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

epcc news

Issue 90 AUTUMN 2021

Data and HPC expertise from Scotland to the world

From our Director

As some, but not all, normality returns to our lives, it's been a great pleasure over the past few months to begin to meet EPCC staff and even welcome some visitors to the Bayes Centre in person.

While we have not yet returned en masse to our offices, around 50% of the team have returned for a few days and on most days the office is now 25% full. It's been a great pleasure to have meetings in person again – online meetings are not a real substitute. Of course our data centre, the Advanced Computing Facility (ACF), has been open throughout the pandemic and I want to take the opportunity to thank everyone who has made this possible, and more importantly safe!

As I hope you will see in this edition of EPCC News, we continue to go from strength to strength with around 115 staff and some new people still to join. This growth has been driven by both the increase in services we are operating – ARCHER2, Cirrus, Tesseract, Tursa, the EIDF etc – but also an increase in the number of data science projects we are involved in.

The original promise of the City Deal's Data Driven Innovation programme is leading to a big growth in such projects. We're also continuing to prepare for the future with new services such as the Cerebras CS-1 system, and more exploratory systems such as the FPGA testbed for the ExCALIBUR project. Most recently we've run a small initial Request For Information to help us prepare the ACF's Computer Room 4 for an Exascale supercomputer service in the future – more on that in future newsletters.

I hope you enjoy this issue and any feedback is always welcome.

Professor Mark Parsons
EPCC Director

Mark Parsons
EPCC Director
m.parsons@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.

Data security: external audit success

We recently passed our ISO 9001 quality external audit and ISO 27001 information security re-certification external audits. We have also passed our external audit for the Digital Economy Act for the National Safe Haven, to continue as an approved data processor.

We rate both the delivery of services for our customers and the secure storage, management and processing of their data as the highest level of importance, and the certifications are very much secondary to our aim of applying best practice. We are however very pleased when we do pass with flying colours!

Anne Whiting, EPCC
a.whiting@epcc.ed.ac.uk

Meet the Advanced Computing Facility interns

Scott Gordon (upper pic)
I started a work placement with EPCC in August 2021. I currently work full-time at the Advanced Computing Facility (ACF) alongside studying at Napier University where I'm pursuing a degree in Cybersecurity & Forensics.

In my short time so far with EPCC I have assisted with a large variety of things, such as helping to diagnose hardware faults with servers, network migrations, and general day-to-day datacentre operations. As I interact with contractors and engineers on-site I have also picked up bits and pieces about HVAC, having witnessed (and been incredibly impressed with) the sheer scale of the infrastructure required to cool and power the ACF and what's contained inside.

One of my main professional interests is process improvement and DevOps practices. Coming from an IT helpdesk background I am aware that time-consuming and repetitive processes can be a real pain point, and I feel that it's incredibly important that a person's job can be done efficiently, not only for the cost and time-saving benefits but also for their overall happiness and wellbeing. This lines up nicely with the way that EPCC is constantly evolving, and I look forward to being part of it.

Kristiyan Tanev
I am a third-year networking student at Napier University. Working as a new ACF intern is a never-ending rollercoaster of emotions, each day unique with new experiences, connections and opportunities to apply my knowledge.

A typical day starts with getting my working space ready. Generally we then ensure the computer rooms are unlocked as required, and contractors are given access.

After our daily team meeting I continue working on my tasks, which include installing and configuring software and hardware systems, designing computing solutions, and applying knowledge in the maintaining of systems. I want to work for EPCC at the ACF because it is one of the few organisations that combines technical knowledge and working with supercomputers with teamwork and a friendly environment, as well as ensuring the quality of work does not drop and every user is satisfied.

I want to gain the skills and experience necessary to become head of networking so that I can build my career and open my own successful company. I have been working and putting all the effort towards my goals and during my internship I have the opportunity to put my knowledge into practice.

The evolution of the ACF
https://edin.ac/3kR27qt

ACF expansion
https://edin.ac/3uoKsJP

Working with industry

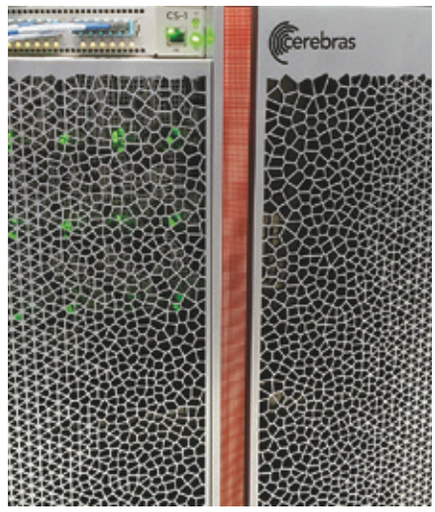

Cerebras CS-1, the world's fastest AI computer, hosted at EPCC's Advanced Computing Facility.

Partnership with industry is a core principle of EPCC and one of the reasons the Centre has continued to grow and develop since its foundation in 1990.

Thomas Blyth, EPCC
t.blyth@epcc.ed.ac.uk

Our industrial collaborations come in a variety of forms: from short-term engagements that are directly funded by industry to longer term research collaborations that include a blend of private and public funding.

Short-term industry engagements are often characterised by working at a high technology readiness level (TRL) and typically include business critical activities that bring rewarding challenges around delivery, pace, and expectation. EPCC's Applications Group structure provides an excellent pathway to manage and deliver these types of projects. It also underpins one of our key advantages which is that we do not need to recruit specifically to resource particular projects.

New collaborations

This year continues to be a busy one for us with several industrial engagements continuing and new projects underway.

We are starting a new collaboration with Danu Robotics (see article opposite), supporting the development of its revolutionary robotic system for recycling and waste management which will significantly increase recycling efficiency.

We are also completing the next phase of work with Bo-Create. This local technology SME in the property & construction sector is now based in the Bayes Centre, the University of Edinburgh's innovation hub for data science and artificial intelligence (and also home to EPCC). We are providing support and expertise to Bo-Create's new product release and additional R&D activities.

We often support industry with respect to cloud architecture and optimising products in this space. One such new collaboration is with Optic Earth, an SME using machine learning to extract value from geophysical data, enabling its customers to make better & more sustainable decisions in the energy and mining sectors. EPCC is providing cloud architecture and data expertise to support Optic Earth's continued development and growth.

HPC on demand

We are continuing to expand our on-demand access programmes for industrial users of world class HPC infrastructure, and with new infrastructure coming online, notably ARCHER2, EIDF and Cerebras CS-1, we expect increased activity in this area. New project collaborations have started across a wide variety of sectors including energy, pharma and digital technology. EPCC's pay-per-use simulation-as-a-service collaborations with ENGYS and Gexcon also continue to thrive, with several new user projects starting this year.

With a return to the Bayes Centre on the horizon we are looking forward to further exciting industrial collaborations.

Our on-demand HPC service

Targeted at engineers and scientists solving complex simulation and modelling problems, our on-demand HPC service gives direct, secure access to EPCC's HPC platforms to deliver the highest levels of performance.

Unlike cloud-based services, there are no inefficient virtualisation techniques. The highest levels of data security are provided, and the service is administered directly by the user. The service is fully supported by an integrated help desk and EPCC staff are available to help with usage problems.

To discuss any of our services for industry, contact our Commercial Manager, Thomas Blyth:
t.blyth@epcc.ed.ac.uk
+44 (0) 131 651 3461

Applying machine learning to the recycling industry

The world generates 2 billion tonnes of domestic solid waste annually but less than ten per cent is recycled because the current recycling process is extremely inefficient. We're working with a start-up that is developing a revolutionary robotic system to significantly increase recycling efficiency.

