Issue 95 SUMMER 2024

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

epcc news

From Edinburgh to the world

Unmatched expertise in supercomputing and data science

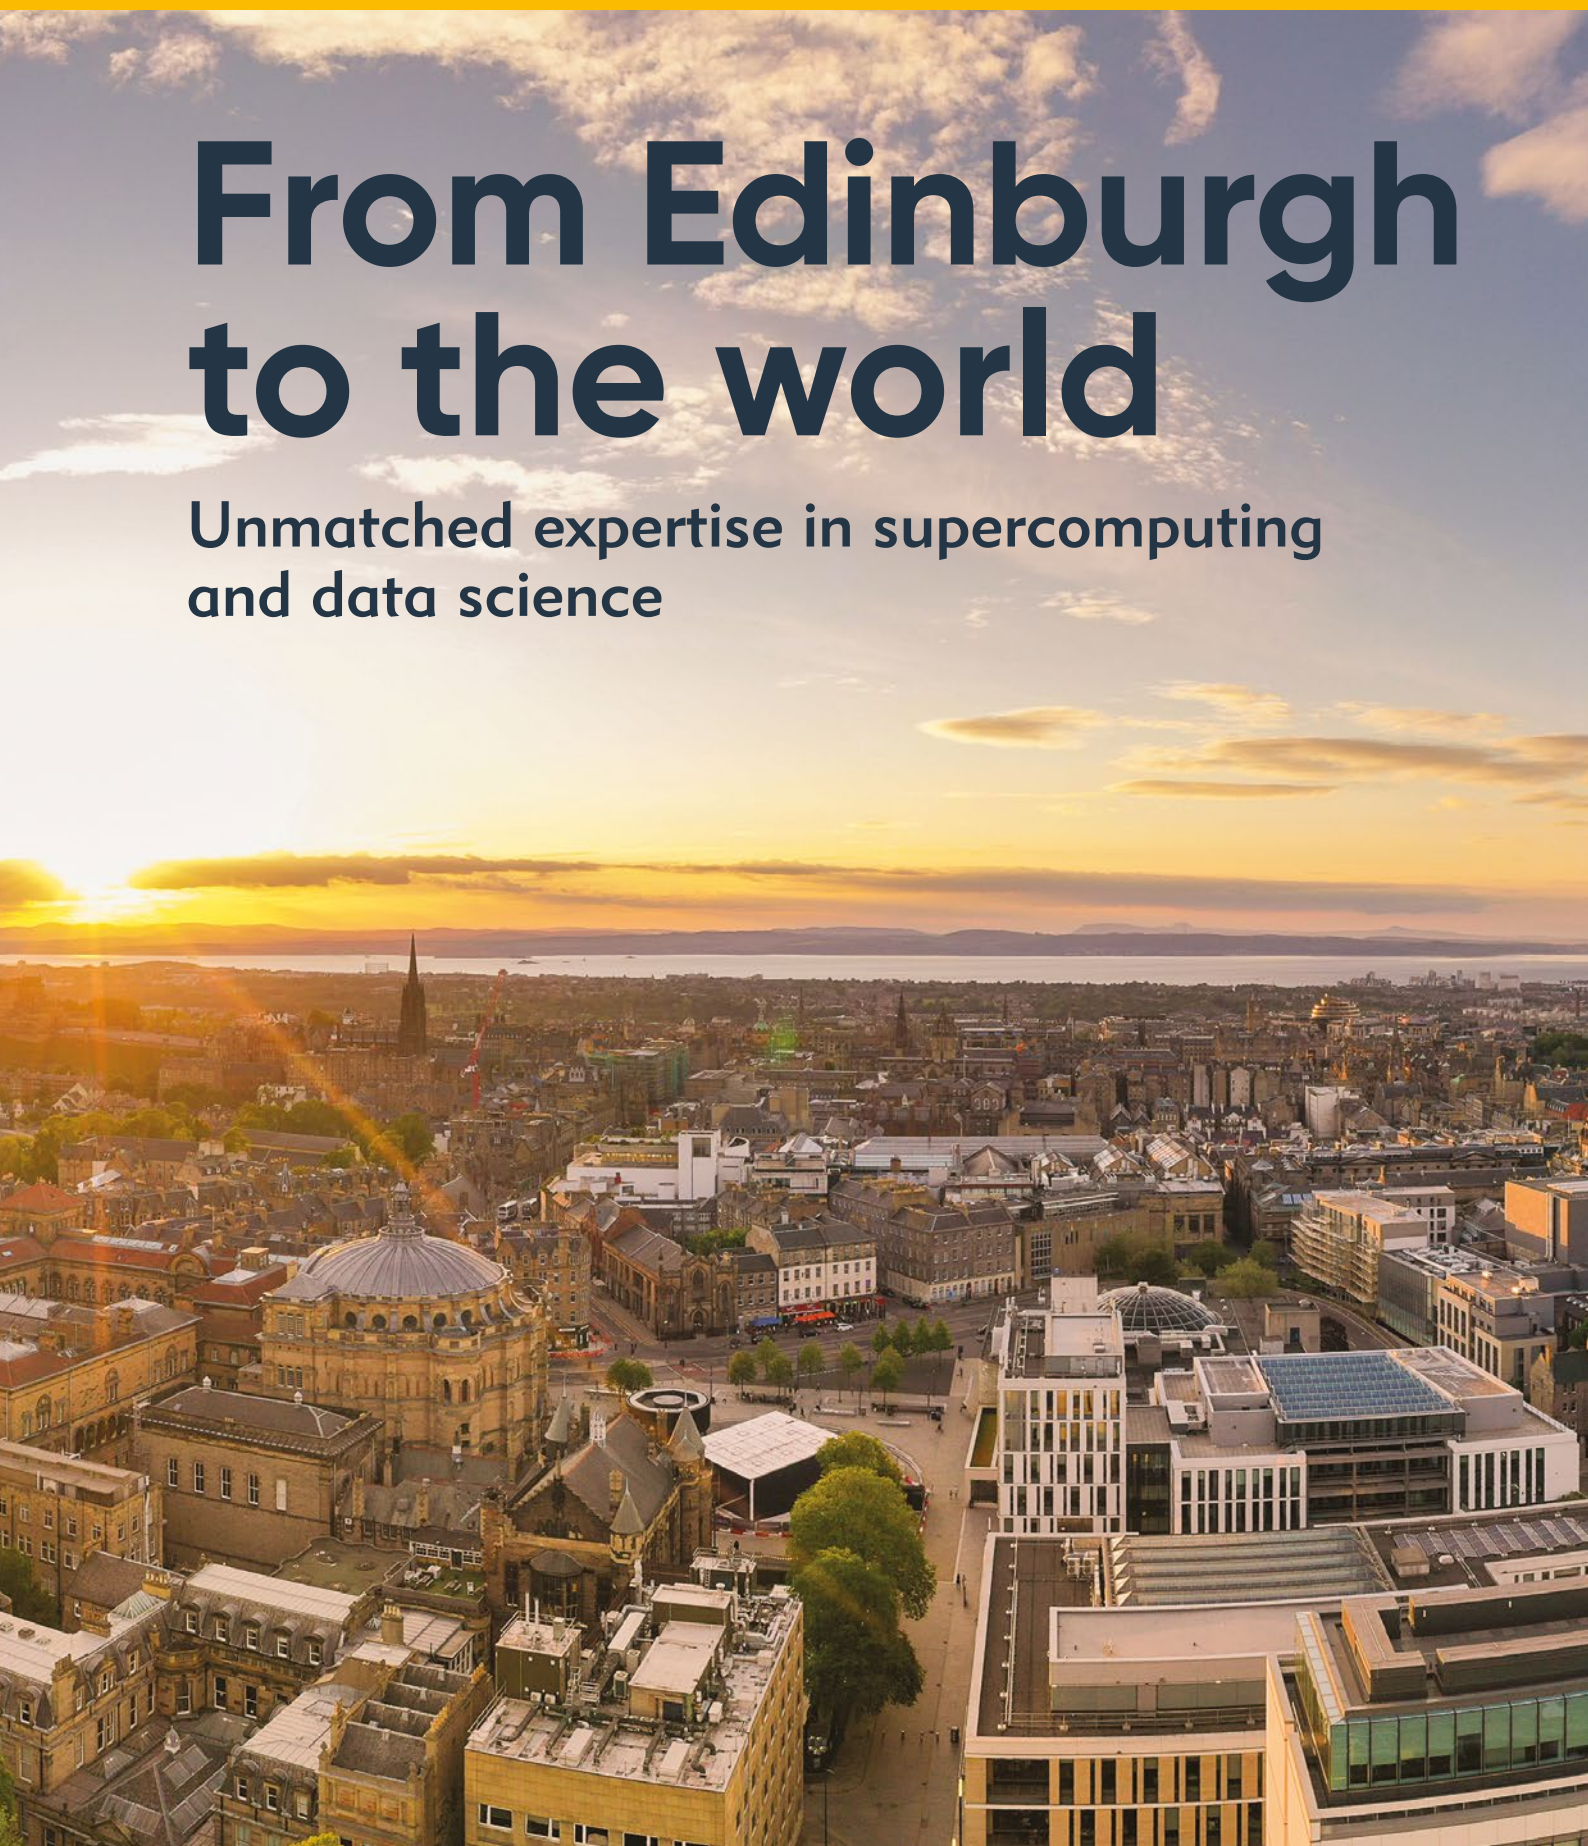

From our Director

Welcome to the Summer 2024 issue of EPCC News.

I was struck when reading this issue by the number of mentions of GPU accelerators and how they are influencing every area of our research. Supercomputing has always progressed through sudden inflection points and we're clearly living through one of those moments now. Artificial intelligence is driving much of the talk about GPUs – largely their availability and their price – but they are also at the heart of the new generation of supercomputers around the world.

You can see their impact throughout our diverse research portfolio, which we document in this issue. They are a key component of our projects and projects that use the Edinburgh International Data Facility. They are transforming traditional modelling and simulation – and are at the heart of most Exascale supercomputers. But because they are also used by the AI community, they are opening up research avenues which bring together AI and modelling and simulation in completely new ways.

