who won the ISC23 Student Cluster Competition: Hristo Belchev, Ikraduya Edian, Jaffery Irudayasamy, Oleksandr Piekhota, and Tomas Rubio Cruz. Their ability to collaborate with their coaches (Spyro Nita and Xu Guo), the student support team, and each other enabled them to rise to every challenge and bring home the ISC23 Student Cluster Competition trophy (see page 3).

Online Learning graduates

The students who will form our first graduating cohort from the Online Learning part-time intermittent MSc programmes must also be commended. We were especially pleased that two were funded to attend the International HPC Summer School, where one was part of the winning team in the Programming Challenge (see p17).

"My experience at EPCC was incredibly reassuring. The interaction with the EPCC teaching staff was abundant, both inside and outside the classroom. The administrative office was always accessible, and the faculty offices surrounded our workspace, making it convenient to seek immediate assistance. What impressed me most was the promptness with which teachers responded to queries, always addressing them efficiently and wholeheartedly. I always felt like I was in good hands; the faculty and staff at EPCC stood with me like a family, ready to support me in any challenges I faced. They are truly deserving of a student's trust."

Weiyu Tu, MSc in High Performance Computing, 2022–23

This year's cohort of trailblazers have provided us with immensely useful feedback to develop our offering for the future. We hope to have further reflections on their journey in the next edition of EPCC News.

Graduations

The successes of the class of 2023 stand testament to the hard work of our students, but also their ability to effectively communicate and collaborate with EPCC staff and classmates. Some students have already started the next stage of their journey with internships, full-time positions, and PhDs, and we wish all our students the very best for their future careers, wherever they may lead them. At the time of writing we are looking forward to the graduations in November when we will host another celebration for students and colleagues.

Masters programmes at EPCC

EPCC offers a suite of two linked Masters programmes. Our programmes in High Performance Computing, and High Performance Computing with Data Science are available both on-campus (full-time and part-time) and online (part-time intermittent).

Both programmes have a strong practical focus, with most courses including lectures supported by practical sessions. Some courses are also available online as stand-alone postgraduate professional development courses. Our MSc in High Performance Computing will provide you with a thorough grounding in HPC technologies and their practical application. It will appeal if you have a keen interest in programming, computer science and would like to learn about HPC and parallel programming.

Ben Morse, EPCC
b.morse@epcc.ed.ac.uk

Looking ahead

No sooner have we said goodbye to one cohort than the next was upon us. Fifty new students joined us in September and we look forward to updating you on their progress.

Masters programmes at EPCC:
https://www.epcc.ed.ac.uk/education-and-training/masters-programmes

The Big Bang Fair 2023: the bigger bang strikes back

In June the EPCC Outreach team represented the ARCHER2 UK National Supercomputing Service at the annual Big Bang Fair, one of the UK's largest science and technology outreach events.

The EPCC Outreach team took a selection of activities to this year's Big Bang Fair (BBF23), all designed to enthuse and inform participants about high-performance and scientific computing.

Ball-sorting

This ever-popular activity demonstrates parallelism through participants sorting different colours of ball by throwing them into buckets. We can then double, triple, and so on the number of participants and, over the course of a day, plot a graph to demonstrate the scalability.

This event draws in those who may be less interested in a computer-based activity, and opens up a conversation about parallelism and scalability. Having discussed these concepts, we can then direct people to our next activity.

Wee Archie

Wee Archie, a cluster of 18 Raspberry Pis, was used to run our Aerofoil demonstration where participants design an aircraft wing by entering parameters into the simulation and submitting the job to Wee Archie.

Once the job is running on the cluster, the first four columns of LEDs on each Raspberry Pi demonstrate computation activity on each of its four cores (helping to explain shared memory programming), while the fifth and sixth columns demonstrate communication in and out of each node/Pi (helping to explain distributed memory programming). The final columns demonstrate memory usage and temperature.

Simulations take only thirty seconds each to run on eight Raspberry Pis/nodes in the two-laptop set-up. Animations then demonstrate to participants whether their plane could take off, and if so how far it could fly.

Hardware and logic puzzles

We also took part of a blade from ARCHER2 and from a Cray XT4, displaying them alongside actual-size images of the front and back of an ARCHER2 cabinet to demonstrate the immense scale of the full-size machine.

Finally, we encouraged our visitors to solve classic logic problems, which are a great way to get participants to think about how to programme such complex problems and to consider the software aspect of HPC.

BBF23's new format, with school groups visiting for a half day only, encouraged more considered and in-depth conversations. All in all it was a very successful and enjoyable experience for everyone involved!