Danu Robotics is an Edinburgh-based clean tech company that is looking to improve the efficiency of recycling through automation. Currently, most recycling centres and plants operate by having human pickers sort through recycled goods on conveyor belts – any recyclable of the wrong category (eg any soda can in the paper recycling) is manually removed from the conveyor belt and either sorted into the correct category or disposed of if not recyclable. This process has many disadvantages: sorting recycling is a thankless and repetitive task; sorters are prone to error (especially late in a shift); and, as sorting rates are determined by the number of people working the conveyor belt, sorting can quickly become a bottleneck in the recycling process.

Here is where Danu Robotics steps in by offering a robotic sorting system that can swiftly seek and sort the recyclables on the belt, with the aim of removing anything that shouldn't be there. To do this, the system needs to know what to pick, and where to pick it. For this, Danu Robotics is developing a machine learning solution – before deployment of the robotic sorting system, Danu Robotics will collect video footage of the conveyor belt with which to train a machine learning agent. This agent will learn to categorise every item on the belt. Once the initial training is complete, this machine learning algorithm will direct the robotic sorting system to sort out waste items efficiently and effectively.

So where does EPCC fit into all of this? Xiaoyan Ma, the founder and CEO of Danu Robotics, is an alumna of the EPCC MSc programme in HPC with Data Science. Early in 2021, Xiaoyan got in touch with us regarding a small project for which she would like help from EPCC. With an Interface Innovation Voucher (matched in time by EPCC) plus support from the EuroCC@UK project (see page 10), we had the seed of an ongoing collaboration.

We have been providing assistance where we can, specifically helping to develop an algorithm that will ensure communication between the machine learning agent and the robotic sorting system goes smoothly, and that the agent does not try to overwork the robotic pickers!

We at EPCC are both proud to see our former students actively engaging to change the world for the better and excited that they want to involve us in bringing forth this change. Plans are already in motion for this collaboration to continue into 2022.

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

Working with small businesses and supporting the development and success of new companies in HPC and data is a key driver for EPCC. We partner with a variety of start-ups and SMEs on a wide range of activities including R&D, innovation, and new product development. A key benefit of collaborating with us is access to our large group of HPC specialist staff combined with the world class computing infrastructure that we host.

DANU ROBOTICS
Green. For Life.

https://danurobotics.com/

Research Data Scotland: harnessing the power of public sector datasets

John Swinney, Deputy First Minister and Cabinet Secretary for COVID Recovery, reflects on how Scotland's economic, social and environmental wellbeing can be improved by maximising use of data for research in the public good.

If we have learned one thing from the past 18 months, as we grappled with the many challenges of COVID-19, it is that nothing is solved in a neat compartment.

The urgency of the pandemic forced us to think differently to ensure services we rely on could adapt and continue at pace, resulting in more efficient ways of working across the public, private and third sectors.

Data has been key to effective decision-making throughout the pandemic. Scotland has a wealth of public sector data but traditionally it has been locked away in lots of individual systems, across multiple organisations, in formats that are difficult to access or compare.

That's why we established a data taskforce in June 2020 to support organisations with evidence-based policy and operational decisions in relation to COVID-19. The taskforce, which subsequently became the Scottish COVID-19 Intelligence Network, brought key datasets on Covid testing and vaccinations together with vital events, hospital care, the 2011 census and schools, in a secure environment to support collaborative research and decision-making.

Research Data Scotland (RDS) has now been established to cement this approach. A Programme for Government commitment – it's mission is to improve the health, economic, social and environmental wellbeing of Scotland and attract investment into Scotland by creating a supportive environment for ethical and secure data-driven research.

As a collaboration between the Scottish Government, public bodies, and Scotland's leading academic institutions, RDS brings together expertise, resources and capabilities from a range of existing data-led programmes across the public sector to help make data-driven research in the public good easier – with a commitment to always handle data legally, ensuring privacy is maintained and that research is carried out transparently.

Better connections between organisations, and the data they hold, will be instrumental to finding innovative solutions for issues such as sustainable employment, financial security for families and low income households, and the wellbeing and mental health of children and young people.

John Swinney
Deputy First Minister of Scotland

Research Data Scotland was established to improve economic, social and environmental wellbeing in Scotland by enabling access to and linkage of data about people, places and businesses for research in the public good.

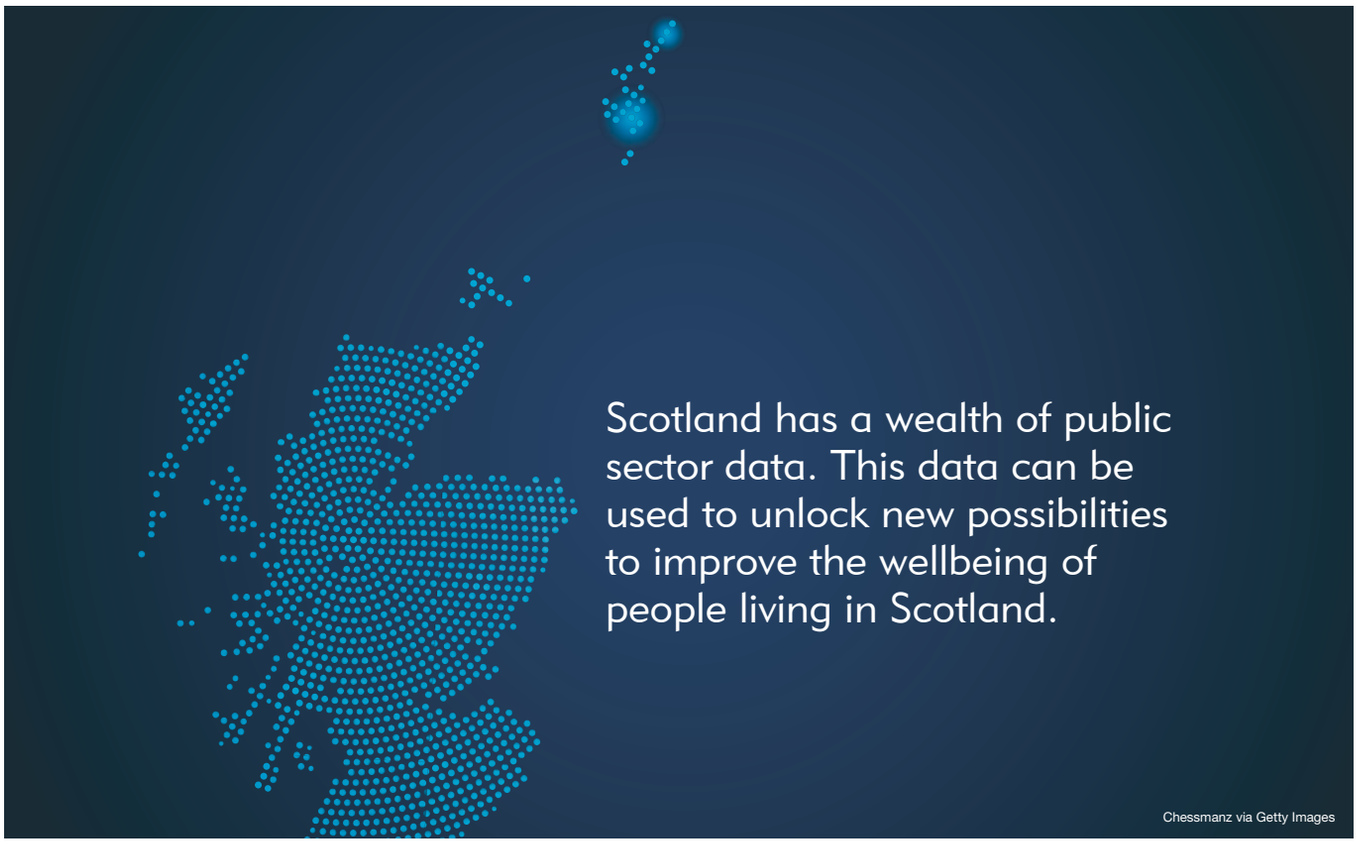

Importantly, it will also inform ongoing work to ensure public services are person-centred – that they are targeted at the individuals they are to support.