Preparations for the UK's first Exascale supercomputer are continuing. The work to extend the power and cooling distribution at the ACF is about to start. But a key preparatory action is now underway – a series of funding calls to support the work to move or rewrite applications so that they can take advantage of the GPU-based Exascale system. These projects will span all of UKRI's research base and those with our industry partners. As a result, our diverse portfolio of activities looks set to grow even further.

It's a very exciting time at EPCC. I hope you enjoy reading this edition of EPCC News.

Professor Mark Parsons
EPCC Director
m.parsons@epcc.ed.ac.uk

EPCC at ISC24

Find us at Booth A10 in the Exhibition Hall, where you can experience our fantastic new virtual reality tour of ARCHER2 and try your design skills with a new demo running on Wee Archie, our mini supercomputer. TeamEPCC will also be competing on-site in the ISC24 Student Cluster Competition, hoping to match the success of last year's team, which was crowned overall winner. We'll have a competition with an exciting prize to win too!

See our website for details of all our activities: **www.epcc.ed.ac.uk**

Contents

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.

Women in HPC: shaping life in high performance computing

Women in HPC (WHPC) is 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.

EPCC is a co-founder of WHPC and is a WHPC Chapter, a community-level group that fulfils the WHPC mission locally. Kicking off 2024 with a bang, the WHPC volunteers at EPCC have already had a busy schedule of events! Here I review some of our most recent events and give some highlights to look forward to at ISC24.

ARCHER2 Celebration of Science

In March, WHPC organised a co-located event at the ARCHER2 Celebration of Science, complementing a programme showcasing the scientific achievements of ARCHER2 users with a session spotlighting the value of diversity in science.

Weronika Filinger, Jenny Wong (2i2c) and I teamed up to lead the WHPC session, opening the floor to the audience to develop ideas for small but meaningful changes to benefit their local communities. We wrapped up the event with an informal networking session to continue the conversations and encourage further connections.

Durham HPC Days

WHPC will also have a spot at the Durham HPC/AI Days in May where we will discuss the challenges facing the compute centre of the future. WHPC will add to the conversation by discussing community building and the importance of allyship in fostering a diverse and inclusive environment.

We have worked closely with the WHPC chapter of the N8 Centre of Excellence in Computationally Intensive Research in Durham to ensure WHPC will make a meaningful impact on important conversations about the future shape of life in high performance computing.

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

ISC24

We are really excited about ISC24!

WHPC is officially collaborating with the ISC conference for the first time to make it an even more inclusive and exciting space for everyone. EPCC has been essential to securing this partnership, with Weronika Filinger alongside Cristin Merritt (Alces Flight) co-leading the organisation of WHPC at ISC24 from the start of the partnership.

This collaboration means even more opportunities for our ISC early-career poster presenters as the WHPC posters are now officially integrated into the ISC24 programme. Our eight fantastic poster presenters will showcase their work through lightning talks alongside the other research poster submissions. We will also award a best Women in HPC poster in-person at the conference.

As the ISC24 WHPC Poster co-Chair, I am excited to see our presenters integrated into the programme and gain even more exposure.

Find us at ISC24!

May 13, 6:45pm–8:30pm:
Highlighting success at the WHPC Solution Forum Takeover.

May 14, 2:15pm –2:30pm:
WHPC poster pitch: our presenters have one minute to showcase their work! Not to be missed.

May 14, 3:30pm–5:30pm:
Come and chat to our presenters at the Poster Reception.

May 15, 9:00am–10:00am:
Learn about the value of mentoring at the WHPC Mentoring Bootcamp BoF session.

Get involved!

Become a member, follow us on social media, join us at events and join your local Chapter or affiliate. We also offer a variety of partnerships, support and sponsorship opportunities for organisations. Find out more on our website: https://womeninhpc.org

New EPCC Directors of Teaching, Research and Data Services

Last year three colleagues took up positions as EPCC Service Directors. Here they each give an insight into their new responsibilities.

Adam Carter
Director of Teaching
a.carter@epcc.ed.ac.uk

This is an interesting time for us, with an ever-growing portfolio of teaching and training activities and some great opportunities to make the most of connections with our exciting work in high performance computing (HPC), Exascale, and data science.

With upcoming Exascale platforms expected to make increased use of GPUs and accelerators, we're putting together a new MSc course devoted to this subject. We're also building on our existing content to create a completely new course on Programming Languages for High Performance Computing.

Running for the first time this year is another new course, Machine Learning at Scale, which addresses the ever-increasing scale of the computations being undertaken to build large machine learning models like the deep neural networks used in large language models. This course also forms an important element of a new joint Imaging, Vision and High Performance Computing MSc programme that we're planning to launch with Heriot-Watt University, led at EPCC by Adrian Jackson.

Our successful MSc programmes in HPC, and HPC and Data Science continue to be led by their long-standing Programme Directors, David Henty (on-campus) and Weronika Filinger (online). I'll also continue in my role as the Programme Director for the University's interdisciplinary online MSc in Data Science, Technology and Innovation.

An important aspect of my new role is to ensure that our multiple education offerings are fully connected with each other and with all EPCC's areas of activity. The expanding number of courses, programmes, and students gives us more opportunities to ensure that EPCC's teaching and training continue to play an important role in the work of the centre.

See page 8 to read more.

Adrian Jackson
Director of Research
a.jackson@epcc.ed.ac.uk

This position includes responsibility for supporting EPCC's academic staff to develop and expand our world leading research. We have a strong base of technical staff actively focusing on research and related activities, alongside colleagues engaged in service provision, application development, and general computational tasks. However, most EPCC staff do get involved with research activities, such as projects that support research, or developing software and/or services that are key for a wide range of researchers.

EPCC's research base has been expanding, most recently with the addition of three new Chancellor's Fellows at EPCC who work on topics from computational fluid dynamics through to quantum computing and large scale data analytics. This has been driven by a new EPCC research strategy to guide the development of our research capability and capacity.

With large government investment in future UK HPC resources, in the form of the AI Research Resources (AIRR) and Exascale systems being developed and installed around the country, we are aiming to further build on EPCC's excellence in large scale computational research to push the boundaries of science in a wide range of areas, supporting the UK academic and industrial computing communities.

We also have a strong climate impact focus in our current research, with a number of Net Zero research projects ensuring EPCC is at the forefront of the drive to minimise the environmental impact of providing computing resources.

These themes, along with the rise of machine learning and associated demand for HPC, promise a busy and exciting couple of years in EPCC's research activities.

See page 10 to read more.

Jano van Hemert
Director of Data Services
j.vanhemert@epcc.ed.ac.uk

I'm thrilled to return to the University of Edinburgh after 14 years in industry working on data science.

My team manages the Edinburgh International Data Facility (EIDF) and it's an exciting time! We're seeing
increased demand for our data and computing services. Customer numbers are growing, and they're using the EIDF for a broad range of research projects.

We're constantly learning from these projects to improve our services. At the same time, we're expanding the EIDF's capacity with more computing power and storage. For more details, see my EIDF article on page 13.

I also look after the team that delivers Safe Haven Services, a secure environment for researchers to analyse sensitive datasets. EPCC provides these secure compute and data environments while conforming to strict information governance processes.

The demand for Safe Havens is growing, both in administrative data and health data. Our upcoming Safe Haven 2.0 version offers researchers more powerful computing resources and greater flexibility in their analyses. Learn more about the details from my colleague Donald Scobbie [1].

As we continue to grow our customer base we will engage to understand the continuous change in requirements. Data scientists make use of a large variety of tools. At the same time we see a demand to operationalise data science pipelines. This year our team will support use cases where we demonstrate how to deliver bespoke data science as a scaleable operation. Watch this space.

Learn with us:
https://www.epcc.ed.ac.uk/education-and-training

Our research:
https://www.epcc.ed.ac.uk/research

Edinburgh International Data Facility:
https://www.epcc.ed.ac.uk/hpc-services/edinburgh-international-data-facility

[1] https://www.epcc.ed.ac.uk/whats-happening/articles/building-epcc-safe-haven-services-20

ISC24 Student Cluster Competition: meet TeamEPCC

TeamEPCC, from left to right: Chi Ho Chan, Zihao Liu, Sean Adrian Quin, Jordan Bell, and Pranav Bharat Chachara.

As I write, we are approaching the ISC 2024 Student Cluster Competition. This year the event will bring together 23 selected student teams from across the world, with eight teams, including TeamEPCC, participating in the in-person competition in Hamburg, Germany.

The ISC 2024 Student Cluster Competition (SCC) is an annual highlight of the conference. Each on-site team aims to demonstrate its abilities in designing clusters, and to achieve the best performance of a set of scientific applications and benchmarks within the power limit of 6KW on its cluster during the three-day competition.

TeamEPCC for ISC24 SCC comprises four students of EPCC's MSc programme in High Performance Computing (HPC), and one student apprentice in EPCC's HPC systems team. With the training from the MSc plus hands-on opportunities provided by EPCC, they hope to excel in the event.

Meet the team

Sean Adrian Quin:
"My aim in joining TeamEPCC is to put into practice the various skills I have gained from the MSc in HPC. This will allow me to grasp a further understanding of HPC outside of the classroom setting. Furthermore, I look forward to meeting industry leaders and other students who share interests in the field."

Chi Ho Chan:
"My goals for joining SCC are to learn beyond the MSc curriculum and apply HPC skills to real-world challenges. Also to interact with other student teams and experts during the competition."

Pranav Bharat Chachara:
"Being a part of the SCC is a great opportunity to apply practices taught to us in the MSc. It also gives us the chance to work with scientific applications and better understand real-world complexities when optimising them. I'm looking forward to meeting the other teams and learning more about the world of HPC at ISC24."

Zihao Liu:
"Participating in the SCC presents a fantastic chance for me to apply my theoretical knowledge in a practical setting and deepen my understanding of the concepts I've studied. Attending ISC24 will also give me the opportunity to engage with experts in the HPC field and to learn about the latest technologies and techniques they employ."

Jordan Bell:
"By taking part in the SCC I will be able to get stuck into compiling, running, and troubleshooting HPC codes as a user, which will allow me to further understand the end-to-end use of HPC systems. I am also excited about meeting our international counterparts and exploring upcoming technologies in our field."

To better understand the hardware components for the competition preparation, TeamEPCC visited EPCC's Advanced Computing Facility (ACF) where the competition cluster and several UK national supercomputing systems, such as ARCHER2 and Cirrus, are located (above and left).

