Issue 97 SUMMER 2025

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

epcc news

Breaking new ground

Largest Cerebras CS-3 cluster in Europe will accelerate AI research and the UK's AI capabilities

From our Deputy Director

Welcome to the Summer 2025 issue of EPCC News.

It occurred to me, whilst reading this issue, how much of what we are involved in is linked to cycles, be those funding cycles, technology cycles, political cycles, or simply the cycles that projects go through, including my own journey at EPCC as I enter my third year here.

This constant of fluctuation, with work progressed through an arc from conception, to design, delivery and review, all lead to a certain level of resilience, in both the system and structures and in the people who enable it all.

In this issue we look back and celebrate the end of the current Cirrus service, which has grown and changed since it was first commissioned in 2016. Over its years of service, it has supported a wealth of research and development for academia and industry.

Looking forward, new opportunities are here now in the form of the Cerebras CS-3 cluster we have deployed – the largest of its kind in Europe (read more opposite) – whilst other opportunities lie on our horizon, in re-establishing our links with peer colleagues in Europe, the delivery of the UK Government AI Opportunities Action Plan and continuing to support advancing science for society, in all its dimensions.

As ever it is an exciting time at EPCC, and I hope you enjoy reading this issue of EPCC News.

Ritchie Somerville
EPCC Deputy Director
r.somerville@epcc.ed.ac.uk

EPCC at ISC25

We will be busy across the entire conference programme. Find us at Booth A10, where you can meet Wee Archie, our mini supercomputer, and find out about the latest developments at the UK's premier supercomputing centre.

See our website for details of all our activities:

https://www.epcc.ed.ac.uk/whats-happening/events/isc-high-performance-2025

Contents

3 **Cerebras CS-3 cluster**
EPCC installs Europe's largest CS-3 cluster

4 **MSc milestone**
Marking 25 years of our MSc programmes

6 **Trusted Research Environments**
Pushing the state of the art

8 **Working with industry**
Digitising 30 million images

10 **Cirrus HPC service**
Evolution of a highly successful resource

12 **Support for science**
Enhancing research software

14 **Research at EPCC**
Advancing imaging algorithms

16 **Women in HPC**
Join us at ISC25!

17 **EPCC HPC Summer School 2025**
Spreading awareness of HPC

18 **HPC careers**
Improving access to skills development

19 **Widening access**
Work opportunities for young people

20 **Growing skills**
Internships and Apprenticeships at EPCC

22 **Software Sustainability Institute**
Update on activities

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.

Cerebras CS-3 cluster launched at EPCC

A groundbreaking Cerebras CS-3 cluster has been installed at our Advanced Computing Facility. Operated by EPCC and part of the Edinburgh International Data Facility (EIDF), the cluster comprises four CS-3s using Cerebras' latest third-generation Wafer Scale Engine processors – the largest commercially available AI chip ever built.

This new service marks a significant milestone in our ongoing commitment to push the boundaries of AI research and development. Leading on from our long partnership with Cerebras and use of their CS-1 and CS-2 systems, the new four-node CS-3 cluster is able to handle compute-intensive workloads at a larger scale, enabling faster training times for large language models, new inference capabilities, and the ability to explore and build novel models and HPC-like accelerated codes via Cerebras' software development kit (SDK).

State-of-the-art architecture

Central to the cluster's power is the Cerebras' WSE-3, which integrates an entire wafer of silicon, containing 900,000 cores and 44GB of on-wafer SRAM, into a single computing unit. Traditional computing approaches typically split silicon into discrete chips, with a single GPU or CPU occupying one die and being connected to memory via a relatively slower interconnect. Cerebras breaks this standard model by harnessing the full surface area of the wafer, drastically increasing the amount of available computing resources, memory bandwidth, and interconnect efficiency. EPCC's adoption of this state-of-the-art architecture reflects our position at the forefront of computing research and enables us to provide our user community with unparalleled access to cutting-edge hardware.

Faster LLM training

One of the major benefits of this system is its ability to train Large Language Models (LLMs) at record-breaking speeds. Researchers and engineers can focus less on orchestrating distributed computing resources and more on refining model architectures and experimenting with novel training strategies. Combined with this is a reduction in energy usage compared to an equivalent GPU resource, working towards our net zero targets by lowering the cost of model training. Rapid iteration is of particular importance in machine learning research, where architectural innovations can dramatically improve the quality and applicability of deep learning models.

Inference engine

We also anticipate researchers using the system as a powerful inference engine, accelerating AI-driven predictions across a range of applications e.g. natural language processing, where it will be used to generate and analyse text at a massive scale. For example EPCC Chancellor's Fellow Rosa Filgueira is investigating the use of the cluster to support high-compute aspects of the Frances AI platform such as:

- Retrieval-augmented generation (RAG) and concept evolution, where large-scale inference is needed to summarise how key terms (like "gravity" or "sugar") evolve across editions of the Encyclopaedia Britannica.

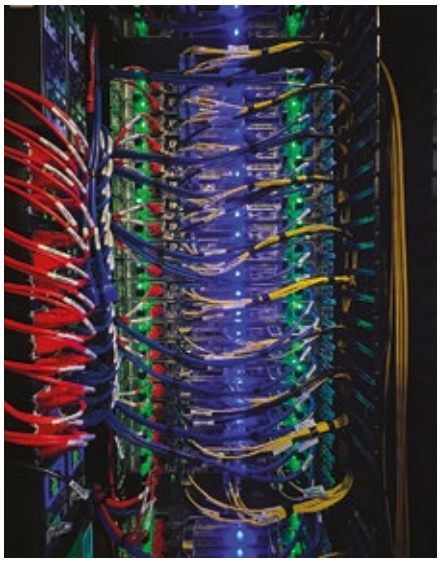

- Off-line named entity recognition (NER) for identifying and linking people, locations, and organisations mentioned in historical texts, so enriching the knowledge graph that underpins Frances.
- Optical character recognition (OCR) cleanup at scale, helping to improve data quality from noisy 18th and 19th Century digitised sources.