Linked data research is already illustrating the practical benefits of this. RDS evaluation of a pilot project led by NHS Fife and Shelter Scotland to coordinate health and housing services for patients admitted to hospital found it reduced the number of people being discharged with no fixed address or into street homelessness, and resulted in better use of healthcare resources, with an average cost saving of £2,422 per person.

That is just one example of the many initiatives underway right now being supported by RDS' researcher support service, which is there to guide people and organisations through what data is available and the information governance needed to access it.

The pandemic has shown us the benefits of doing data right, doing data together, and doing data better. Faster analysis and evidence gathering to support decision-making means better outcomes can be delivered more quickly.

Now more than ever it is vital we maintain this momentum to deliver a fair and sustainable recovery for the people of Scotland.

Research Data Scotland's Seven Principles

1. RDS will only enable access to data for research that is for the public good

2. RDS will ensure that researchers and RDS staff can only access data once an individual's personal identity has been removed

3. RDS will ensure that all data about people, businesses or places is always kept in a controlled and secured environment

4. RDS will only create a dataset if requested for a research study or programme in the public good

5. All income generated by RDS will be re-invested into services to help researchers continue to access data

6. Firms that access public data for the public good through RDS will share any commercial benefits back into public services

7. RDS will be transparent about the data it provides access to and how it is being used for public benefit.

This article first appeared on www.scotsman.com

Research Data Scotland

RDS is a collaboration between the Scottish Government, Scotland's leading academic institutions and public bodies which aim to facilitate insight from data to support the public good. Partners include:

- The Scottish Government
- Public Health Scotland
- The Scottish Centre for Administrative Data Research
- National Records of Scotland
- EPCC
- Administrative Data Research UK
- Grampian, Tayside, Lothian and Glasgow Regional Safe Havens

To find out more, contact Ian Green at Research Data Scotland:
ian.green@gov.scot
https://researchdata.scot

What makes a Safe Haven safe?

For many years EPCC has operated the technical end of the Scottish National Safe Haven (NSH), the Trusted Research Environment which underpins nationwide public-benefit research with sensitive data in Scotland. Recently we have compiled a toolkit that enables us to create additional Safe Haven services as part of our work in building the Edinburgh International Data Facility.

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

Our Safe Haven services follow two sets of defining principles: the Five Safes model [1] designed by the Office for National Statistics, and the Scottish Government Charter for Safe Havens [2] . Both sets of guidelines highlight something which is absolutely central: a Safe Haven service comes in two separate parts, a secure system (run by one organisation, us, perhaps) and an information governance function (IG, performed by a different organisation)—a "gate" and a "gatekeeper".

EPCC provides the secure systems, but we don't mark our own homework when it comes to data security. Information governance does that. For the Scottish NSH service, we provide the system and the eDRIS team within Public Health Scotland provide the IG. They are the service owners, we are the operators.

Providing Safe Havens for enabling research with public data is all about information risk management and assurance, and is essentially a records management exercise with sufficient detail logged at appropriate stages of the data life cycle. The Five Safes model provides a way to balance risks of one kind (eg technical) against risks of another (eg "bad actors"). The model breaks down the decisions surrounding data access and use into five related but separate principles, usually framed as questions, like this.

Safe people

Access to sensitive data is usually restricted to "approved researchers"; how do we define an "approved researcher"?

What are the appropriate skills and credentials required for researchers to access, manage and process your data? What about the system's administrators?

What evidence is needed to demonstrate these skills, and do they need to be refreshed or renewed on a regular basis?

Safe projects

Who approves the kinds of research that can and can't be done with each data set?

What constraints have the data providers placed on data use?

Are there specific contractual obligations regarding data use?

What has been done to ensure that the proposed project use of the data complies with constraints?

Are there specific requirements on transparency, citation, publishing and access to aggregations of the data?

Can management of the data be delegated to the project or must it be actively managed by a dedicated independent data administrator?

Where is this all written down and approved!?

Studies regularly show that the public trust (university) researchers to "do the right thing" with public data but only if we do it safely, explain what we're doing and don't get complacent. Models like Five Safes provide a framework for thinking about data safety and doing the right thing—and remind us that we can never take our eyes off the ball.

MCCAIG via Getty Images

Safe settings

What are the security concerns and specific requirements for storing and processing the data? What do the data providers need?

Are access restriction controls and specific login processes required?

Can the data potentially be downloaded to a local device or uploaded to a remote location? For instance, upload and download of data from the Scottish NSH by researchers is not only not allowed, it is technically disabled.

Safe data

Are the data accessible by the researchers proportionate to the approved project requirement, in line with GDPR requirements?

Where personal data are involved, have the data been treated so as to minimise risk of disclosure of any individual's information?

Safe outputs

How will output from the project research be published and tracked?

What approvals and controls are required for project output to be extracted from the safe environment?

How will FOI requests regarding the project be received, tracked and fulfilled?

Is a public or private website or a Virtual Data Room [3] required?

Safe Havens at EPCC

The systems we operate as part of our Safe Haven services are safe settings in the sense above, and the procedures are defined. These procedures also support elements of the other "safes": eg safe people, for instance (who has administrator permission to which part of the system and have they been suitably trained?). But the security of the system alone is not enough to implement a Safe Haven service, it must be supported by additional IG procedures (the gatekeeper). IG is typically responsible for safe data, safe projects and safe outputs, and the user-side of safe people.

The Scottish Government Charter for Safe Havens provides a useful template that builds on the Five Safes model, covering division of responsibilities, Service Provider authority, data safety, data sensitivity reduction, data separation, staff training and managed collaborations. This last point refers to collaborations with commercial entities—firms with ideas for new products and services which could offer public benefit.

Scotland doesn't sell public data, but neither does it want to lock out innovators from creating valuable products. Managed collaborations are best, with approved researchers taking the lead and working closely with firms to maintain the public trust that is absolutely vital to any kind of research with public data.

The Edinburgh International Data Facility, which is hosted and managed by EPCC, provides Safe Haven services to health and government users. Safe Haven services can also be created for organisations wishing to host and govern access to their data assets in a highly secure environment.

[1] The Five Safes model was devised in 2003 by Felix Ritchie at the UK Office for National Statistics (ONS).

[2] A Charter for Safe Havens in Scotland, 2015, ISBN: 978-1-78544-496-8 (web only), https://www.gov.scot/publications/charter-safe-havens-scotland-handling-unconsented-data-national-health-service-patient-records-support-research-statistics/pages/1/

[3] https://www.investopedia.com/terms/v/virtual-data-room-vdr.asp(link is external) for a definition of a Virtual Data Room.

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

EuroCC@UK: the UK's National Competence Centre in HPC, HPDA and AI

EPCC leads EuroCC@UK in collaboration with STFC Hartree. Both centres are drawing on their complementary expertise to deliver the programme's goals.

The UK is an advanced country in the context of HPC, with a long history of infrastructure provision and computational science research by academia and industry. Both EPCC and the Hartree Centre have significant experience in supporting and collaborating with academia and industry. Since it was established in 1990, EPCC has collaborated with UK science and industry, as well as providing high-quality accessible training courses. The Hartree Centre is funded to work primarily with industry and has undertaken more than 160 collaborative projects with companies since its foundation in 2012, with its training activities focused on meeting the needs of UK industry.

During EuroCC@UK's first year, both centres have been active in a range of areas including training, industry engagement, facilitating scientific and technical expertise transfer, and competence mapping. Here is a selection of our activities.

Training and skills development

EPCC is a recognised leader in high-performance computing (HPC) training and innovative training delivery mechanisms, while the Hartree Centre has a strong focus on training for industry. Over the past year the team at EPCC has investigated innovative training methodologies and accessibility: see article on p22.

The Hartree Centre has also been exploring different ways of engaging industry in training activities, for example business webinars for SMEs on the topics of high-performance data analytics and artificial intelligence techniques.