Last year TeamEPCC took overall first place in the Student Cluster Competition at ISC23, beating another seven teams in the conference hall and 15 teams who competed virtually.

Preparation

TeamEPCC has designed its cluster with support from team sponsor Hewlett Packard Enterprise (HPE) and the EPCC systems team. They have opted for a hybrid architecture with a combination of GPUs and CPUs to maximise performance.

The current design of the cluster includes two distinct nodes:

- The GPU node will contain eight NVIDIA H100 GPUs, paired with two Intel Xeon Platinum CPUs and backed with 1TB of RAM
- The CPU node will house two AMD EPYC 9334 CPUs alongside 384GB of RAM.

SCC benchmarks and applications

In addition to HPCC, HPCG and HPL, the scientific applications for the on-site competition this year include RegCM, CONQUEST and Neko. Prior to the ISC24 SCC, TeamEPCC has been focusing on porting and optimising the announced benchmarks and applications with different optimisation techniques on the competition cluster.

The members of TeamEPCC are greatly enthused about applying their skills and knowledge of HPC in the real world at the ISC24 SCC! They are also looking forward to learning from other teams as well as meeting many leading industrial partners and academics in the HPC field at ISC24.

Xu Guo, EPCC
x.guo@epcc.ed.ac.uk

Hewlett Packard Enterprise

We would like to thank HPE for once again acting as TeamEPCC's sponsor for the ISC24 SCC.

ISC24 SCC:
https://www.isc-hpc.com/student-cluster-competition-2024.html

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

Teaching at EPCC: preparing for next-generation systems

Recent announcements about major investments in UK computational capacity raise the question of how computational and data scientists will learn the skills required to exploit the power of these new hardware resources.

Although next-generation supercomputers may offer a revolution in performance, they represent more of an evolution in how we program them. Exascale systems still use the same distributed memory approach of many independent computing nodes connected by a fast network that has been the dominant high performance computing (HPC) architecture for decades. The difference is that the nodes are much more powerful than previous generation systems as they contain many accelerators such as GPUs, making them much faster (and more power-efficient) than a pure CPU node.

Frontier, the first computer to break the Exaflop barrier, has over one-and-a-half times as many nodes as our current ARCHER2 system. However, Frontier has more than 50 times the performance of ARCHER2 (as reported in the top500 list) because it contains almost 40,000 GPUs while ARCHER2 is a pure CPU system.

There are already a number of standard approaches to programming CPU-based distributed memory systems (eg using MPI) and to offloading computationally intensive operations from the CPU to the GPU (eg many codes use OpenaCC, or CUDA on NVIDIA GPUs). The evolution in programming models for Exascale is to combine efficient accelerator offloading with very large scale message-passing in the same application.

With the Exascale era upon us, we have been looking at the content of our MSc in HPC programmes to make sure that we are covering all the right material. The basic parallel programming models of OpenMP (shared memory) and MPI (distributed memory) still apply to the CPUs of an Exascale machine. We already cover GPU architectures and a range of accelerator programming models, but we'll be looking to give these more emphasis and also keep an eye on which of the current candidate models becomes the most useful in practice.

OpenMP itself includes directives for GPU offloading, which has the advantage of being an architecture-neutral standard. This is already covered in the MSc but there are a number of other portable accelerator programming models which are potential candidates for use on Exascale systems. These are often designed to be called from C++ programs, so we plan to place more emphasis on teaching C++ in the future.

This year we introduced the new "Machine Learning at Scale" course which addresses how to exploit large scale accelerated HPC systems for our data science students, and we will also add a related course focusing more on computational science for our MSc in HPC programme.

The UK's large investment in Exascale hardware is expected to open up new employment opportunities for both computational and data scientists. We believe that our MSc in HPC programmes already provide a solid grounding in all the core tools and techniques to exploit such systems, and planned future developments in the curriculum are designed to keep pace with a rapidly changing field.

MSc in Imaging, Vision, and High Performance Computing

We are planning to launch a brand new MSc programme in September 2024, leveraging our existing expertise in high performance computing and moving into new areas of knowledge.

This joint MSc will be run in collaboration with Heriot-Watt University, and combines the knowledge required to understand and efficiently exploit large scale computers with advanced science in the imaging and vision domains.

At the intersection of computational imaging algorithms, computer vision, and parallel computing, the programme will provide a grounding in these and related topics, along with hands-on experience of developing and using vision and imaging models and algorithms on large scale computing platforms.

Machine learning has swiftly become a key tool for vision and imaging approaches, underpinned as it is by large scale computing platforms. EPCC is strongly placed to support and develop machine learning at scale on our GPU computing platforms such as the Cirrus and Edinburgh International Data Facility systems, as well as through access to novel machine learning computing resources such as the Cerebras CS-2 and Graphcore systems that EPCC hosts.

The future deployments of the national AI Research Resource (AIRR) and Exascale systems will also provide large scale resources for students to test, develop, and explore large scale machine learning and computational simulation models.

With computational approaches proliferating in science and industry, demand is rapidly growing for graduates who can both efficiently exploit large scale computing systems, either using machine learning approaches, or more traditional computational simulation, and who understand and can develop the complex computing systems now being commonly used. The skills and knowledge that this new MSc will provide should ensure the student graduates are well positioned to be at the forefront of computing in the coming years.

Adam Carter, David Henty, Adrian Jackson, EPCC
a.carter@epcc.ed.ac.uk
d.henty@epcc.ed.ac.uk
a.jackson@epcc.ed.ac.uk

We are a leading provider of high performance computing and data science education and training in Europe.

If you're interested in finding out more about our MSc programmes, please visit our webpages:
https://www.epcc.ed.ac.uk/education-and-training/masters-programmes

Research at EPCC: spreading the benefits of computing at the leading edge

Rawf8 via Getty Images

EPCC has had a vibrant research culture since its inception, from helping define and create global standards such as MPI and OpenMP, through to developing new hardware architectures and software approaches for academic computing.

As a large centre sited at the interface between a number of different domains and involved with enabling research for many disciplines, EPCC is well placed to spot the gaps in computing where research can bring benefits, and implement and provide those benefits for a wide range of communities.

We have a plethora of active research projects, ranging from very large and long lasting activities such as the ASiMoV Strategic Prosperity Partnership, which EPCC leads with Rolls-Royce, through to small and short collaborations or projects, such as the recent Net Zero for Digital Research Infrastructures sandpit activities.

If you're interested in the full range of research EPCC has been involved with, please see our website for details and also links to all our recent publications. To give a flavour of the breadth and types of research we undertake at EPCC, I will briefly describe three current research projects.

CONTINENTS

EPCC has recently been jointly awarded an International centre to centre research project with the National Centre for Atmospheric Science (NCAS) in the UK, and the NSF National Center for Atmospheric Research (NCAR) in the USA. Such collaborations aim to provide leading UK research groups with the opportunity to work in partnership with the best international researchers.

The International Collaboration Towards Net Zero Computational Modelling and Simulation (CONTINENTS) project will build on the expertise of all three centres in sustainability and power/energy efficiency of computational modelling and simulation.

We have defined an ambitious programme of research that will drive innovations in data centre and system operation, optimal exploitation of hardware, machine learning applied to data analysis and numerical modelling, and software design and implementation strategies. The ultimate goal is to significantly reduce the climate impact of high performance computing (HPC).

Quantum computing

The range of active projects in quantum computing at EPCC currently includes exploring the application of quantum computing to end user challenges, how quantum computers can be programmed efficiently and effectively in collaboration with traditional high performance computing systems, and investigating where the quantum advantages are going to be. To this end we currently have a number of active research projects.

At the applications level we have two projects funded by Innovate UK, the first being a collaboration with HSBC and Rigetti that is