Access the CS-3 service

The launch of the Cerebras CS-3 cluster at EPCC underscores a major step forward in EIDF's computing capabilities. We invite interested researchers, engineers and potential collaborators to reach out for more information.

Nick Johnson, EPCC
n.johnson@epcc.ed.ac.uk

Availability

This service is generally available for DDI Programme partners. Please see the EIDF Portal for information on how to apply for access:

https://docs.eidf.ac.uk/access/project

Silver Jubilee of MSc programmes at EPCC

The opening of the new academic year in September 2025 will mark the twenty-fifth run of our on-campus MSc in HPC.

It makes me realise what an old-timer I am that most of our new students won't have been born when the MSc launched back in 2001! Although we started with a small class of half a dozen students, those first graduates have gone on to make an impact, with one being the current Director of Research at EPCC and another leading the Environmental Modelling Research Group at the University of York. With this anniversary approaching, I thought it was a good time to reflect on the history of the MSc and its future.

In the beginning

Back in the late 1990s we became 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, 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.

We initially followed a "short and fat" approach and shared courses with the training programmes for users of our HPC systems. Courses ran for three full days a week spread across each semester, meaning external users could attend an entire course with a single trip to Edinburgh. 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 standard University format of two lectures and a practical session every week, although some of our more practical-based courses have now reverted to a full half-day format.

Expansion and evolution of programmes

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 2020 both of EPCC's MSc in HPC programmes have been made available for part-time online study. The timing was fortuitous as 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. Most recently, in 2024, we launched a new joint programme with Heriot-Watt University in Imaging, Vision and HPC.

Our courses have changed greatly since the early days, with printed handouts entirely replaced by the University's virtual learning environment and lectures now

automatically recorded to help students when revising. Perhaps the biggest change in hardware has been the emergence of GPUs as a mainstream HPC technology, and we have recently developed two new courses on GPU programming and using large GPU clusters for machine learning.

Although hardware has evolved at an incredible rate – the current ARCHER2 is almost a million times more powerful than our Cray T3D system of the 1990s – the fundamental concepts of HPC and computational science have stood the test of time. As a result our core content has remained the same: we have always taught fundamental principles such as parallelism, performance programming and good software development techniques. We believe that this prepares our graduates for success in a wide range of future careers in academia and industry.

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

"The MSc in HPC was transformative for me and I remember my time with EPCC fondly. The experiences I gained on the MSc helped me obtain my PhD in geological modelling at University of Edinburgh, helped me get my first postdoc at Imperial College London in ocean modelling, and are still important today in both research and teaching as an academic at the University of York. Without the understanding of HPC gained during the MSc I would not be doing the research I am today."

Dr Jon Hill,
Senior Lecturer, Environment and Geography, University of York

Masters programmes in High Performance Computing (HPC) and HPC with Data Science

Learn from the experts how to write code that's not just correct, but maintainable, efficient and scalable to the world's most powerful supercomputers. Add a specialism in Data Science to learn how data can be analysed at scale using powerful machine learning and AI algorithms.

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

Trusted Research Environments at EPCC

Here at EPCC we have several exciting projects that are advancing the possibilities of Trusted Research Environments (TREs). Each one is pushing the state of the art in its specific discipline while at the same time developing the ability to federate research across two or more TREs. Such research requires access to pseudo-anonymised sensitive data which must be combined before it can be analysed, requiring advances on both the technical front and in information governance. In addition, these advances must be aligned to enable studies to be performed effectively and efficiently. Here I describe four projects that make use of TREs operated by EPCC.

Financial data service (FINDS)

We're part of a £3million collaboration between Smart Data Foundry and the University of Edinburgh to operate a new financial data service (FINDS), which will enable more researchers to study the financial health of millions of households across the UK by providing secure access to financial behaviours, economic resilience, and regional economic activity.

The new service will be part of a network with five other data services across the country. Providing safe and efficient ways for researchers to access and use the smart data generated through everyday interactions with the digital world, FINDS will provide unprecedented insights into the economic health of the UK through secure access to de-identified banking and finance data from millions of households and businesses. The service will enable a transformation in the UK's understanding of how economic changes and policy interventions affect different communities, helping policymakers design more targeted and effective responses to economic challenges.

https://smartdatafoundry.com/solutions/financial-data-service

SteatoSITE

EPCC's Safe Haven Services host the NHS data that underpins SteatoSITE, a unique Scottish project developing new tests and treatments for people with non-alcoholic steatohepatitis (NASH), the progressive form of non-alcoholic fatty liver disease (NAFLD).

Bringing together data, storage and computing systems, SteatoSITE is an integrated gene-to-patient Data Commons whose users will include researchers, clinicians, and potentially patients and charities too. SteatoSITE contains NHS clinical and pathology data, including images of liver tissue, and has generated RNA sequencing data from tissue provided by the NHS Scotland Biorepository network. As more health data is added, the NASH Data Commons will evolve into a smarter, more comprehensive knowledge system that will be used to make new discoveries to understand and treat this disease better.

https://steatosite.com

Connect4

EPCC leads the Connect4 project, which is working to identify impediments to research that requires data from more than one national TRE. Connect4 has developed a roadmap to federate data access among the devolved UK nations. The roadmap outlines a list of required changes to information governance, describes at a high level a potential shared service model, and promotes the value of metadata enhancements so a researcher can efficiently discover, apply for, and analyse data held at one of the UK's four national TREs through a single access point.