Many HPC cycles are now consumed by users of pre-compiled research software packages, which requires a different approach from standard HPC training. In response, EPCC has delivered a set of courses relating to package usage to gain feedback on best practice in this area.

Technology transfer and business development

The partners have focused on ensuring that EuroCC@UK has the correct mechanisms in place to support technology transfer in the UK, in terms of HPC know-how, applications, tools and mechanisms. Key components of

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

EuroCC will raise participating countries to a common high level in high-performance computing (HPC), high-performance data analytics (HPDA), and artificial intelligence (AI).

National Competence Centres will be responsible for surveying and documenting the core HPC, HPDA, and AI activities and expertise in each participating country. The ultimate goal is to make HPC available to different users from science, industry, public administration, and wider society.

shulz via Getty Images

this approach are being implemented, providing a supportive environment for organisations to explore the latest digital technologies and skills, develop proofs-of-concept, and apply them to industry challenges.

Collaboration with industry

EPCC and Hartree both already have a range of mechanisms for collaborating with industry: from simple cycle sales, consultancy, and development projects, to full-scale collaborations. Two examples of first-year outcomes are given below.

An awareness campaign has been started by EPCC, aligned with its normal commercial outreach activities, to promote the opportunities that EuroCC@UK can provide to UK industry. We are also assisting companies in preparing bids for further funding – a key objective of the programme

EPCC has also initiated a pilot study under EuroCC@UK, with additional support by the Scottish Funding Council, which is investigating the application of novel machine learning techniques to the recycling of packaging materials (see p4). A second pilot study has been confirmed and we have a pipeline of interested companies with whom we are discussing future pilots.

Access to scientific and technical expertise

EuroCC@UK has worked to facilitate access to scientific and technical expertise and knowledge pools. Examples include:

- Active contribution by EPCC to the UK and global HPC community, working to promote an open, cohesive, and collaborative UK HPC community that facilitates access to technical expertise and knowledge pools.
- Participation by EPCC in a comprehensive review of the skills required by Research Software Engineers (RSEs) in HPC and their future education and training needs, focusing on the role of RSEs as we prepare for the arrival of Exascale supercomputers.

An initial competence mapping, carried out by EuroCC@UK, to gain an overview of the existing HPC, Big Data and AI competencies in the UK.

EuroCC@UK is jointly co-ordinated by EPCC, the Science and Technology Facilities Council (STFC) Hartree Centre. It is funded by the European Commission's Horizon 2020 as part of the EuroHPC Joint Undertaking and through national funding programmes within the partner countries.

EuroCC website
www.eurocc-access.eu

Enabling machine learning for exascale simulations

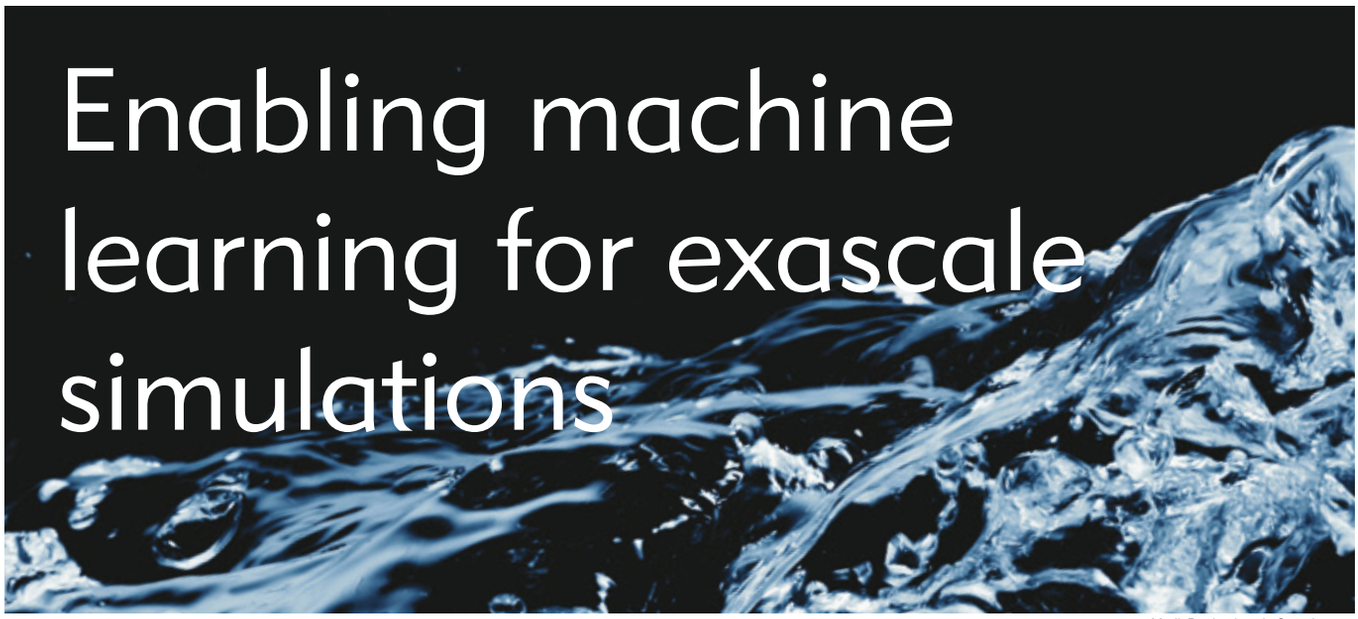

MediaProduction via Getty Images

SiMLInt is an interface that increases the efficiency of machine learning (ML) techniques when solving large-scale physical simulations by consuming fewer resources without sacrificing precision.

In a recent breakthrough, Kochkov et al. at Google Research demonstrated the successful use of Convolutional Neural Networks (CNNs) to solve systems of partial differential equations (PDEs) that describe the evolution of complex physical systems, such as turbulent flows, while using considerably less computational resource.

The underlying idea is that, if trained properly, the CNN can learn basic rules that define flows and use this knowledge to describe a new flow and provide better approximation in terms of coefficients of the PDEs describing it. However, given the high pace with which the field of ML evolves, the domain experts would need to develop a second specialisation to set up the models, both in terms of design as well as the tools and libraries to be used.

We are building on the domain expertise of the School of Maths to help us identify the most robust and general ways of enabling ML algorithms to assist in solving systems of PDEs. Our focus is on simulations that are considered high priority and which are known to require extensive resources in order to be solved, such as turbulent flows or fusion modelling.

Given EPCC's experience with software development and knowledge transfer, we are well placed to navigate the "zoo of tools" the current ML landscape offers and consider their strengths when run on a variety of HPC machines.

Based on this detailed comparison of various ML libraries and toolkits, with focus on CNNs and their performance in terms of accuracy as well as resource consumption, in different hardware set-ups including some of the testbeds such as Cerebras, we will present the user with an optimal configuration no matter where they choose to run their simulation.

Moreover, SiMLInt's interface will ensure that basic ML operations can be implemented straightforwardly, and the accompanying explanatory and training materials will ensure that the domain specialists can set up their ML models correctly and efficiently, without having to build up a full ML expertise.

The project is at an early stage and will use community engagement in the form of workshops and other knowledge exchange activities to inform SiMLInt's development to ensure the result is useful for the scientific community.

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

SiMLInt is being developed in a close collaboration between EPCC and The School of Mathematics at the University of Edinburgh under the ExCALIBUR funding, which is funded from the Strategic Priorities Fund and is dedicated to enhancing high priority computer codes and algorithms in line with latest progress/development in hardware, software and algorithmic tools.

Paper: Machine learning–accelerated computational fluid dynamics
www.pnas.org/content/118/21/e2101784118.short

Evaluating parallel AI training using the Cirrus HPC facility

Doug Carson, Illuminate
contact@oneilluminate.com

At Illuminate, we provide data engineering expertise to help the analyst access timely, insightful and accurate data and make their job easier. Central to this workflow are the AI models to triage and highlight items of interest which must be continually trained to adapt to new threats and eliminate any biases that could impact accuracy. Accelerating this decision process therefore requires that AI models are deployed and operated at maximum efficiency.