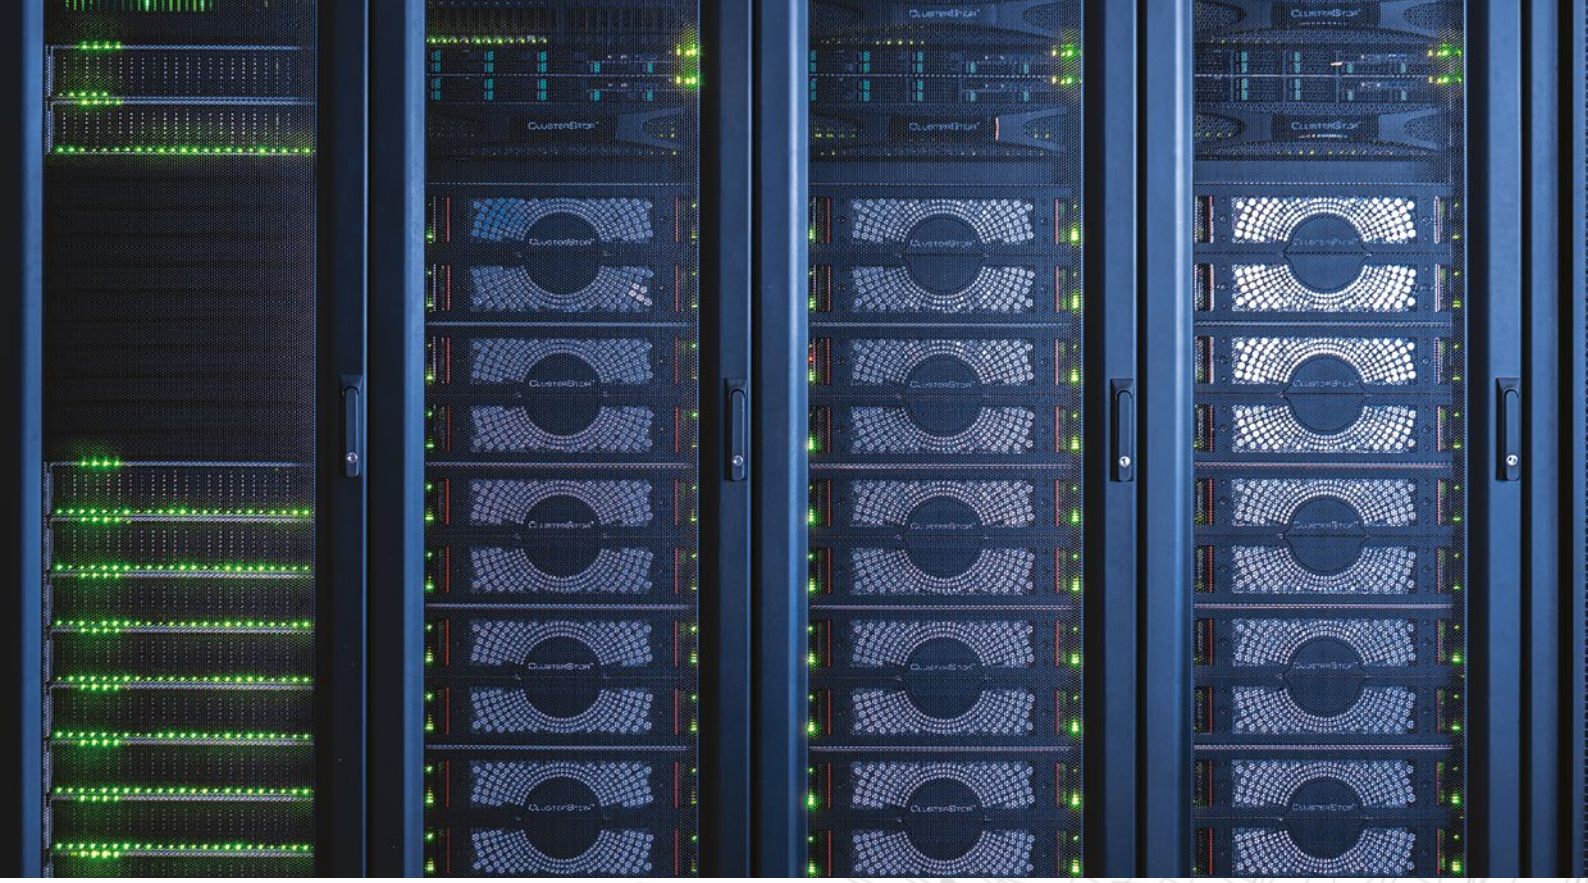

exploring how quantum machine learning (QML) could be used to detect money-laundering.

The second is exploring the applicability of QML to reducing carbon emissions in aviation, in collaboration with Quantum Base Alpha and NATS.

When it comes to the question of programming a quantum computer, EPCC presented work on quantum task offloading with OpenMP at the 2023 Supercomputing international conference in Denver, USA, and we now have a PhD student (co-supervised by STFC Hartree Centre) who will be further exploring programming models for hybrid HPC and quantum computing. Underpinning all this work is our ability to run large scale classical simulations of quantum hardware.

We recently presented results from studying and optimising the energy consumption of an exact 44-qubit simulation on 4,096 nodes of ARCHER2 using QuEST. Following on from that we are currently in the middle of a project to scale QuEST on multiple GPUs, complemented by an ARCHER2 eCSE project looking to scale a statevector simulator using one-sided MPI communications.

Hardware testbeds and services

There is a wide range of computing systems and hardware available at EPCC beyond the large scale systems such as Cirrus and ARCHER2. We have novel machine learning hardware, such as the Cerebras CS-2 and Graphcore systems, which we are leveraging to investigate machine learning performance and bottlenecks for a wide range of applications.

In the same vein, we have FPGA and RISC-V testbeds, provided as UK academic resources for researchers, that enable research on alternative computing approaches to be explored by EPCC researchers, collaborators, and the wider computing community.

Adrian Jackson, EPCC
a.jackson@epcc.ed.ac.uk

To keep up to date with EPCC research activities, please visit our website and browse our latest publications. If you have ideas for research collaborations don't hesitate to get in touch with me.

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

INchendio via Getty Images

Sharing next-generation software research

ExCALIBUR is a UK research programme working to enable the exploitation of future supercomputers by the next generation of high performance simulation software.

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

The ExCALIBUR programme is oriented around five key strands: high priority use cases; emerging requirements for high performance computing (HPC); cross-cutting research; novel architecture testbeds, and knowledge exchange. Over 45 projects have been funded so far and a wide range of activities undertaken.

The knowledge exchange strand is a major part of the programme, both by connecting all the individual projects together to ensure they are sharing information and common work, as well as positioning ExCALIBUR within the global computing community.

This is crucial because as more Exascale supercomputers become available, the community must have the applications, tools and algorithms necessary to make the most effective use of these instruments. Ensuring that advances made during ExCALIBUR are exploited globally and complement efforts across the world will help us advance further into the Exascale era.

Knowledge Exchange Coordinator

Each ExCALIBUR project has a Knowledge Exchange Coordinator, and there has been an active network since the programme began. However it was felt that in this final phase of the programme, as many of the projects are now delivering significant insights and advances, it would be valuable to have someone with time dedicated to coordinating this network and helping organise knowledge exchange at the programme level.

I have taken on this coordination role, and it has been extremely interesting to become involved in the range of activities that are progressing across the programme. While I was already involved in several ExCALIBUR projects and felt I knew the programme well, I have been surprised at the breadth of successes and activities across the projects.

We have a busy programme of events planned in ExCALIBUR. To date we have participated in the SC23 conference, and an ExCALIBUR focus session and several main session talks at Computing Insight UK, and most recently had an ExCALIBUR stand and poster at the ARCHER2 Celebration of Science. In the coming months the programme will have significant involvement in HPC Days Durham, ISC, and PASC, as well as many other smaller conferences, workshops, and events. See the ExCALIBUR website for details.

The Exascale Computing ALgorithms & Infrastructures Benefiting UK Research (ExCALIBUR) is a UK research programme that aims to deliver the next generation of high performance simulation software for the highest priority fields in UK research.

ExCALIBUR began in 2019 and will run through until March 2025, redesigning high priority computer codes and algorithms to meet the demands of both advancing technology and UK research: https://excalibur.ac.uk

The two Cerebras CS-2 Wafer-Scale units available to users of EIDF.

New services for data science research

In 2023, we launched a GPU Service on the Edinburgh International Data Facility (EIDF) which provides access to a cluster of 112 NVIDIA A100 GPUs via Kubernetes. We have just added five NVIDIA H100 GPUs and will add six more soon.

Large language models

Seventy projects have used EIDF services since 2022. For example Bea Alex from the School of Literatures, Languages and Cultures at the University of Edinburgh and her team have used EIDF to fine-tune and prompt large language models for clinical classification.

Here's what Bea said about her experience of using the GPU Service:

"My students and I have been extremely delighted to have access to this invaluable computing resource. It has enabled us to conduct rigorous analyses of LLMs (large language models) as part of a group project in the CDT for Biomedical AI and to participate in a competition on natural language inference for clinical trial data."

On the subject of large language models, we have seen an explosion in the popularity of generative AI in recent years. At EPCC we anticipate researchers will need to tune ever larger language models.

Cerebras CS-2

To accommodate this, we now offer access to two Cerebras CS-2 Wafer-Scale units coupled with a high-speed network.

This setup enables researchers to handle models with more than one billion parameters through a new operating mode called weight streaming, which does not require all weights to be in memory all the time. See the link to Nick Johnson's article to read about the advantages of this approach [1].

Cerebras has trained seven models, ranging from 111 million to 13 billion parameters, using the Pile dataset, which was used to create the open source Pythia models. To enable researchers to get started with LLMs on Cerebras, the company has shared these models under the Apache 2.0 open source licence as Cerebras-GPT [2].

Data-focused infrastructure

EIDF continues to be a data-focused infrastructure. We are developing services to make it easier to work with data inside and outside EIDF. We will enable S3 to make data transfers easier and more interoperable. I hope to provide updates on this front soon.

Jano van Hemert, EPCC
j.vanhemert@epcc.ed.ac.uk

EIDF is governed by the Data-Driven Innovation (DDI) initiative with capital investment from the UK Government. DDI, part of the Edinburgh & South East Scotland City Region Deal, comprises a cluster of innovation hubs that bring academic disciplines together to delve into some of the world's most pressing challenges, using data to innovate.

Anyone can apply for any EIDF service. Please see:
https://portal.eidf.ac.uk

Data-Driven Innovation initiative:
https://ddi.ac.uk

Edinburgh International Data Facility:
https://edinburgh-international-data-facility.ed.ac.uk/

[1] https://www.epcc.ed.ac.uk/whats-happening/articles/faster-large-language-model-training-cerebras-cs-2-wafer-scale-cluster

[2] https://github.com/Cerebras/modelzoo

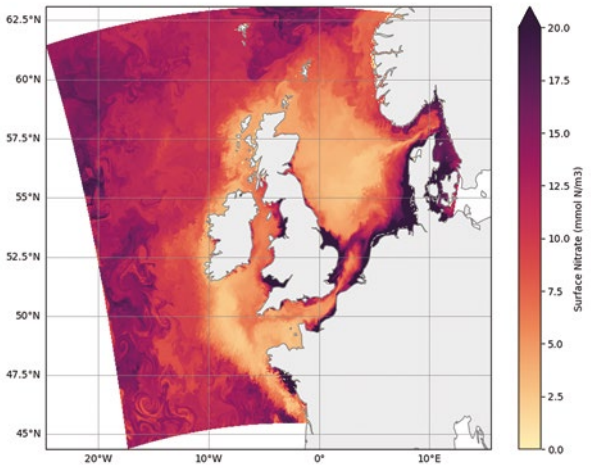

From the ARCHER2 eCSE report "Optimising NEMO-ERSEM for High Resolution", Jerry Blackford (Plymouth Marine Laboratory) et al.

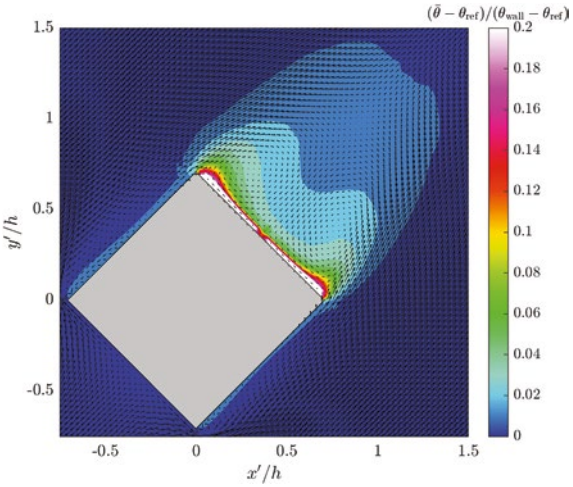

From the ARCHER2 eCSE report "uDALES: towards exa-scale simulation of air quality and microclimate in urban environments", Maarten van Reeuwijk (Imperial College London) et al.