The project partners are working towards more concrete suggestions to improve information governance and towards a shared model of operation, both in the context of federated data access across UK national TREs.

EPCC is leading on a technical activity to enhance metadata towards easier identification of datasets in TREs. We also led a trial project requiring linked data across organisations, with an overall intention of enabling new policy-relevant insights while laying the path for ongoing federation across UK national TREs.

https://www.epcc.ed.ac.uk/whats-happening/articles/connect-4-improving-data-access-between-four-uk-nations

NEURii

NEURii is a two-year research collaboration that is developing scalable digital solutions to transform care for people living with dementia. The data that underlies the project is held in the Scottish National Safe Haven, which provides a secure platform for the research use of NHS electronic data. This resource is commissioned by Public Health Scotland and hosted by the Edinburgh International Data Facility through EPCC.

Pilot projects include SCAN-DAN, which aims to improve dementia diagnosis and treatment by using cutting-edge AI and brain imaging technologies to predict dementia risk. In the NeurEYE project, data scientists and clinical researchers are working with high street opticians for the first time to develop a digital tool that can predict a person's risk of dementia from a routine eye test.

https://www.lifearc.org/project/neurii/

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

EPCC Trusted Research Environments

A Trusted Research Environment is a highly secure data and compute infrastructure where de-identified sensitive data is made available for analysis. In this environment, organisations can create and control their own Safe Haven Services to provide access to sensitive data. Each Safe Haven is isolated from all other Safe Havens. A Safe Haven owner can run independent projects where they control who has access to which projects.

The EPCC Trusted Research Environment is covered by EPCC's ISO27001 accreditation for information security practices. It is self-certified under Cyber Essentials and NHS Digital's Data Security and Protection Toolkit. In addition, the Scottish National Safe Haven Service is accredited under the Digital Economy Act 2017 by the UK Statistics Authority and we operate all our Safe Havens to the same standard.

Contact

To discuss access to our Safe Haven services, please contact us at:

eidf@epcc.ed.ac.uk

Vertigo3d via Getty Images

Turning 30 million old photos into digital gold

A major collaboration is digitising one of the world's largest collections of aerial photos. Builders, the economy and even unsolved crimes could benefit.

National Collection of Aerial Photography

Harnessing digital, data and AI

The National Collection of Aerial Photography (NCAP), based in Edinburgh, holds a collection of over 30 million images from across the world. The collection is a vast treasure trove of high-resolution data capable of tracking climate change, uncovering unexploded bombs, reheating cold cases for the police, and helping planners make better decisions for economy-boosting land development.

NCAP is working with EPCC to unlock the archive's potential, and convert crumbling celluloid into valuable data.

Declassified material

NCAP's archive includes photographs from the pioneers of aerial photography in the 1920s, vast military collections produced during the Second World War and the Cold War, as well as commercial photography of places around the British Isles and Commonwealth. The collection continues to grow as the United Kingdom's Ministry of Defence declassifies more military-related aerial photography.

NCAP was established in the early 1960s to collect and preserve these records and make them as accessible as possible. As the large majority of the collection is in physical 'analogue' format, and some of it in need of careful preservation and storage, NCAP has been running a programme to digitise all the items in its care.

The collection is of interest to many sectors, from documentary film makers and computer software companies, to researchers with an interest in land use change. It has even been used by Police Scotland when reviewing cold cases and in long-running investigations.

Its largest customer base comes from the Explosive Ordnance Disposal sector in Europe. Its Second World War post-strike aerial photographs are vital to pinpoint the entry point of unexploded bombs in the ground and this is essential information for land developers in their risk assessments.

NCAP Head, Allan Williams, says: "The collection contains more than 150,000 boxes of prints, photographic films and associated records, and because much of it is very old there is the risk of deterioration. So we've now got the challenge of extracting the information from the analogue record and converting it into a high resolution digital format."

Supporting NCAP in its digitisation programme is Julien Sindt, Commercial Manager for EPCC, who has been looking at the most data-efficient way of digitally storing such a vast collection of imagery. Julien has helped NCAP optimise its robotic system to automate the digitisation of flat photographs on an array of 100 bed scanners, which can perform the task 24 hours a day.

Julien says: "We developed a computer programme that was able to reduce the time it took to crop a scanned photograph – the process of removing/adjusting the edges of an image to improve its framing – from two minutes per image to a couple of seconds per image, thereby minimising the amount of data that each scanned image takes by 40 per cent.

"It does this by cropping out all the extraneous information collected from the scanning bed, which is not the actual photograph, and this helps to reduce the image size. Typically, there's a drop of 100 megabytes per image, which doesn't sound like much until you realise that NCAP has 30 million images to store. Uncropped, the dataset would require 7,500 TB of storage. Using our programme we can reduce this to 4,000 terabytes of data so effectively we are reducing the cost of storage of these digital assets by 40%."

Decades-long process

Allan says NCAP is making progress but there is a long way to go: "We've currently got roughly 2.8 million images converted to an optimum high resolution format out of 30 million but our collection continues to grow. I think it's now realistic to say that we should be able to convert the whole collection in 15-20 years."

In addition to providing online access to the collection, the digitised photography will also provide further value to researchers, particularly those interested in changes to land use. Allan explains: "You can overlap the images to create a seamless mosaic that covers a large geographical area, and as there was a government programme to systematically resurvey the majority of the Commonwealth, from the 1940s up until the 1980s, this can provide a snapshot of the regions which can have significant research implications, especially for things like deforestation or climate change."