In recent years there has been a convergence between HPC facilities and hyperscale data centres towards heterogeneous systems that are radically different from homogeneous systems found in enterprise server farms. The facilities use high-efficiency interconnects, GPU accelerators, S3 storage and batch schedulers. A workload that is particularly amenable to heterogenous systems is AI training.

HPC facilities can provide a springboard to evaluate system configurations for AI training which can then be smoothly migrated to hyperscale facilities for production pipelines. This avoids the time and cost associated with repeatedly building and debugging systems in the cloud or on bare-metal servers. To investigate this, we leveraged the expertise of our frequent partners at EPCC in Edinburgh and its Cirrus HPC facility.

The model we selected for benchmarking was a fully convolutional network for semantic segmentation of images with over seven million parameters coded in TensorFlow.

One of the trade-offs we explored was TensorFlow distributed training vs the Horovod framework. On Cirrus, Horovod achieved 80% efficiency in a 4 node/16 GPU configuration and offered a ten per cent boost over distributed TensorFlow. This can be attributed to the efficiency of the Cirrus interconnect which was used by Horovod. The expertise from EPCC ensured the code modifications and launch scripts were straightforward.

Our company's customer commitment to accelerate informed decisions has many aspects and this study is an illustrative example. During the study we accelerated efficiency by 90% by leveraging our system engineering expertise to make technology choices and optimise them on Cirrus.

Overall, we accelerated the training process by 24 times, using EPCC facilities and expertise to create a highly parallel AI training system configuration. These learnings will be applied to ensure that our data pipelines continue to run as efficiently as possible and at peak performance.

Cirrus is a state-of-the-art HPC system designed to solve computational, simulation, modelling, and data challenges. Users can run their own codes or access a range of commercial software tools on an on-demand basis. The system has been designed to support a wide range of industry uses and EPCC can provide both access to HPC resources and consultancy.

To access Cirrus and EPCC expertise contact Thomas Blyth: t.blyth@epcc.ed.ac.uk

Cirrus
https://edin.ac/3umKGRI

Illuminate
https://oneilluminate.com

Developing a unique FPGA testbed for UK researchers

EPCC is hosting a new FPGA testbed and running it in collaboration with UCL and the University of Warwick. Funded under the ExCALIBUR H&ES programme, the system will be available for developers to experiment with novel hardware for their workloads to enable them to understand the role that these technologies might play in future exascale machines.

The testbed will act as a "one-stop-shop" where developers can run their code on FPGAs, with all the tooling and necessary licences already installed and available to lower the barrier to entry as much as possible. In the few months since my previous EPCC News article (see link below), we have installed our first set of FPGAs into the testbed and made them available to users.

There have been numerous exciting activities undertaken on the testbed, for example accelerating the UK Met Office's MONC atmospheric model across both Xilinx and Intel FPGAs. The Met Office is one of the major use cases of ExCALIBUR, and so it was natural to select its high resolution atmospheric MONC model as a test case. An advection kernel of this model was selected and ported to both Xilinx Alveo U280 and Intel Stratix-10 FPGAs. This illustrates a major benefit of the testbed, where not only are users able to explore the algorithmic transformations necessary for their code to run effectively on FPGAs, but also able to compare the technologies of different vendors to understand the consequences of different choices.

Interestingly the different vendor offerings seemed to be suited for different aspects of the workload in this instance. At a single kernel level, the Intel toolchain seemed able to perform more automatic optimisation and thus performed best. However, the disadvantage of this was that the programmer had less explicit control, which then became more important as the number of kernels was scaled up, with the Xilinx toolchain ultimately being able to accommodate more kernels at a higher clock frequency across the chip.

The optimised FPGA versions were then compared against a 24-core Xeon Platinum Cascade Lake CPU, and V100 GPU. Both the FPGAs and GPU significantly out-performed the CPU, however the GPU was far more challenging to beat, ultimately winning out on performance. However, when considering power draw and power efficiency, which is another important dimension for exascale computing, the significantly reduced power draw of the FPGAs, especially the Xilinx Alveos, resulted in significant wins here.

This is just one early success story of the testbed. In the short time that it has been available users have explored numerous HPC simulation codes and optimised existing FPGA libraries. We are now moving into the second phase of the testbed, where we will install Xilinx's Versal FPGAs and look to further increase the user base by running training courses.

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

The FPGA testbed will be a first step towards building a future community and ecosystem around the role of FPGAs in HPC, data science, AI, and machine learning workloads in the UK.

It will be made publicly available and will form a unique resource within UK academic computing.

More information about the testbed can be found at https://fpga.epcc.ed.ac.uk

Read Nick Brown's previous article https://edin.ac/3Be6YYK

ExCALIBUR programme https://excalibur.ac.uk

AMD Instinct GPUs: an introduction

Regular technical training is essential to ensure EPCC provides the best service to our partners. A recent example provided by AMD is described below.

EPCC's technical staff engage in a wide range of HPC, HPDA, and AI-related projects, all underpinned by an intimate knowledge of the hardware and software technologies that make up the underlying computing ecosystem.

It is therefore critical that our knowledge of this constantly evolving landscape is current and so, in addition to direct involvement with projects that explore novel computing technologies, EPCC staff regularly undertake technical training.

One notable recent HPC trend has been the procurement and upcoming availability to researchers of some of the largest HPC machines (pre-exascale and exascale) in the US and Europe, including the "Frontier" and "El Capitan" systems in the US, and EuroHPC's "LUMI" in the EU. These feature AMD CPUs and GPUs, signalling hardware diversification, especially in HPC accelerators which have been dominated by NVIDIA GPUs.

We therefore welcomed the recent provision of technical training by AMD on its GPU hardware and associated software during August and September, and we continue to benefit from ongoing opportunities for direct consultation.

AMD experts gave us an overview of the silicon architecture and thread-based parallel execution model of AMD's Instinct GPUs and shared valuable insights into the resulting implications for the performance of scientific applications. Programming for AMD GPU acceleration was covered, including HIP (originally designed as AMD's alternative to CUDA) and other programming models such as OpenMP.

The ROCm software stack, which provides compilers, drivers, handles runtime execution and includes high-performance numerical libraries, was demonstrated, including illustrated examples showing usage of the latest capabilities of the rocprof profiler tool for analysing the performance of GPU-accelerated execution. Access to relevant hardware provided remotely by AMD enabled EPCC staff to gain valuable hands-on experience.

Such high-level HPC training supports EPCC's expertise in developing and optimising software to make use of upcoming hardware in research and collaborative partnerships, and helps us to better support and train the scientific research communities that use our and other HPC services.

Arno Proeme, EPCC
a.proeme@epcc.ed.ac.uk

Thanks to this AMD GPU training, EPCC staff who are engaged in BioExcel – the EU-funded Centre of Excellence for Computational Biomolecular Research – will be better positioned to provide guidance to its user community in using CP2K (a quantum chemistry package) to perform biomolecular simulations on AMD GPUs. This will include LUMI, a EuroHPC pan-European pre-exascale supercomputer, once that machine's GPU partition becomes available.

AMD-powered Exascale supercomputers
www.amd.com/en/products/exascale-era

Lumi
www.lumi-supercomputer.eu/may-we-introduce-lumi/

BioExcel
www.bioexcel.eu

A brief history of the ARCHER2 Image and Video Competition

The competition gathers together some of the best images produced by users of the ARCHER2 UK national supercomputing service, which is hosted and managed by EPCC. Clair Barrass, who has been responsible for the smooth running of the competition since its launch in 2014, gives an insider's view.

Clair Barrass, EPCC
c.barrass@epcc.ed.ac.uk

Back in autumn 2014, just a couple of months after I had joined EPCC, I was invited to set up and launch an image competition to showcase the work being done on ARCHER, the UK's national supercomputing service at that time.