Launch of eCSE programme to support GPU-based software development

EPCC has run an eCSE software development programme for over a decade now, first under the ARCHER UK national supercomputing service and more recently under its successor ARCHER2.

Chris Johnson and Catherine Inglis, EPCC

c.johnson@epcc.ed.ac.uk
c.inglis@epcc.ed.ac.uk

eCSE (embedded computation science and engineering) calls have mainly focused on the development of software running on those two systems and have therefore centred on developing software for CPU architectures. As we move towards the Exascale era, architectures are likely to be much more GPU-based and there is still a lot of preparatory work to do – not only to move currently CPU-based software to GPUs and optimise existing GPU software, but also to develop the skills needed to run efficiently on such architectures. The new GPU eCSE calls plan to address these issues.

The first GPU call opened in January 2024, and proposals were invited for development of software enabling research in any area within UKRI's remit. This now opens up the programme to developers of software used in research funded by the Economic and Social Research Council and the Arts and Humanities Research Council, for example, as well as in scientific fields which have traditionally embraced research software.

The first batch of successful projects will start in late spring or summer 2024. Each project can have a maximum duration of two years and can fund up to 36 person months of effort, allowing significant development work to be carried out.

Examples of the work which might be funded are the development effort to take a CPU-based code and transform it into a GPU-based code, or the improvement of an existing GPU code – for example to enable effective use of an increased number of GPUs, or to improve the sustainability and maintainability of such a code.

We also run an observer programme allowing early career researchers the opportunity to attend a panel meeting where the funding decisions are made. By giving them a better insight into how applications are evaluated and decisions are made, we believe they will be better prepared to write their own funding proposals in the future. Typically we have around three such observers at a panel meeting and feedback suggests early career researchers have found the experience valuable.

Examples of work carried out under the ARCHER2 eCSE programme, with an outline summary of the work and a link to the technical report for each project, can be seen on the ARCHER2 website. Recent projects include improving codes used for modelling air quality in urban environments, carrying out earthquake simulations, and investigating atomic or ionic diffusion in the search for novel materials for use in batteries and fuel cells.

ARCHER2 eCSE reports:
https://www.archer2.ac.uk/ecse/reports

Byte-Sized RSE: essential skills training for busy researchers

GCShutter via Getty Images

Accessible fundamental research software training instils best practices and cultivates the skills and confidence needed for users to pursue advanced courses, thus widening access routes and dismantling entry barriers.

A recent survey found that a lack of time available for training was one of the biggest barriers to training across all domains sampled. Many researchers would like to access training but are simply unable to dedicate large chunks of time to it.

To address this issue the UNIVERSE-HPC project, which is funded through the SPF Exascale Computing ALgorithms & Infrastructures Benefiting UK Research (ExCALIBUR) programme, has developed and is running a programme of short-form training via a new series called “Byte-Sized RSE”.

The Byte-Sized RSE programme takes core research computing topics and provides a simple, accessible introduction in short sessions of between one hour and 90 minutes. While this is a great way for newcomers to the research software and research computing space to learn key skills, our later sessions have also been tackling more advanced topics.

Twelve online events have been run so far on topics including version control (Git), code review, Integrated Development Environments, and Continuous Integration. Most of these events are run standalone but two were run within the Software Sustainability Institute’s Research Software Camps to try to reach new audiences.

Every Byte-Sized RSE session is followed up by a professionally produced podcast episode, released through the Code for Thought series. This explores the topic in more depth, helps to reinforce learning and allows a deeper understanding of the material covered. In addition, as the podcast can be downloaded and listened to at any time, it allows people the flexibility to access it when their schedule allows.

The first series of Byte-Sized events was organised by members of the UNIVERSE-HPC team: project partners Jeremy Cohen (Imperial) and Steve Crouch (University of Southampton).

For the second series the team is encouraging people from across the community to propose topics and to develop and lead sessions, thus providing the wider community with an opportunity to promote and teach topics that they see as important, while following the successful Byte-Sized RSE format.

By providing these fundamental, accessible courses we aim to provide a route for people to build knowledge, experience and confidence so that a wide range of people are equipped with the skills required to exploit the huge opportunities and the potential advantages offered by Exascale infrastructure.

Kirsty Pringle, EPCC
k.pringle@epcc.ed.ac.uk

What our participants say:

“Absolutely brilliant series - great structure and a good covering of topics. I feel much more confident in implementing things myself.”

“A useful quick summary plus an activity within an hour worked well rather than taking multiple hours.”

“I liked the live coding session and low barrier to entry.”

“I have really enjoyed this series of training sessions, I am pleased that more are being planned. The format, frequency and short time frame make them very accessible for me.”

Code for Thought:
https://codeforthought.buzzsprout.com

Byte-Sized RSE:
https://www.universe-hpc.ac.uk/events/byte-sized-rse/

Building a science platform for the Legacy Survey of Space and Time

LSST camera and SLAC camera team. Image: Jacqueline Ramseyer Orrell/SLAC National Accelerator Laboratory

After nine years of construction, the Vera C. Rubin Observatory is almost ready to begin its state-of-the art Legacy Survey of Space and Time (LSST) which – over the next ten years – will employ the world's largest digital camera to image the entire southern sky with an unprecedented level of fidelity and sensitivity.

Operations are on track to commence in early 2025, from which point around 20PB of data will be served each year from the host site in Cerro Pachon (Northern Chile) via dedicated data facilities in the United States, France, and the UK to an international community of around 10,000 astronomers.

The Simonyi telescope, which is the heart of the Rubin Observatory, has an extremely large field of view and can manoeuvre very quickly, allowing much faster imaging of the sky. Over the ten years of the survey, the telescope will photograph every point in the southern sky an average of 800 times.

Taking many images of each patch of sky means the survey is very sensitive and able to detect even the faintest objects, measuring an estimated 20 billion galaxies and a similar number of stars, most of which have never been observed before.

However, visiting every point in the southern sky every three to four nights means it will also detect subtle changes, such as the beginning of a supernova, a huge explosion that marks the death of certain types of star and helps us to understand the expansion of the universe. The telescope is expected to identify millions of changes in the sky each night.

Science potential

UK-based astronomers are very excited about the science potential of Rubin and, recognising this, formed the LSST:UK Consortium in 2013 to coordinate involvement in the telescope and secure long-term funding for UK astronomy priorities.

Rubin is as much a computing facility as it is a telescope. The sheer volume and complexity of data that will be captured makes the traditional approach to survey astronomy impractical and will challenge the capabilities of research computing worldwide.

Every step, from the initial processing and cleaning of survey images, to the measurement and preparation of survey products, and the end-user analysis of the survey by researchers, requires an international-scale computing resource, and EPCC, working with the LSST:UK Consortium, is a key contributor to this.

Data Access Centre

EPCC, in collaboration with the Institute for Astronomy at the University of Edinburgh, has a critical role in the design and operation of a UK-based science platform (one of only three) called a Data Access Centre or DAC for short, to connect astronomers with at least 200 Petabytes of survey products and provide advanced tooling, adapted from other HPC and data-intensive science domains, to handle these products.

UK alert-broker service

EPCC staff, in collaboration with research software engineers (RSEs) in Queen's University Belfast are also developing a UK alert-broker service called Lasair, which will receive, analyse and filter the tens of millions of nightly alerts in pseudo-real-time. Already, for the precursor Zwicky Transient Facility, Lasair is providing a valuable science platform with more than 100 users and acknowledgements in more than 80 academic publications.

Rubin Observatory. The 8.4-metre telescope at Rubin Observatory, equipped with the highest-resolution digital camera in the world, will take enormous images of the southern hemisphere sky, covering the entire sky every few nights. RubinObs/NSF/AURA/H.

Somerville Research Cloud Service

At the heart of the UK DAC is a research cloud service called Somerville, which is hosted in EPCC's Advanced Computing Facility (ACF). Built on top of open source cloud and virtualisation technologies such as OpenStack and Kubernetes, during Rubin operations Somerville will support around 1,000–1,500 astronomers internationally to pursue their science ambitions.

The DAC provides three different user interfaces to LSST: a web-based science portal for interactive enquiries, a notebook-based engine for more substantial and scripted analysis workflows, and a high performance computing/batch interface for the most ambitious and computationally intensive research campaigns. These interfaces are backed by custom services, engineered within the Rubin project to the requirements of the survey. They include a bespoke distributed database system called Qserv, able to serve LSST catalogues containing hundreds of billions of objects and to satisfy the demands of hundreds of concurrent queries.

Already, Somerville hosts simulated survey databases containing more than 120 billion records occupying more than 100TB: a record size for EPCC.

Where possible, Rubin has drawn on experience of other computing infrastructures to reinforce its platforms and, in particular, we are working with the high-energy physics community at CERN for large scale data movement and complex workflow tooling. The Rucio distributed data-management system, in part developed by EPCC, is the backbone for Rubin data management, responsible for moving tens of petabytes of data each year between North America and Europe, as well as on to end-user astronomy groups around the world.

With a little under a year to go until the beginning of operations, it is a busy time for all involved in the Vera C. Rubin Observatory, not least, the EPCC team who are working to ensure the UK Data Access Centre, hosted in EPCC's ACF, is ready to play its part in Rubin early science.

Open and standardised are priorities for the Rubin community, both for the software developed inhouse and the software solutions sourced from the wider research community. All Rubin software is open source and wherever possible adopts the standards from International Virtual Observatory Alliance (IVOA), an organisation to which EPCC has contributed since its inception in 2002.

George Beckett, EPCC
George.Beckett@ed.ac.uk

LSST:UK Consortium

Vera C Rubin Observatory:
https://rubinobservatory.org

LSST:UK Consortium:
https://www.lsst.ac.uk

The Edinburgh International Data Facility: open for business

The Edinburgh International Data Facility promotes and enables innovative research and development projects from companies in the data science and machine learning sectors. EPCC is proud to provide access to this world-class private research cloud computing infrastructure.