Above: Images show the site of the Bayes Centre, now home to EPCC, in 1960 (left) and 1990 (right).

Images: NCAP/ncap.org

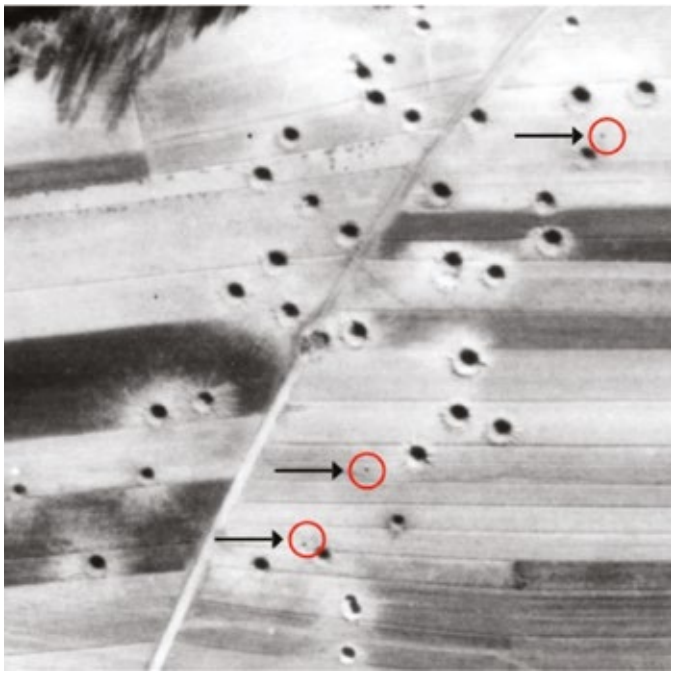

Above: Bomb craters during WW2.

Robotic arm using suction to gently lift print from automated scanning workstation.

Julien is also excited about integrating the data further: “There is an ongoing collaboration between NCAP and our MSc programmes. For example one student has been working on automatically enhancing the digitised images’ clarity and contrast to make features of interest easier to spot, while another student is developing an automatic process to ‘stitch’ overlapping images together.”

The other project is to enhance the 3D processing of these photographs. Since each one is taken at a very slightly different angle as the aircraft travels forward there is the potential to take advantage of the stereoscopic effect to create a 3D image of the surface.

This article originally appeared on the University of Edinburgh’s Edinburgh Impact website:

https://impact.ed.ac.uk

CONTACT

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

To explore how EPCC can support your business goals, please get in touch:
Commercial@epcc.ed.ac.uk

Cirrus: a decade of service

As Cirrus moves into a new phase of operation, we look back on the first chapter of this highly successful service.

Since its launch in 2016, Cirrus has been a flexible, state-of-the-art, high performance computing (HPC) system designed to solve computational, simulation, modelling, and data science challenges. The system was a national service, part of the Tier-2 HPC network, and co-funded by UKRI's EPSRC and EPCC. From 2017 to April 2025, over 2,600 people ran jobs on the system, benefiting over 600 projects. This breaks down to approximately 469 million CPU hours and 2.4 million GPU hours.

Cirrus and EPCC's MSc programmes

Cirrus has been a fantastic workhorse for our MSc programmes. It has served as the primary home for all our students, hosted almost all our taught courses, and has been the target platform for hundreds of MSc dissertations.

A national HPC service is a great teaching asset for both students and staff. Students can access the latest hardware and software, it is well documented, and most of the administrative load of managing student accounts and access to resources is automated via EPCC's SAFE framework. The introduction of GPUs in 2020 allowed us to support a wider range of machine learning projects, and the more recent upgrade to the Lustre filesystem was useful for parallel I/O benchmarking.

During the current academic year, on-campus students alone have already run almost twenty-five thousand individual batch jobs on the Cirrus CPU partition and ten thousand jobs on the GPUs.

Cirrus and research

Cirrus has supported research in areas including physics, chemistry, and engineering. For example the system was used in the design of rotor blades for Ingenuity, the NASA Mars helicopter, which successfully carried out the first ever powered flight on another planet. EPCC's Quantum Applications Group has also employed the system to investigate the use of HPC for efficient simulation of quantum circuits and algorithms. As part of this work, the cuQuantum and QuEST libraries were made available as modules for Cirrus users. Cirrus has also supported research led by the University of Edinburgh to train a neural network to recognise fields from satellite images, increasing the accuracy and decreasing the costs of delineating agricultural field boundaries.

The system was fundamental to integrating HPC into biomedical practice. In 2017, EPCC collaborated with UCL and used Cirrus to develop and host the "HPC for Medics" course. Despite having no command-line experience, the medical students learned the basics, submitted their own batch jobs, ran scaling tests, and used scientific results to determine their hand's microbiome. Now part of several university curricula, it has been adapted to various HPC systems and has trained over 3,000 medical students.

Cirrus and industry

Cirrus has supported a wide range of industry applications in sectors such as automotive, aerospace, energy, oil and gas, general engineering, life sciences, and financial services.

Almost since its launch, Cirrus has acted as a software-as-a-service (SaaS) business solution for ENGYS and its users. ENGYS specialises in the development and provision of computational fluid dynamics (CFD) software products and services. Its HELYX product is a general purpose CFD software solution and its users often require access to expensive HPC hardware to run their simulations. ENGYS partnered with EPCC to provide its clients with access to a cloud-based SaaS that enabled them to run the large-scale high-definition simulations they require for their workflows.