Having been immersed in the world of HPC for only a few weeks, I honestly had no idea what to expect. My assumption was that computational research, scientific simulations and crunching of complex equations didn't really equate with aesthetically pleasing images, artwork and things of beauty. My work centred mainly on service users' helpdesk, helping to set up projects and passing on problems with codes that weren't running as hoped to our technical team, and also with administering our training programme, getting new users up and running. I hadn't really seen much of the kinds of outputs the successful work could lead to.

As soon as the entries started coming in, it was clear that my expectations were completely off-base; the images were fascinating, beautiful, and engaging, and the accompanying text genuinely started my learning journey of what the users of ARCHER were achieving with their work.

In order to make the judging both fair and easy to conduct, I created an online gallery of all the entries, anonymised of course, along with the all-important descriptive text. The gallery would allow the judges to review all the entries easily, and after the judging was completed, the gallery was opened to the public so everyone could enjoy them.

The next challenge was to assemble a judging panel. We wanted to include people from a variety of backgrounds; some HPC experts who would appreciate the technical achievements of the work, but also non-HPC experts who would be focus on the impact for a general audience.

The final team for the first competition comprised ten people and included EPCC staff, ARCHER's academic and industrial partner organisations, a graphic design artist who often works with

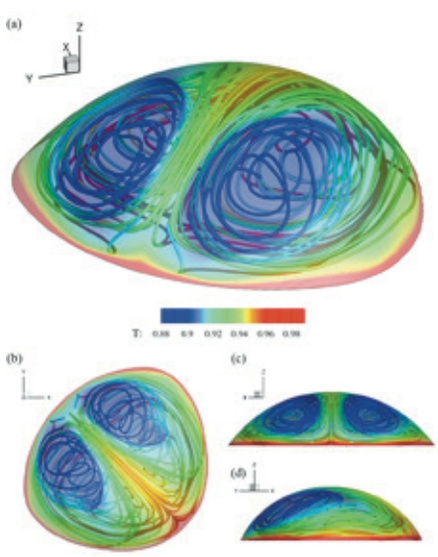

2014 winner: "Streamlines illustrating the temperature distribution and emergence of azimuthal currents within the flow of an irregular 3D liquid droplet undergoing phase change" by Dr Pedro J. Sáenz, Institute for Materials and Processes, University of Edinburgh.

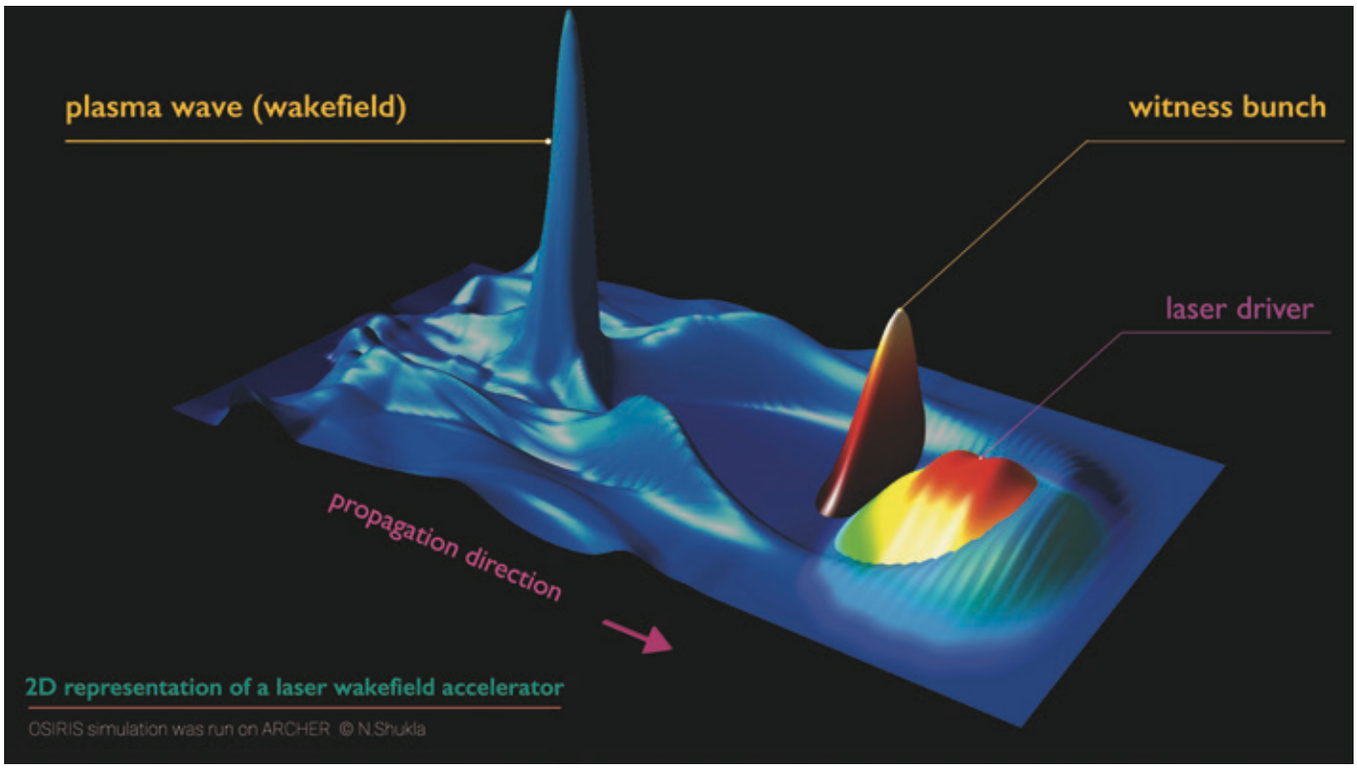

Last year's winner: "Particles surf on plasma waves excited by high-power lasers" by Nitin Shukla, Instituto Superior Technico

EPCC on our Outreach and display materials, and members of the University of Edinburgh Outreach team. None of the judges knew anything about the images they were judging other than seeing the image itself and the accompanying text.

The judges were asked to select their top five images, rated first down to fifth, rating them on how they appealed aesthetically, how well they demonstrated "ARCHER enabling research", and how well they conveyed the importance of ARCHER and HPC to the audience.

My very pleasurable job was then to tally up the judges' marks, and then let the winning entrants know how to collect their prize money.

We have run the competition every year since, and the range of entries and the science areas represented have continued to grow. In 2017 we extended the competition to include video entries alongside the images. We have also always offered a prize to the highest rated entry from an Early Career researcher, as well as the overall highest rated image and video.

The competition is thus a wonderful showcase for the work being done on ARCHER and ARCHER2 and the researchers doing that work, and is also a great opportunity for Early Career researchers to both show off their own work and to see what others are doing, and what opportunities may be open to them.

Looking at the galleries from the last seven years shows just how stunning and engaging the output of HPC can be. The competition has not changed a great deal over the years from an administrative point of view; the numbers of entries grew over the first few years, but we have always had plenty of stunning images and videos to choose from. The judging panel is made up of a mix of some of the same people each year and some new faces, drawing from the same sources.

This year marks the eighth year of the competition, and the first year in which entries may come from work done on ARCHER2, so a whole new exciting range of work and images may be arriving in my inbox soon. I can't wait to turn them into a new gallery to share with everyone once more.

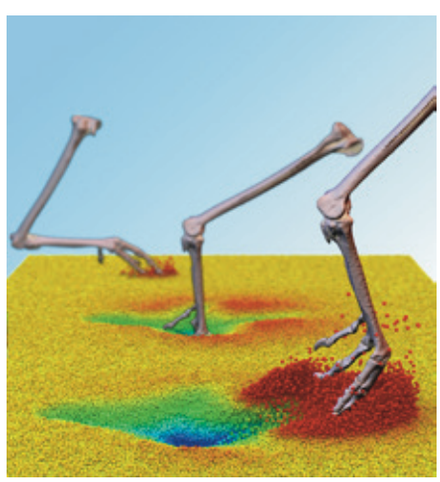

2016 winner: "The birth of a footprint" by Dr Peter Falkingham, Natural Sciences & Psychology, Liverpool John Moores University.