Regular readers of EPCC News will have noticed a growing number of articles about the Edinburgh International Data Facility (EIDF for short). The EIDF is EPCC's first foray into providing large scale compute infrastructure specifically to promote and enable data science research.

The impressive infrastructure of EIDF facilitates innovative data-focused industry research and promotes collaboration between industry and academia. It brings together and makes available relevant and valuable datasets with the compute infrastructure required to extract this value.

We are continually looking to improve and curate our collection of datasets – if you are looking to access a specific dataset, please get in touch to find out whether it is either already available or could be added to EIDF's ever-growing collection.

EPCC charges a transparent monthly fee calculated on the requirements of individual users prior to any commitment to use EIDF.

See page 13 to read more about the Edinburgh International Data Facility.

EIDF includes:

- An openstack virtual machine (VM) server that can host both small VMs for day-to-day tasks and more specialised, high-memory VMs for data-intensive work.
- A Kubernetes GPU server with 112 NVIDIA A100 GPUs, which is expected to double in size and incorporate NVIDIA H100 GPUs.
- Large scale backed-up storage solutions to access and share datasets with other groups or researchers in a secure and efficient manner.
- Access to a centrally-maintained catalogue of valuable datasets (such as the entirety of the European Space Agency's Sentinel data) to facilitate your research and development activities.
- AI-ready systems (such as two Cerebras CS-2 systems) that can significantly speed up the training time for your machine learning models.
- Access to EPCC's highly reputable user support and advice services to minimise the time from access to actual research.

If EIDF sounds like the infrastructure your company has been looking for, we would be more than happy to discuss how it can be accessed.

Julien Sindt, EPCC
J.Sindt@epcc.ed.ac.uk

EIDF
Edinburgh International Data Facility

To find out more, contact:
commercial@epcc.ed.ac.uk

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

Processing images on a petabyte scale

The National Collection of Aerial Photography is working with EPCC to reduce the expected data footprint of digitising its ever-growing collection of aerial photographs from an estimated 7 petabytes to an estimated 4 petabytes.

The National Collection of Aerial Photography (NCAP) is the custodian of one of the world's largest collections of aerial photography. It is in the process of digitising its ever-growing collection of over 30 million aerial photographic images, which date from as far back as the 1920s, to ensure that these images are preserved for future generations.

The physical photographs are stored on multiple media (photographic films or prints using different technologies across the decades), each of which requires different specialised equipment to digitise. Frequently the resulting scanned images contain extraneous data – for instance, photographs digitised on a scanner will contain "borders" of the total scanned area. This additional data can significantly increase the size of the resulting digital file (sometimes by up to 65-70%).

NCAP and EPCC are working together to develop an algorithm that automatically extracts the original photographic image from the full digital file and discards the extraneous data.

Our initial efforts produced a cropping algorithm that finds and selects the largest contiguous area of a different colour to the back-plate of the scanners. Nine times out of ten, this algorithm accurately identifies the original photograph within the large digital file, and will automatically save it as a new file with significantly reduced size (on average, the initial digital image is around 240MB whereas the cropped aerial photograph is around 140MB in size).

However there are certain cases, especially night-time photographs, where the image is hard to distinguish from the scanner plate's black background. In these cases, the algorithm is quite likely to fail.

NCAP and EPCC are exploring methods for overcoming this situation, including investigating whether adding a pattern to the scanner plate then using a pattern detection algorithm to remove this data could have the desired effect.

While we are continuously evolving this image processing software, NCAP has already begun to process a backlog of over one million scanned digital images on the Edinburgh International Data Facility (EIDF).

EPCC has provided NCAP with access to a data-processing virtual machine that enables NCAP to process images in an average of 30 seconds (over 2,500 images per day). NCAP and EPCC are developing an automated workflow that will transfer digital image files to the EIDF immediately after scanning for them to be processed and transferred to NCAP's data storage facility. This will ensure that the full digitisation of the NCAP dataset will happen within a reasonable time, thereby preserving this invaluable dataset for future generations.

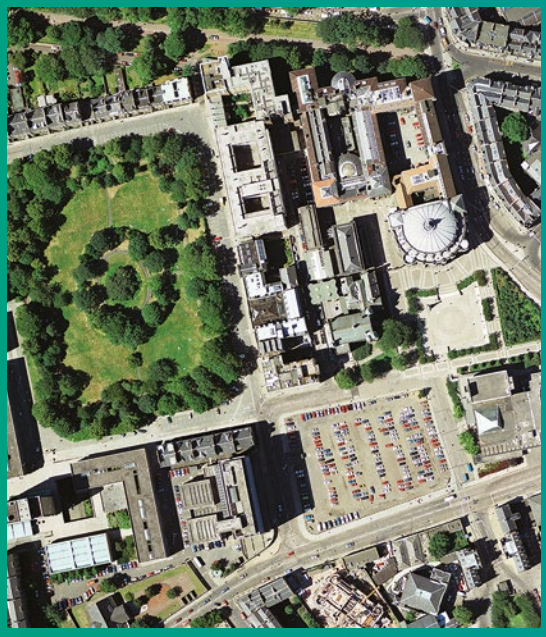

This old image shows the central campus of the University of Edinburgh, including the future site of EPCC's home in the Bayes Centre. Image courtesy NCAP.

Julien Sindt, EPCC
J.Sindt@epcc.ed.ac.uk

"The partnership between NCAP and EPCC has transformed our capacity to process millions of high-resolution aerial images that record places throughout the world and key moments in history."

Dr Allan Williams
Head of NCAP

The Edinburgh International Data Facility (EIDF) provides the enabling data infrastructure platform for the Data-Driven Innovation initiative:
https://ddi.ac.uk/

National Collection of Aerial Photography:
https://ncap.org.uk

Edinburgh International Data Facility:
https://www.epcc.ed.ac.uk/hpc-services/edinburgh-international-data-facility

Building the Virtual Human

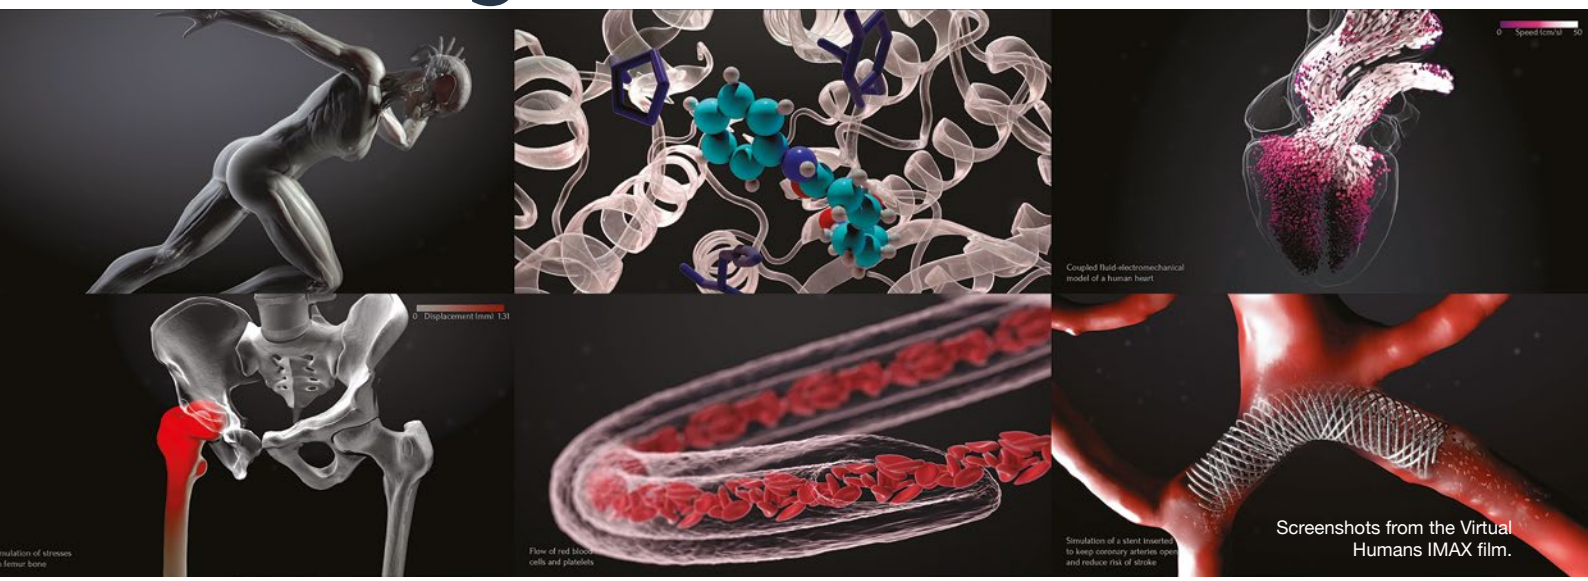

Screenshots from the Virtual Humans IMAX film.

EPCC was a core partner of the recently ended CompBioMed Centre of Excellence, a European Commission-funded Centre of Excellence focused on the use and development of computational methods for biomedical applications, which worked towards truly personalised medicine: the Virtual Human.

As part of this seven-year endeavour, the CompBioMed Centre of Excellence (CoE) began optimising its applications for Exascale in partnerships between application code developers, high performance computing (HPC) centres, and hardware vendors, using co-design principles.

As CPU speeds are no longer increasing, Exascale architectures must aggregate large numbers of multi-core nodes with accelerators, such as GPUs. To exploit these new platforms, we must optimise strong and weak scaling for applications where spatial size or resolution are key and combine applications through coupling and complex workflows. These three compute patterns are as follows.

- Monolithic: deployment of a single computational job spread over a substantial fraction of the compute resources of a single supercomputer, including the pipelines necessary to manage such large computational jobs.
- Coupled: deployment of multiple, communicating subcomponents each assigned to a sub-section of compute resources, in total comprising a substantial fraction of the supercomputer.
- Ensemble: multiple instances of an application launched in parallel with different input data; each such instance may be a monolithic parallel code, itself running at extreme scale. These are now becoming part of highly complex data intensive workflows combining conventional HPC applications with machine learning (ML) components. Moreover, ensemble simulations are required for validation, verification, and uncertainty quantification; these are increasingly becoming essential for uptake of in silico methods in clinical medicine.