Founded by a graduate of the EPCC MSc programme, Danu Robotics is an Edinburgh-based SME that aims to revolutionise waste recycling by introducing fully automated robotic waste-sorting. This requires highly accurate state-of-the-art computer vision solutions to correctly identify the component parts of recyclables (e.g. plastics, glass, etc). The Cirrus GPU service provided Danu Robotics with the resources to train its computer vision models to the required level of accuracy before testing its prototype in real-life environments.

Our collaboration with the National Collection of Aerial Photography (NCAP) began on Cirrus. With over 30 million historical photographs, NCAP holds one of the largest collections of aerial photographs in the world and has been exploring methods for automating the process of digitising these images. NCAP has worked closely with EPCC to develop a fast, scalable solution for processing newly-scanned images, thereby minimising the digital footprint of this dataset. NCAP and EPCC developed software on Cirrus that enabled NCAP to process images in parallel at a rate significantly higher than they are being digitised. See page 8 to read more about this unique project.

Future plans

We are working to replace the current ICE XA system with a new CPU-based one, which will be available to our users within the Edinburgh and South East region as part of the Edinburgh International Data Facility.

Jo Beech-Brandt, David Henty, Gavin Pringle, Julien Sindt, EPCC

j.beech-brandt@epcc.ed.ac.uk
d.henty@epcc.ed.ac.uk
g.pringle@epcc.ed.ac.uk
j.sindt@epcc.ed.ac.uk

Find out more:

To read more about the Cirrus service, see:

www.cirrus.ac.uk

HPC services at EPCC:

www.epcc.ed.ac.uk/high-performance-computing-services

GeoChemFOAM on ARCHER2: enabling larger meshes and advanced simulations

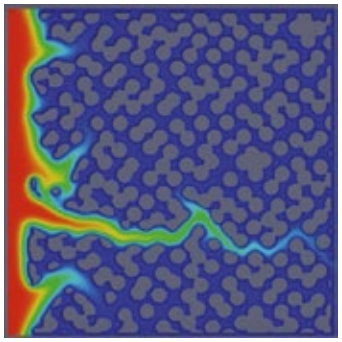

GeoChemFOAM was designed specifically for researching pore-scale processes related to energy transition and net zero, making it a powerful tool for clean energy research. A recent ARCHER2 eCSE project increased the scale of simulations it can perform.

GeoChemFOAM (GCF) provides state-of-the-art solvers for heat transfer, transport, multiphase flow, and chemical reactions in porous materials. The software is continually evolving, with ongoing developments in areas such as nuclear waste disposal, building materials and fuel cell technology.

A recent ARCHER2 eCSE project focused on enhancing GCF's capabilities to support research in large multi-scale materials such as carbonate rocks and fuel cells, where both large domain sizes and high mesh resolution are vital for accurate predictions. The research conducted by the GCF team has been limited by the serial nature of OpenFOAM's pre-processing tools, specifically blockMesh and decomposePar, which are used to create a distributed mesh. These tools must loop over the entire mesh and manage I/O operations, making them inherently serial and difficult to parallelise. This limitation restricts the scale of simulations that can be performed, as the memory of a single node becomes a bottleneck. This eCSE project aimed to overcome these limitations by enabling far greater scales through parallelisation and custom code development.

The project made the latest GeoChemFOAM available to all ARCHER2 users, along with a 64-bit version of OpenFOAM (v2212) to enable larger mesh sizes than previously possible. Comprehensive user documentation is provided on the GCF wiki [1].

Project challenges

The project faced significant challenges in adapting the inherently serial pre-processing tools in OpenFOAM. Various approaches, such as using existing parallel mesh generators or virtual memory, were considered but found to be complex and potentially inefficient. The typical workaround, to decompose a coarser mesh and then refine it, is not possible as the initial mesh must batch the rock image to separate the pore-space from the solid part with enough accuracy. Profiling blockMesh and decomposePar revealed significant variations in execution times due to I/O operations, further complicating the parallelisation efforts. The extensive and intricate source code required careful modification to remove unnecessary computations while ensuring correct output, making the process labour-intensive and reliant on trial-and-error.

Given the challenges with parallelising blockMesh and decomposePar, the team developed a custom Python routine, createMesh.py. This routine leveraged the simplicity of cubic cells and a cuboid mesh, allowing for straightforward division and distribution across multiple processors. Initially designed as a serial routine, createMesh.py was later parallelised using mpi4py, enabling concurrent creation of processor directories and their associated polyMesh files without writing a global polyMesh directory.

Profiling efforts revealed that memory usage and execution times varied wildly in the shared environment of ARCHER2. Despite these inconsistent results, key findings indicated that createMesh.py is I/O-bound, with performance heavily dependent on the I/O subsystem's load.

Several strategies were implemented to reduce the memory footprint and improve the performance and scaling of createMesh.py. These included inlining subprocesses, removing unnecessary arrays, preloading Python libraries, optimising loop invariants, introducing directional padding, and modifying data reading to a layer-by-layer approach. These optimisations significantly improved performance and parallel efficiency, which increased from 40% to 100% on 128 cores.