All previous entries can be viewed in our gallery at https://edin.ac/3nLZvfy

Look out for this year's winners in the next issue of EPCC News.

Updates from the Software Sustainability Institute

The Software Sustainability Institute (SSI) facilitates the advancement of software in research by cultivating better, more sustainable, research software to enable world-class research. The Institute is based at the Universities of Edinburgh, Manchester, Oxford and Southampton, and has funding from all seven research councils.

Jacalyn Laird, EPCC
j.laird@epcc.ed.ac.uk

Research Software Camps

This year the Software Sustainability Institute started running Research Software Camps – free online events running over two weeks which explore topics around research software. The first Camp in February focused on helping researchers make their research accessible. The Camp involved keynote talks, workshops, guides and blogs on the website and engaging social media content.

The next Camp, Beyond the Spreadsheet, will run from 1–12 November and explore the uses of spreadsheets in research and first steps into further use of software in research. Sign up for updates on this and future Camps: https://edin.ac/3lS2RLp

Hidden REF awards

The SSI supported the Hidden REF Awards which took place in September, celebrating all research outputs and everyone involved in their creation. The Hidden REF committee will continue to lobby for a broadening of research recognition, and will work on publicising its approach to the assessment of novel research outputs and hidden research roles. You can watch the award ceremony, and hear from all of the winners, on YouTube: https://edin.ac/3kx9zHi

Simon Hettrick, SSI Deputy Director and Chair of the Hidden REF said: "The Hidden REF awards showed that a volunteer-run campaign could recognise everybody who is vital to research. Our next goal is to help organisations adopt our methods so that they can also recognise the full breadth of contributions people make to research."

Research England has invited Simon Hettrick and two other Hidden REF committee members to take part in a scoping meeting about the future of research evaluation.

GitHub citation

Researchers can now easily cite the software they use with GitHub's new built-in citation support, giving proper credit and recognition to those who develop research software. The SSI is proud to have supported the new Citation File Format (CFF).

SSI's Director, Neil Chue Hong, a co-author of CFF and co-chair of the FORCE11 Software Citation Implementation working group, said: "Initiatives like the Citation File Format, and the support for it from GitHub, Zenodo and Zotero will encourage visibility and adoption of software citation by the mainstream of researchers by removing barriers to citing software directly."

Software is fundamental to research. Seven out of ten researchers at fifteen Russell Group universities reported that without it their work would be impossible.

Intermediate Software Carpentry course

The SSI has been developing and piloting a new Intermediate Software Development training course over the past year and a half. The new course provides intermediate research software development skills as a next step to courses like Software Carpentry. It teaches skills like developing software using feature branch workflow in Git, unit testing and Continuous Integration with GitHub's Actions, object-oriented and functional programming, practical software architecture and design, and releasing your code to the world.

The material will be made available to the community for comments via the Carpentries Incubator in October 2021. We expect it to graduate into a more stable stage by the end of the year and be ready for reuse and teaching by the wider community beyond the material creators.

UKRI Innovation scholarship collaborations

The SSI is proud to be involved in three successful bids to the UKRI Innovation Scholars: Data Science Training in Health and Bioscience initiative.

The programme's key objective is to produce training opportunities for researchers in areas ranging from bioinformatics to the social sciences, to give them the self-confidence and skills to manage and analyse their data. Collaborating with these projects will bring together the SSI's extensive training expertise towards improving research software practices within the Health and Bioscience domains.

Read more about the three projects at: https://edin.ac/3zyxkD1

Fellowship programme shortlisting underway

The Institute's Fellowship Programme provides funding for individuals who want to improve how research software is used in their area of work. Each Fellow is awarded £3,000 to spend over fifteen months. This funding can be used for any activities that meet both the Fellow's and the Institute's goals, such as travel to workshops, running training events, nurturing or contributing to communities of practice, collaborating with other Fellows, or for any other activities that relate to improving computational practice or policy. The Programme's next successful applicants will be announced on 13 January 2022.

Better software, better research

The Software Sustainability Institute is a national facility for building better software. We help researchers introduce software into their research or improve the software they already use.

The Institute's vision is to create a world where software is treated as a first-class citizen and is sustainable, enabling better research. Sustainability means that the software used today will be available – and continue to be improved and supported – in the future.

Software Sustainability Institute

Software Sustainability Institute
www.software.ac.uk

Two decades of MSc programmes at EPCC

David Henty, EPCC
d.henty@epcc.ed.ac.uk

Back in the late 1990's we had become aware of a need for postgraduate level training in high-performance computing (HPC) and parallel programming. We already had a good portfolio of material from the courses we delivered as part of the UK national supercomputing services (at the time it was based around the Cray T3D/T3E), and when the Engineering and Physical Sciences Research Council (EPSRC) offered Masters Training Programme grants in 2000, we took the opportunity to build on this material and successfully applied for funding to develop an MSc in HPC.

The biggest challenge in launching a new higher education course lies in getting it off the ground. However, the EPSRC pump-priming funding meant we could start small and not worry too much about initial student numbers. Course overheads are another potential barrier – developing and delivering a new course is as much work for 5 students as for 50 – but we overcame this by initially opening course registration to both MSc students and general HPC participants. Broadening participation gave us time to build the MSc's reputation while covering our costs.

We initially followed a "short and fat" approach, offering courses as three full days a week over six weeks each semester, which suited our external, non-MSc attendees. However this changed once numbers grew and we switched to MSc-only classes, moving to half-day training blocks spread over a semester. As the course became more established we adopted the now standard University of Edinburgh format of two lectures and a practical session every week, and became fully integrated with the University's teaching timetable.

In 2014 the MSc programme was augmented with a new option to specialise in data science. The balance is now split roughly 60:40 between students of the MSc in HPC with Data Science, and the MSc in HPC alone.

Since 2016 we have offered online courses as part of Edinburgh's interdisciplinary DSTI MSc, and since 2020 both of EPCC's MSc in HPC programmes have been made available for part-time online study.

Dissertation projects often focus on real-world problems in collaboration with external partners. Recent examples in academia include spiNNaker boards with the University of Manchester, and molecular dynamics simulations with the University of Nottingham. Our students have also worked with public bodies such as Historic Environment Scotland and the National Library of Scotland, Societies such as the Marine Biology Association, as well as Government organisations. Industrial projects have included Scottish companies as well as international ones such as Xilinx (inventor of FPGAs), Intel, Virgin Money, and the Renault FP1 team.

Left: the class of 2013, including Weronika Filinger (front row, third from left), who is now EPCC's Programmes Director for our online MSc programmes and Chair of Student Programming at SC21, and Larisa Stoltzfus (front row, far left), now an Applications Consultant at EPCC.

Our preparations for the online programmes were invaluable in enabling us to switch to online teaching and student support at the start of the COVID pandemic.

Our courses have changed greatly since the early days with printed handouts entirely replaced by the University's virtual learning environment, and lectures automatically recorded to help students when revising The curriculum has evolved - original courses included Grid computing and high-performance Java - but our core content remains the same: we have always taught fundamental principles such as parallelism, performance programming and good software development techniques. Although hardware has evolved at an incredible rate – the new ARCHER2 service will be almost a million times more powerful than the Cray T3D system from the 1990's – the fundamental concepts of HPC and computational science have stood the test of time.