The CoE's monolithic codes include the lattice Boltzmann method, the underlying fluid solver algorithm used in HemeLB, Palabos and HemoCell (via Palabos) for hemodynamics, and the Finite Element Methods used in HemeLB, Palabos, HemoCell and Alya, and their coupling for multiphysics problems and complex workflows.

The CoE's IMPECCABLE workflow couples ML and physics-based methods to accelerate the drug discovery process, which employs RADICAL-Cybertools to enable the portable execution of ensemble-based applications at extreme scale.

Co-design

EPCC was also part of the EXCELLERAT CoE, in which we led the co-design task. Here we employed an "indirect" co-design paradigm: code authors gained access to the early release of state-of-the-art hardware, which is available typically due to a close working relationship between the vendor and HPC centres. This avoids the "direct" co-design paradigm, given vendors typically will not alter their hardware to benefit a small set of applications.

Naturally, lessons learned were shared with CompBioMed and, using co-design principles, the CoE ported Alya to energy-efficient RISC-V vector units, and developed GPU-enabled versions of HemeLB, Palabos and Alya, and a GPU-prototype of HemoCell.

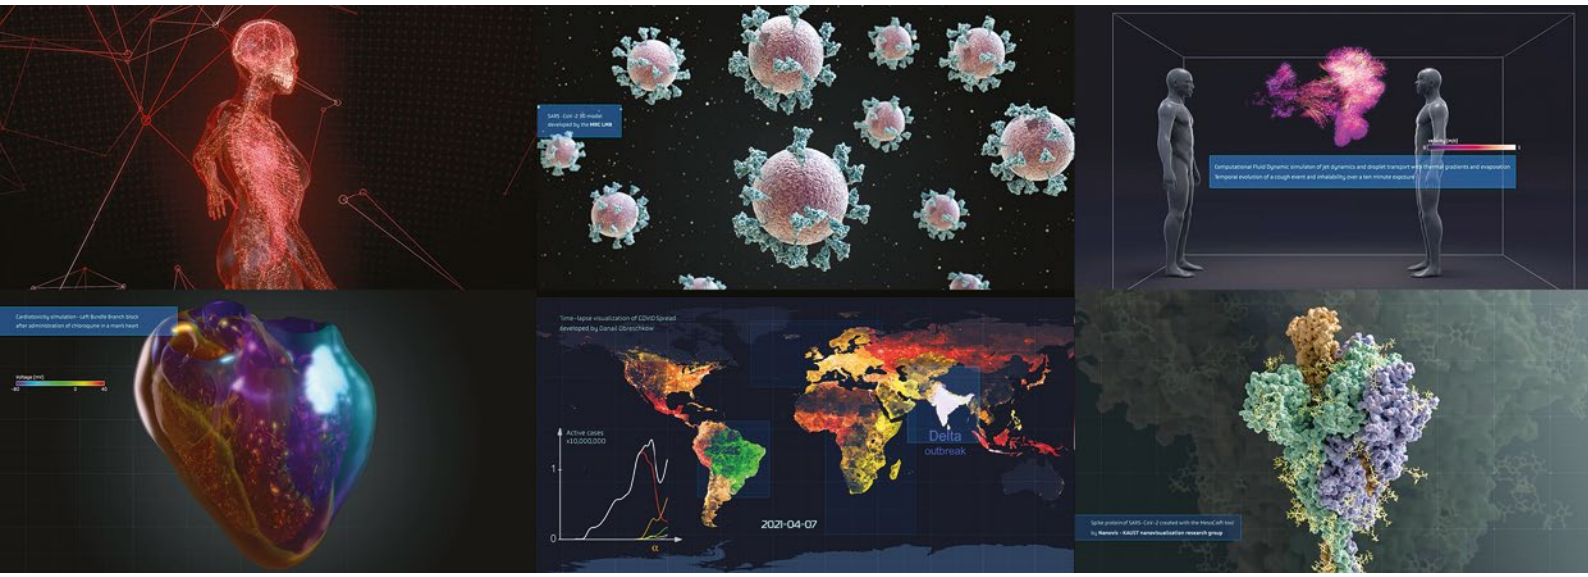

Screenshots from The Next Pandemic IMAX film.

The porting efforts to various GPU programming models ultimately depend on the longevity of the API, which might be abandoned due to insufficient adoption or lack of strategic benefit for the vendor. As such, the CoE looked to widely supported parallel programming models such as OpenMP, or the use of portability frameworks such as Kokkos or SYCL, or via standard parallelism in C++ or Fortran.

The CoE received early access to machines at Oakridge Leadership Computing Facility (OLCF). The CoE's HPC partners (EPCC, LRZ, SURF, and Atos) also offered early system access to a diverse collection of hardware, including ARM clusters, NVIDIA, and AMD GPUS, and FPGAs. Further, to test applications at full scale, the CoE was provided the entirety of supercomputers.

Such access usually occurs before service downtime, when it is tolerable to put the machines under higher stress than usual, and success is only possible via close coordination between application developers and HPC staff. These tests also expose bottlenecks in I/O and pre- and post-processing requirements.

Resilient urgent computing

Exascale machines have a very high node count, and, while individual nodes may have a reasonable mean time to failure, a collection of hundreds of thousands of nodes together in a single system will dramatically decrease the overall mean time to failure.

One solution is to continue failed simulations to complete elsewhere, and EPCC was instrumental in developing a Resilient Exascale Workflow for time- and safety-critical biomedical simulations, in collaboration with LEXIS and EXA4MIND, where results are required at the surgical operating table in faster than real-time.

CompBioMed was funded by the European Commission under H2020, grant agreement no. 675451, H2020-EINFRA-2015-1, and no. 823712, H2020-INFRAEDI-2018-1.

Gavin J. Pringle, EPCC
g.pringle@epcc.ed.ac.uk

As part of CompBioMed, EPCC led the creation and maintenance of a crib sheet for code authors to prepare their applications for Exascale: https://www.compbiomed.eu/rough-guide-to-preparing-software-for-exascale

For more details and other personal CompBioMed highlights, please see my PubMed article PMID: 37702935.

CompBioMed website: https://www.compbiomed.eu

ACF update: undertaking major works in a "live" environment

Here at EPCC we have been making plans for the extension of the site space and power capabilities at the Advanced Computing Facility for some time now. Our plans are progressing well, but we still have a lot of work to do to bring them to fruition. This article describes a recent project undertaken as part of this long-term programme of works.

To ensure we are ready for any site expansion, there must also be extensive reconfiguration project works in another part of the Advanced Computing Facility (ACF) site, to accommodate any relocation of equipment from Computer Room 4 (CR4). This room currently houses equipment relating to the Edinburgh International Data Facility.

Fortunately, we have adequate floor space available in Computer Room 1 (CR1) for relocation of equipment. However, this space would then require a new dedicated mechanical and electrical (M&E) infrastructure to support the additional equipment.

Therefore, a major M&E project was required to be scoped, designed, costed, and then installed before any relocation could even begin to happen – and this work had to be done without disrupting the operation of the existing systems in CR1.

A decision was taken to expand use of the existing HPE ARCS cells (Adaptive Rack Cooling System) already in use in CR1 to accommodate the relocated equipment from CR4. This type of system was selected due to its high efficacy, low noise, and other benefits that it provides.

The plan to use ARCS also drove the scope of the M&E project works, which were a prerequisite to allow the installation of these cooling distribution units (CDUs) and cabinets.

Once the room layout had been agreed, the ACF HPC Systems team engaged with the University of Edinburgh's Estates department, RSP M&E Consultants, and M&E contractor A McKay to take the project forward and, as usual with projects at the ACF, it moved at an accelerated pace.

The excellent working relationships we have with these stakeholders allowed us to get straight on with the job in hand, starting with the M&E design phase to verify and work out a non-disruptive way forward which would also meet requirements and time expectation.

Once the design was agreed and completed, University of Edinburgh Estates Projects worked on the costing, programming and contract award with contractor A McKay, who mobilised swiftly to order materials and arrange resource to start on-site.

The process was made easier as we were fortunate to use an M&E team who we have worked with previously at the ACF. This was a great benefit because the contractors' site familiarity and knowledge of safe working practices in a "live" computer environment meant we were able to hit the ground running and commence installation works quickly.

All of these installation works had to happen under the raised computer room floor, which is only 60cm deep and which also has existing "live" power, water, network, and fibre infrastructure.

Access to this void is limited as it is only possible to lift a maximum of six floor tiles in a straight line at any one time to access the void! Not easy when trying to install long chilled water pipework runs.

There was also a need to remove some redundant computer room air conditioning (CRAC) units and relocate and re-engineer the computer room gas suppression bottles and control system to free up valuable computer cabinet floor space. This intrusive work was difficult to accomplish in a "live" computer room. Works also included the replacement of floor and ceiling once the equipment had been removed.

Although this was a complex and sometimes difficult job, all CR1 project works were completed on time and on budget and, more importantly, without any negative or detrimental effect on the other "live" systems or networks.

We are now starting to populate CR1 to house the additional equipment. This is in advance of any major Computer Room 4 project works which would be required to support future anticipated growth.

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

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

HPE five-cell ARCS cabinets at the Advanced Computing Facility.

CRAC units removed and gas suppression bottles relocated to single location, freeing up floor space.

ARCS cells underfloor pipework.

Before: large CRAC units and gas bottle occupying valuable floor space.

"This complex and difficult job was completed on time and on budget."