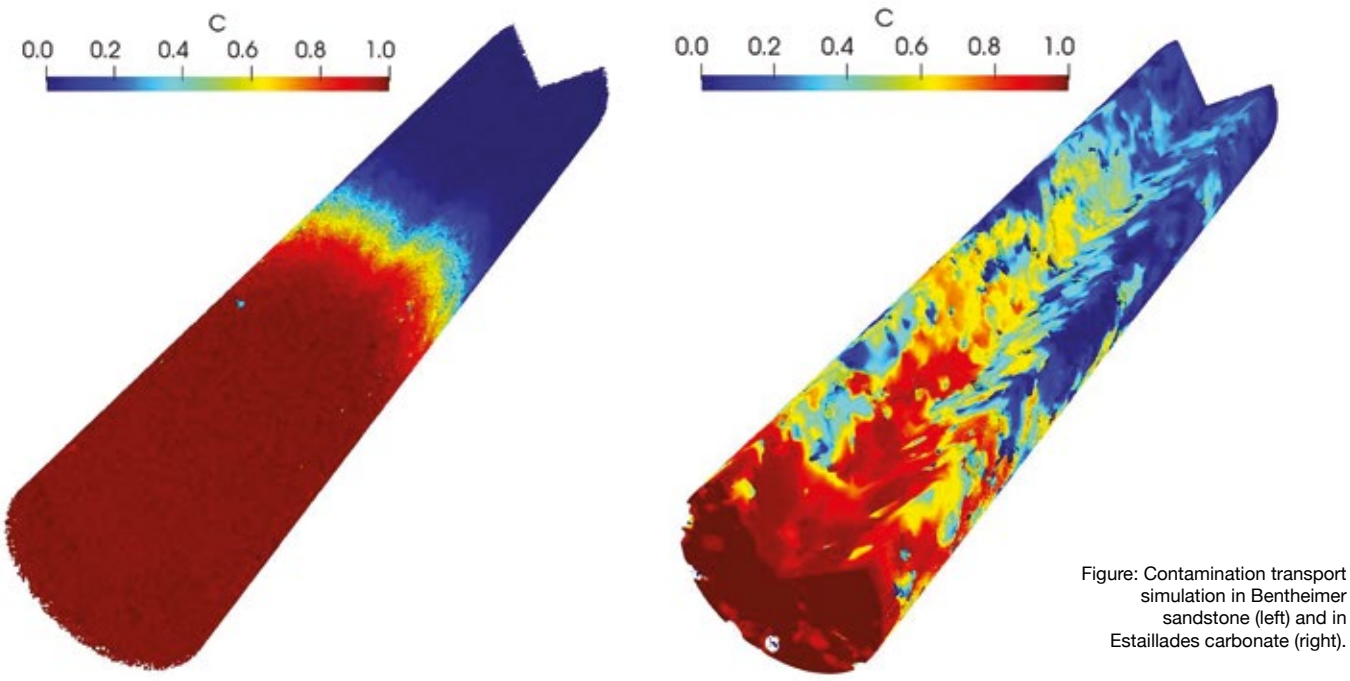

Figure: Contamination transport simulation in Bentheimer sandstone (left) and in Estaillades carbonate (right).

Achievements

The project demonstrated that createMesh.py enables the generation of previously prohibitive mesh sizes, regardless of any efficiency considerations when scaling to hundreds of nodes. The team successfully investigated flow and transport in x-ray computed microtomography scans of two rock samples with unprecedented sizes. For the first sample a Bentheimer sandstone with a simple pore structure, a dataset of 1950x1950x10800 voxels at a resolution of 6 microns, was obtained from the Digital Rock Portal online repository. The second sample, a carbonate rock with a complex pore structure and multiscale features, was represented by a smaller dataset of 1202x1236x6000 voxels at a resolution of 4 microns, obtained from the British Geological Survey online repository.

A complete 3D mesh was generated for each sample, consisting of 50 billion and 20 billion grid blocks using 540 and 80 nodes respectively. Flow and transport simulations were performed using GCF solvers, completing the entire process in under 20 hours for both cases.

The eCSE project has enabled vastly larger meshes by making the memory of multiple nodes instantly available, facilitating simulations on larger datasets. This advancement leads to more sophisticated research opportunities for GCF users and the broader OpenFOAM community, especially those on the road to Exascale computing. The successful installation of the latest version of a 64-bit GeoChemFOAM, along with the development of createMesh.py, have significantly enhanced the capabilities available to users, paving the way for groundbreaking research and innovation.

[1] GeoChemFoam on ARCHER2: https://github.com/GeoChemFoam/GeoChemFoam/wiki/GeoChemFoam-on-ARCHER2

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

Julien Maes,
Institute for GeoEnergy Engineering,
Heriot-Watt University
J.Maes@hw.ac.uk

Hannah Menke,
Institute for GeoEnergy Engineering,
Heriot-Watt University
H.Menke@hw.ac.uk

GeoChemFOAM

GeoChemFOAM is an advanced open source pore-scale numerical simulator based on OpenFOAM. It was developed in the DigiPorFlow Group at Heriot-Watt University. Founded in 2019 by Dr Julien Maes and Dr Hannah Menke, the GeoChemFOAM team collaborates with various partners, including EPCC.

ARCHER2 eCSE programme

The ARCHER2 eCSE programme supports the development of research software for use on the ARCHER2 UK National Supercomputing Service, which is operated by EPCC. Project reports can be viewed on the ARCHER2 website:

https://www.archer2.ac.uk

EIRA: improving generated images for astronomy and medical imaging

The end of 2024 saw the EIRA (Extreme-scale precision Imaging in Radio Astronomy) project drawing to a close. EIRA was a three-year research endeavour focused on advancing imaging algorithms and their associated software implementations for radio astronomy and medical imaging, using a combination of machine learning and optimisation algorithms to significantly increase the resolution and precision of generated images.

In collaboration with Professor Yves Wiaux's group from Heriot-Watt University, EIRA built on expertise in applying optimisation approaches to image reconstruction. These optimisation algorithms define objective functions as a sum of a data-fidelity term and a regularisation term, promoting a given generic signal model. This approach had previously been demonstrated to generate images with higher dynamic range (difference between the brightest and faintest parts of the image) and precision (accuracy of features within the image) than traditional approaches to image generation, enabling fainter signals to be accurately reconstructed and creating more robust images. Ultimately this facilitates new science on existing radio telescopes as well as potentially providing new images from older datasets.