Ben Morse, EPCC
b.morse@epcc.ed.ac.uk

The EPCC team at BBF comprised myself, Jemma Auns, Oliver Brown, Sam Haynes, Kara Moraw, Spyro Nita, Luca Parisi, James Richards, Darren White, and Chris Wood. We represented ARCHER2, the UK National Supercomputing Service.

Organised by Engineering UK, The Big Bang UK Young Scientists & Engineers Fair was held at the NEC, Birmingham, in June.

The Big Bang Fair
https://www.thebigbang.org.uk

Discover and Learn, our website for Outreach
https://discover.epcc.ed.ac.uk

Nuffield Research Placements at EPCC

This summer EPCC supervised three student visitors as part of the Nuffield Research Placements (NRP) programme, which provides hands-on research projects for 16 and 17-year-old school students.

Three secondary-school students joined us for a three-week research visit in August, each investigating a separate topic proposed by EPCC. They were supervised by EPCC technical staff, principally Joseph Lee, Lorna Smith, and James Richings. During their time with us, the students were also given a tour of our Advanced Computing Facility and the facilities we host there.

Following their research visits, each student produced a report and an academic-style poster to be showcased at an event organised by the Nuffield Foundation.

Student projects

Machine Learning: supervised by James Richings.

The project was designed to enhance the student's understanding of the Python programming language and the basics of machine learning, starting with training a model for the binary classification of images according to the type of object depicted. The key challenges were to understand the principles of neural networks and how to tune the parameters to train the system efficiently.

Navier Stokes Solver: supervised by Mario Antonioletti, Luca Parisi and Lorna Smith.

This project focused on the code employed to demonstrate the principles of high-performance computing at outreach events. It uses a simple Navier Stokes solver to simulate the amount of lift generated by modifying three parameters for the wing of an aeroplane, and has been adapted to use GPUs. The project tested the code's performance across a number of different GPU and CPU systems.

Benchmarking A Quantum Chemistry Library: supervised by Joseph Lee.

The focus of this project is the benchmarking of a quantum chemistry library and a computational fluid dynamics library on Cirrus, the Tier-2 supercomputer hosted by EPCC.

These codes are heavily used by scientists to accurately simulate different physical systems, and are computationally intensive. Through this research project, the student learned about the various components of supercomputing, including hardware architecture and software development practices.

"I thoroughly enjoyed the placement in Edinburgh, it gave me great insight into multiple careers and the EPCC staff made the whole process enjoyable; they were extremely warm and welcoming. I developed numerous computing skills and the support was fantastic."

Nuffield student visitor

Nuffield Research Placements

Chris Wood, EPCC
c.wood@epcc.ed.ac.uk

The Nuffield Foundation
https://www.nuffieldfoundation.org

Nuffield Research Placements
https://www.nuffieldresearchplacements.org

Master's degrees in High Performance Computing (HPC) and in HPC with Data Science

**EPCC is the UK's leading supercomputing centre.
We are a major provider of HPC training in Europe, and have
an international reputation for excellence in HPC education and research.**

Our MSc programmes in High Performance Computing (HPC) and HPC with Data Science have a strong practical focus and provide access to leading edge systems such as ARCHER2 (the UK's National HPC Service), and Cirrus (an EPSRC Tier-2 National HPC facility including over 150 GPUs).

MSc students have the opportunity to undertake their dissertations as an industrial project, building on EPCC's strong business links. Recent project partners range from start-ups to multinationals.

"The year I spent pursuing my MSc degree at EPCC was an enriching chapter in my life. The curriculum is designed to combine theory and practice. After attending the interesting lectures, the challenge of implementing the programs by myself stimulated my creativity, and I really enjoyed this learning experience."

Hanhui Chen
MSc in HPC student, 2022-23

Programmes can be undertaken on-campus (full-time and part-time) and online (part-time intermittent). Optional course choices include modules from the School of Informatics and the wider College of Science and Engineering.

Our graduates are in high demand in both academia and industry in the UK and abroad.

The University of Edinburgh is ranked in the top 30 universities in the world by Times Higher Education World University Rankings 2023, and 15th by QS World University Rankings 2023.

"I was initially attracted to the course as it allowed me to continue working full time. It provided relevant, detailed content that is well executed, with staff constantly willing to receive feedback on how best to deliver a programme online to part-time students."

Harry McHugh
MSc in HPC with Data Science student, 2020-23

www.epcc.ed.ac.uk/msc