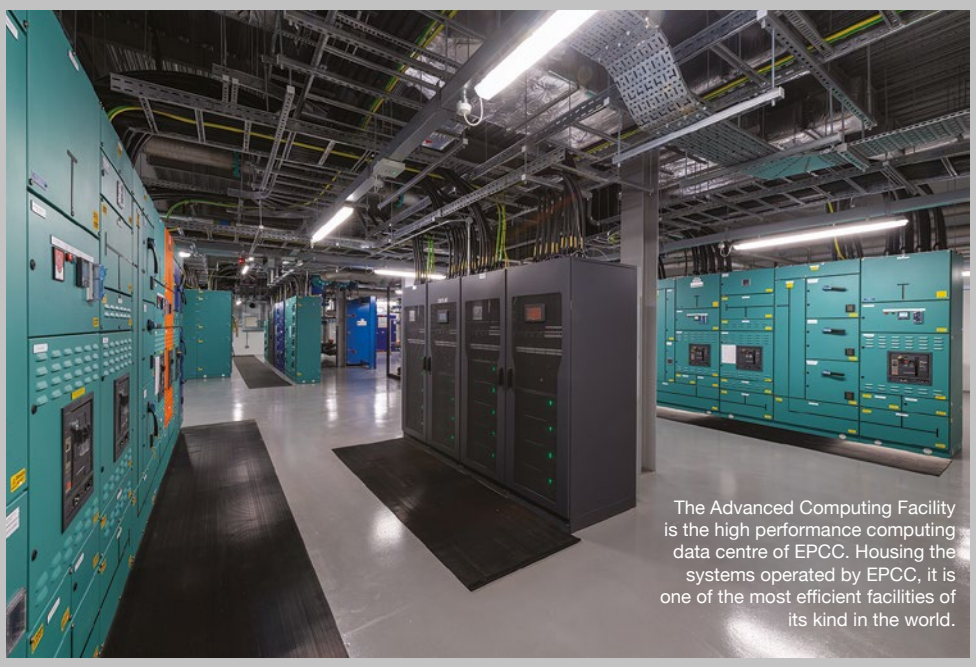

The Advanced Computing Facility is the high performance computing data centre of EPCC. Housing the systems operated by EPCC, it is one of the most efficient facilities of its kind in the world.

Better software means better research

EPCC is a founder member of the Software Sustainability Institute (SSI), the first organisation in the world dedicated to improving software in research. SSI was founded on the premise that helping individuals and institutions understand the vital role that software plays in research would accelerate progress in every field of scientific and academic endeavour.

Collaborations Workshop 2024 (CW24)

CW24 will take place as a hybrid event online and at the University of Warwick from April 30 – May 2. The event, which brings together researchers, developers, innovators, managers, funders, publishers, policy makers, leaders and educators to explore best practices and the future of research software, has three themes.

Environmental sustainability: The environmental impact of hardware and code, training and education about environmentally sustainable behaviour, and activism highlighting potential improvements to existing infrastructure, processes and systems.

AI/ML tools for science: including benefits and limitations of using AI/ML tools to write your code, ethical considerations when using these tools, and training and education around the usage and impact of AI/ML tools.

Citizen science: including designing participatory research projects, training and education around citizen science, and the role of research software in citizen science projects.

CW brings together all members of the research software community to explore and discuss specific themes. It is also an ideal opportunity to form new collaborations and to discuss topics proposed by attendees.

RSE Conference 2024 (RSECon24)

RSECon24, which will take place in Newcastle in September, is the UK's Annual Conference for Research Software Engineering. The event, the first and largest in its field, provides an unparalleled opportunity to network with Research Software Engineers (RSEs) – the technical decision makers in research – from around the world.

EOSC EVERSE

Director Neil Chue Hong, Project Officer Giacomo Peru, and Project Manager Kirsty Pringle will participate in the European EVERSE project representing EPCC and the University of Edinburgh. The project started in spring 2024 and will run for 36 months with 18 partners across ten countries.

Coordinated by the Institute of Applied Biosciences of the Centre for Research and Technology Hellas, EVERSE aims to create a framework for research software and code excellence, collaboratively designed and championed by the research communities across the five European Open Science Cloud Science Clusters and national Research Software Expertise Centres, in pursuit of building a European Network of Research Software Quality and setting the foundations of a future Virtual Institute for Research Software Excellence.

SSI runs various activities and programmes designed for researchers, research software engineers, and virtually anyone working with research software from all domains and career stages.

We offer two annual events: Collaborations Workshop and the Research Software Camps, as well as ad-hoc smaller events and workshops throughout the year. See our website for details.

Building on its leadership of the SSI and work in FAIR-IMPACT on software metrics and metadata, the team at EPCC will focus on developing and implementing processes and tools that support the assessment and verification of code quality, based on established best practices and standards across scientific communities, as well as building a collaborative community for evaluating, verifying, and improving the quality of scientific software and code, and collecting and consolidating best practices and standards from across scientific communities on research software.

Report on "Software and skills for research computing in the UK"

Published earlier this year, this report provides an analysis of the transformative role of software and the individuals behind it in reshaping the research landscape and delivers crucial insights into the challenges and opportunities that lie ahead for research computing in the UK.

The study was funded through the UKRI Digital Research Infrastructure (DRI) programme and will contribute to the understanding of the evolving landscape of DRI in the UK. It includes a comprehensive landscape analysis comparing the UK to current international approaches, a survey of over 400 researchers/research-adjacent staff, eight focus groups, and DRI stakeholder interviews. It was undertaken between December 2021 and October 2022 by the SSI, with researchers based at the University of Edinburgh and the University of Southampton, in collaboration with Dr Michelle Barker, Director of the Research Software Alliance.

Recommendations in the report are aimed at maintaining the UK's world-leading position in research software, including the recognition of diverse roles, professionalising RSE roles, and developing a national roadmap for research computing infrastructure.

Denis Barclay, EPCC
d.barclay@epcc.ed.ac.uk

CW24:
https://www.software.ac.uk/workshop/collaborations-workshop-2024-cw24

RSECon24:
https://rsecon24.society-rse.org

EOSC EVERSE:
https://www.software.ac.uk/news/eosc-everse-ssi-team-represent-university-edinburgh

"Software and skills for research computing in the UK" report:
https://www.software.ac.uk/news/new-report-software-and-skills-research-computing-uk

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

Photographer: Ludovic Farine.

Delivering public engagement

EPCC has been actively involved in public engagement for many years. Beginning with colleagues giving talks at local schools, today this has grown into a coordinated programme of activities.

Our public engagement activities aim to demonstrate the societal benefits of EPCC's wide range of high performance computing (HPC) platforms and data facilities to the general public. We look to explain the purpose of these facilities and show that they are a valuable use of public funds.

A key factor in their development has been funding from ARCHER2, the UK's National HPC service, and from its predecessor ARCHER. This funding has allowed us to build a coordinated set of activities that target specific audiences and offer benefit across the UK, enthusing young people about STEM (science, technology, engineering and mathematics) and showcasing the potential rewards a career in computational science can bring.

Activities

We regularly attend large science festivals covering different areas within the UK. For example, the Big Bang Fair in Birmingham, New Scientist Live in London, and the Edinburgh Science Festival.

At these events we offer a range of fun, engaging hands-on activities for participants. For example, Wee Archie, our suitcase-sized supercomputer, allows participants to use a parallel computer, learn about the type of science running on our supercomputers and understand more about performance and power (see opposite page).

We provide a set of logic puzzles to encourage participants to apply their problem-solving skills and introduce the skills required to develop algorithms.

We have lots of different practical hands-on activities designed to introduce the concepts of parallelism. For example, a fast and furious parallel ball-sort activity has always been popular.

We also offer activities designed around understanding how computers themselves work. For example, our "make your own binary keyring or bracelet" illustrates how binary numbers are used to store data on computers.

Work experience

We offer work placements for schoolchildren, providing opportunities for students to experience the world of work and explore careers in computational science. Students usually visit EPCC for one to three weeks and undertake a small project while working closely with a member of staff. They also have the opportunity to visit the Advanced Computing Facility to see the computer and data systems and plant rooms we operate.

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

Our most recent activity is a virtual reality tour of the ARCHER2 system, allowing members of the public to see what a supercomputer really looks like.

Online resources

We have a wealth of online resources targeted at both schoolchildren and lifelong learners. In addition to worksheets and games, our timeline of computing facilities at EPCC provides an insight into how supercomputers have evolved over the years.

Please visit Discover and Learn, our website devoted to outreach:
https://discover.epcc.ed.ac.uk

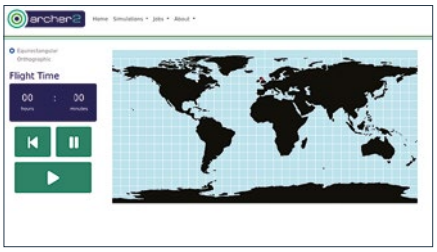

These images show the enhanced Wee Archie aerofoil demo, which has been developed to give an example of the power of computer simulation.

Wee Archie turns eight!

Our bite-sized supercomputer has outlasted many larger systems and is still going strong.

The Wee Archie Raspberry Pi cluster has become a staple of our public engagement activities and has travelled all over the world. As with other supercomputers, the system has gone through various upgrades and is currently looking towards a significant redevelopment to incorporate different types of processor and offer enhanced monitoring.

When launched in 2016, the system was a great success with young people. The design, which was created and built by the University's science outreach group, FUSION, in collaboration with EPCC, has proved key, with an eye-catching appearance that allows all its boards and components to be seen. The LEDs on each Pi show how much work is being done as a simulation progresses.

The earliest demo involved simulating races between on-screen dinosaurs. Users modified the structure of dinosaurs' muscles and joints, altering their ability to run. Wee Archie then tested each of the configurations quickly, and presented the results as an on-screen race. This linked directly to science being run on the ARCHER system, the then UK National Supercomputing Service.

Our most recent demo for Wee Archie involves the user creating their own aerofoil, experimenting with different angles, curvatures, and thicknesses and seeing the impact on take-off and fuel consumption for an aeroplane. As well as calculating the lift and drag properties of each design, the game makes a small animation, demonstrating whether an aeroplane with this design could take off (or not!) and how far it could fly on a full tank of fuel.

Wee Archie is currently undergoing a significant upgrade, with the aim of introducing small graphics cards into the setup, which will allow users to see the impact of different types of hardware on performance. There will also be enhanced monitoring on the system, enabling users to look closely at performance and scaling across multiple cores. Finally, as with other larger systems, Wee Archie will be able to report on power consumption.

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

Wee Archie:
https://discover.epcc.ed.ac.uk/hardware/activities/2023/01/10/wee-archie/

Study HPC with us

epcc

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.

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 2024, and 22nd by QS World University Rankings 2024.

"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

"Studying at EPCC came with significant benefits. I gained access to the UK Tier-1 and Tier-2 supercomputers ARCHER2 and Cirrus. Learning from renowned researchers in HPC, the dedicated study area for MSc students at EPCC, and being part of a global community were enriching experiences."

Jaffery Irudayasamy
MSc in HPC with Data Science graduate 2023, now EPCC employee

www.epcc.ed.ac.uk/msc