However, these optimisation approaches are computationally expensive compared to the simpler imaging algorithms routinely used nowadays. Therefore, EIRA focused on maintaining this imaging quality while reducing the computational cost of such advanced image construction approaches. This is essential not only to reduce the cost and time required to create images from collected data, but also to support new telescopes and medical scanners, which will generate significantly more data and therefore enable much larger images in the near future. An example of the challenging type of instrument for which we designed these imaging approaches is the Square Kilometre Array (SKA), which when finished will be capable of imaging the sky at much higher resolution, with much higher sensitivity, than current instruments over wide fields of view. In this context, wide-band (multi-frequency) image cubes will exhibit rich structure and can reach sizes between 1 Terabyte and 1 Petabyte, while the associated data size required to reconstruct the image will reach the Exabyte scale.

One of the key approaches used in EIRA was the replacement of the regularisation term in the optimisation algorithm, essentially an image denoising step, with a machine learning based approach. As this denoising step can be trained purely on images, we can create a ground truth dataset (a set of known images) and then an associated noisy image dataset (those images with defined noise added to them) that can be used to train the AI approach independently from the imaging domain being considered (ie independent of the radio astronomy or medical imaging data). This enables the creation of AI approaches to denoising that are robust in their image creation even in a domain such as radio astronomy, where the correct answer (or generated image) for a given set of measurements of the sky is not known.

EIRA used images from optical telescopes as the ground truth of the types of images that are of interest in the types of imaging under consideration, as well as images from MRI scanners (an imaging method that uses similar reconstruction techniques to radio telescopes). The resultant imaging approach, AI for responsible imaging (AIRI), gives the same imaging quality as prior optimisation approaches but at half the computational cost (once the AIRI denoiser has been trained), significantly improving the overall image reconstruction process.

Alongside the learned denoiser approach, we also developed domain decomposition methods to enable the parallelisation of image generation across larger amounts of computational resources, thereby enabling image reconstruction to be undertaken quickly, an important requirement when data is collected in large volumes or images are required immediately (as is the case for medical imaging approaches). Instrument data is decomposed into blocks, and the output image cube is decomposed into small, regular, overlapping 3D facets. Facet-specific regularisation terms and block-specific data-fidelity terms can then be computed in parallel using proximal splitting optimisation methods, enabling more computing resources to be used to speed up image generation.

A reconstruction of the ASKAP telescope Early Science and Evolutionary Map of the Universe Pilot survey observation SB9442–35, including the "dancing ghosts", using the advanced optimisation approaches developed in EIRA.

Read more about dancing ghosts at:
https://academic.oup.com/mnras/article/523/2/1933/7148138

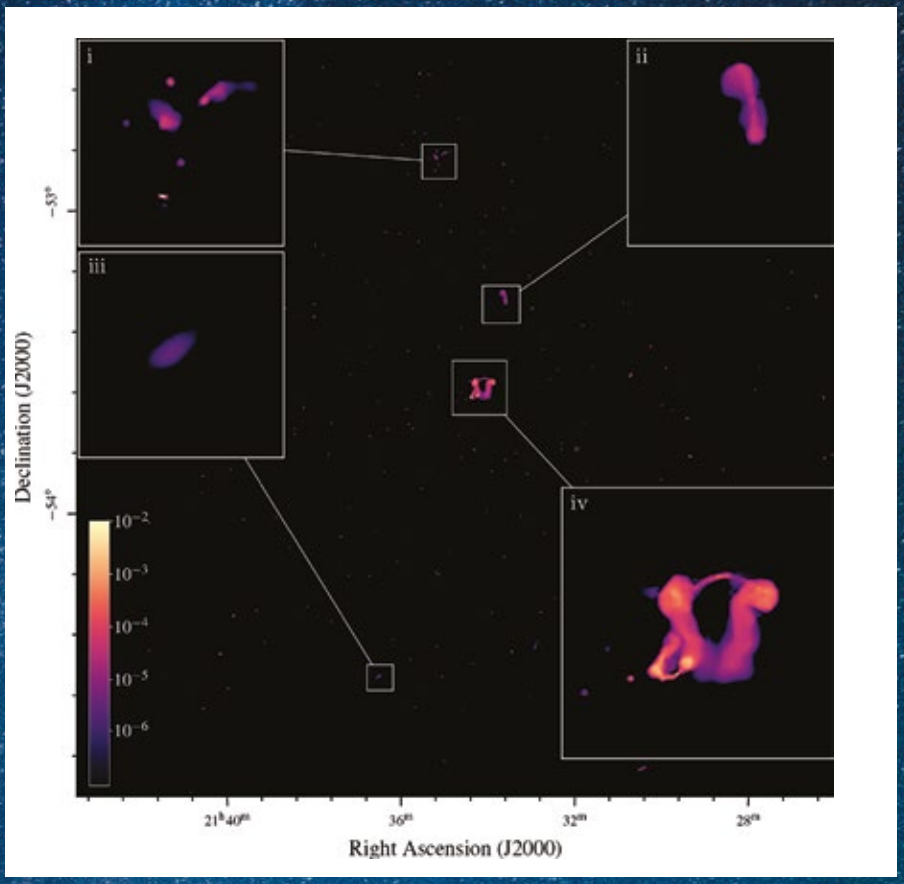

Future applications

The developed approaches can also be applied in the medical imaging domain, particularly for MRI image reconstructions. PhD students at EPCC and Heriot-Watt University are exploring further optimisations and applications of the AIRI approach, either to hyper-spectral images (images constructed from a range of frequency data) or to medical imaging challenges. There is also active work exploring newer machine learning techniques to investigate whether more of the optimisation algorithm can be sped up while maintaining image quality and robustness.