Interdisciplinarity is an essential element of our MSc programmes, and we encourage our students to work with other groups within the University and with industry. We have established a strong international reputation, which has been reflected in the steady increase in overseas students. Our teaching programmes are a major element of EPCC's activities with around half of our 120 staff actively involved in lecturing or supervising student projects. We aim to limit each year's cohort to about 50 to ensure the quality of each student's experience, particularly during their dissertation. Students benefit from time spent with experts who are working in the field of study, while staff gain from refreshing their knowledge. The regular cycle of the MSc helps us keep all our training materials fresh, with our training courses for the national HPC services benefiting from the MSc and vice-versa

Academic governance is overseen by the School of Informatics at the University of Edinburgh.

Although the HPC landscape has changed substantially since we launched our first MSc programme, our goal remains the same: to produce graduates ready to apply their learning to solve real world problems wherever they are.

ISC Student Cluster Competition

For many years our students have successfully participated in the international ISC Student Cluster Competition, where student teams build small ad-hoc clusters and compete to demonstrate the best performance across a series of benchmarks and applications without surpassing a 3KW power limit for the system.

In 2014 EPCC's team achieved the highest-ever LINPACK score. The pic above shows the winning EPCC team at ISC14.

If you have any questions regarding our MSc programmes, please get in touch:

msc@epcc.ed.ac.uk
+44 (0) 131 651 3398
www.epcc.ed.ac.uk/msc

Investigating accessible training

One of the aims of the EuroCC project is to explore the benefits of innovative training methodologies for HPC. A key component of this is ensuring barriers to learning are minimised.

Stephen Farr, EPCC
s.farr@epcc.ed.ac.uk

Accessibility has two meanings in the context of the EuroCC project (see article on p10).

Firstly, training material must be accessible in terms of how it can be obtained, ie is membership of a specific teaching institution required, is it behind a paywall, can it be downloaded easily, what file types are used, and is it compatible with different types of computers?

Secondly, the training material needs to be accessible to people who are visually impaired, have hearing problems, physical disabilities, or different learning style preferences.

To address the first goal we will serve the training materials up as static webpages hosted on GitHub, this means the webpages are publicly hosted — anyone can visit the site, and no specific login details are needed. Furthermore, the source code and media for the site (markdown files, HTML, images, etc) will be hosted on the corresponding GitHub repository, which is also publicly hosted and freely downloadable. This helps mitigate the problem of privately hosted materials being deleted when old websites are no longer maintained.

Our second goal is also partially addressed by using webpages as the main content delivery mechanism because HTML can be read by screen readers and is easily enlarged to facilitate access by visually impaired people. Further considerations are to ensure that images and diagrams are accompanied by adequate alternative text, and videos have an audio description option. Similarly, to be accessible to people with hearing problems, lecture videos need to have accurate subtitles and audio clips must be transcribed.

We have chosen colour schemes and design layouts that are high contrast, clear, and consistent. The webpages themselves make use of existing template designs such as MkDocs and Software Carpentries which are tried-and-tested using web design best practices. For example, they render appropriately on different sized screens and can be navigated by only a keyboard. To enable different learning styles we are incorporating text articles, video lectures, lecture slides, and practical exercises into the material, being careful to maintain clear navigation and flow of the course content.

The accessible materials will be initially used in a new “Introduction to HPC” course that is currently under construction. Feedback and lessons learned from this course will be filtered and refined to further increase the accessibility of HPC training throughout the UK.

EuroCC is a 33-state European project to establish national Competence Centres in high performance computing (HPC), high performance data analytics (HPDA) and artificial intelligence (AI). In the UK it is jointly coordinated by EPCC and the Science and Technology Facilities Council (STFC) Hartree Centre.

EuroCC project
www.eurocc-access.eu

EuroCC in the UK
www.epcc.ed.ac.uk/research/computing/computing-infrastructures/euroccuk

PRACE Summer of HPC

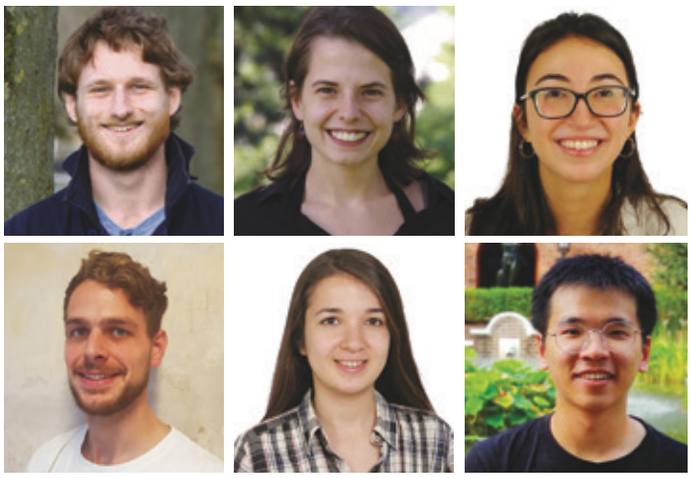

Clockwise from top: Jonas Alexander Eschenfelder, Carla Nicolin Schoder, İrem Okur, Jiahua Zhao, Aybüke Özçelik, and Alejandro Dinkelberg.

Since 2013 the Partnership for Advanced Computing in Europe (PRACE) has run a summer internship programme, the Summer of High Performance Computing (SoHPC), where final year undergraduate or master students undertake high performance computing (HPC) projects at PRACE HPC centres throughout Europe.

Normally students are given basic HPC training before travelling to an HPC centre to work on a project under the guidance of a local expert. However the pandemic meant PRACE SoHPC 2021 was a virtual event.

EPCC proposed three projects and was assigned six students as follows:

Investigating scalability and performance of MPAS atmosphere model (project 2113), managed by Dr Evgenij Belikov, EPCC. This project looked at the performance and scalability of the Model for Prediction Across Scales (MPAS) Atmosphere Model on ARCHER2, attempting to determine its performance characteristics and carrying out some optimisation of the code.

Students: Jonas Alexander Eschenfelder, Geophysics, Imperial College in London and Carla Nicolin Schoder, masters student in Astrophysics, University of Vienna.

Re-engineering and optimizing software for the discovery of gene sets related to disease (project 2114), managed Dr Mario Antonioletti, EPCC, assisted by Dr Pau Navarro, Institute of Genetics and Cancer, University of Edinburgh, and Dr Claudia Cabrera, Centre for Translational Bioinformatics, QMUL. The project looked at understanding and improving the performance of the genomicper R package available through CRAN, which performs pathway analyses to find associations between genes and disease. The project aimed to increase the size of data sets handled by the package and to process existing data sets faster. Students: İrem Okur, Computer Engineering student, Dokuz Eylül University, and Aybüke Özçelik, Electronics and Communication masters student, Engineering Department, Izmir Institute of Technology.

Performance of Parallel Python Programs on ARCHER2 (project 2115), managed by Dr David Henty, EPCC. This project investigated and improved the performance of a computational fluid dynamics of fluid flow in a cavity developed by a SoHPC project last year.

Students: Alejandro Dinkelberg, PhD candidate, University of Limerick, and Jiahua Zhao, geophysics masters student, University of Padova.

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

The 2021 Summer of HPC Programme ran between July and August, with 66 students distributed virtually across 33 projects.

All students have produced a video presentation: https://www.youtube.com/user/SummerofHPC

Project reports will be available here: https://summerofhpc.prace-ri.eu

Image: Paul Dodds

Study HPC with us

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.

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 both Times Higher Education World University Rankings 2021 and QS World University Rankings 2021.

"The MSc in High Performance Computing course had a really good practical focus which allowed me to develop my HPC programming skills and get hands-on experience working on applications using the skills I learned during the MSc on a daily basis."
Holly Judge
MSc student 2018/2019
Now an Applications Developer at EPCC.

"The quality of the courses provided is exceptionally high, all lecturers are the field experts in HPC. Additionally the support that EPCC provides to students is second to none, all staff are extremely helpful and supportive."
Xiaoyan Ma
Graduating MSc class of 2021
Now CEO Danu Robotics.

www.epcc.ed.ac.uk/msc