The tools developed within EIRA and follow-on projects can help generate images for use by clinicians, and a key aim is to ensure these images are as detailed and robust as possible. The use of optimisation approaches can enable uncertainty quantification on generated image features (another area EIRA has researched), with the AI approaches employed by EIRA used to improve image quality in a closely controlled and defined manner. Such funded research is key in ensuring our use of AI for science is grounded in reality and ethically justifiable.

EIRA was funded by EPSRC (EP/T028351/1).

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

CONTACT

AIRI:
https://arxiv.org/abs/2312.07137

To learn more about research at EPCC see:
https://www.epcc.ed.ac.uk/research

real444 via Getty Images

Connecting communities and celebrating diversity: WHPC at ISC25

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

All of us at Women In High Performance Computing (WHPC) are excited to continue our collaboration with the ISC High Performance conference series! Let's dive into where you can find us at ISC.

Tech Talks

This year, under the overarching theme "Connecting the Dots", we are focusing on "Connecting Communities". Throughout the evening, we will showcase a dynamic mix of invited and lightning talks, highlighting the many ways our community collaborates and grows together. Thanks to the generosity of ISC25, this event is open to all Conference and Exhibition Pass holders attending ISC25 in person.

Diversity Day

Come and celebrate diversity with the WHPC community! This initiative is all about supporting the visibility of women and other under-represented groups in HPC. This year, Diversity Day is on the 11th of June 2025, and we're excited to announce that several supporters will be distributing ISC25 WHPC swag leading up to the day. For more details on when and where to grab yours, check out our website. On Diversity Day, wear your WHPC swag from this year or previous years, or simply show your support for diversity in your own way. Let's celebrate together!

Posters

As part of our partnership with ISC, the WHPC Posters are officially integrated into the main ISC Poster Programme. 2025 marks the second year of this collaboration, and we were thrilled to double the number of posters we were able to accept, thanks to an overwhelming number of exceptional submissions.

Join us in celebrating our sixteen early-career poster authors at the Poster Pitch Lightning Talks and Poster Networking Reception. Hear from the next generation of HPC talent, explore their research, and network over drinks and nibbles. Supported by ISC and our sponsors, these events are open to Exhibition Pass holders and above.

Birds of a Feather session

Co-organised by the Central European, MAR, and JuWinHPC WHPC Chapters, join us for the "Super(computing) Heroes" Birds of a Feather session! This popular event returns as part of our main programme and offers an opportunity to meet influential female HPC "superheroes" from academia, research labs, and industry, providing inspiration and career guidance.

This partnership has been made possible thanks to Weronika Filinger (EPCC) and Cristin Merritt (Alces Flight), co-leading since 2024. I have taken on the role of Partner Manager for WHPC at ISC25, ensuring the continued success of this collaboration and an outstanding lineup of WHPC events!

Everyone is welcome to join WHPC!

Become a member, join your local chapter or affiliate, and connect with us at our events. Want to support WHPC? We offer a variety of sponsorship opportunities to suit all organisations. Find out more:
https://womeninhpc.org
Or get in touch:
info@womeninhpc.org

WHPC
WOMEN IN HIGH PERFORMANCE COMPUTING

Where to find us at ISC25

- **Connecting Communities: Tech Talks**, June 10
- **WHPC Posters on Display**, June 10–12
- **Celebrate Diversity Day,** June 11
- **WHPC Poster Pitch**, 13:00–13:15pm, June 11
- **Poster Reception**, 2:15–4:15pm, June 11
- **WHPC Super(computing) Heroes Birds of a Feather**, 10:15–11:15am, June 11

For announcements, more information, and where to find us during ISC25, check out https://womeninhpc.org

EPCC HPC Summer School 2025


Image: Paul Dodds

By the time this issue is printed, the EPCC HPC Summer School 2025 will be just about to start. Running for the last two weeks in June, the summer school is an opportunity for undergraduates at UK universities to learn about all aspects of high performance computing (HPC) including modern software engineering tools, writing efficient code in compiled languages, different models of parallelism and how to use GPU accelerators.

Our HPC Summer School is a residential programme which sees a cohort of around 15 students housed in University accommodation in central Edinburgh and taught in venues in and around the Bayes Centre, where EPCC is based. All accommodation and travel costs are covered by EPCC, and students also receive a daily stipend to cover subsistence and other incidental expenses. We particularly wish to attract students from diverse backgrounds and who represent a range of academic disciplines. In this way we hope that students will learn as much from each other as they do from EPCC staff.

The main aim of our HPC Summer School is to make undergraduates aware of the central importance of computer simulation and parallel computing in all areas of science and industry. We hope that this will encourage them to learn more about the subject in the future, perhaps choosing computational options in their own degree programme or seeing HPC as a potential future academic or industrial career. This does seem to be working as very recently the ARCHER2 helpdesk received a query from one of last year's students, now studying a PhD, asking some interesting questions about visualising their simulations!

We also want students to develop practical skills, so a typical day will see lectures in the morning followed by working on practical examples in the afternoon. Students receive accounts on one of our national HPC systems, giving them access to modern hardware and system software.

This is the third run of the HPC Summer School, and we will keep the same general format as last year, which included a very popular trip to the University's Advanced Computing Facility to see real HPC systems close-up (see lower right photograph).

If the students enjoy their two weeks in Scotland's capital city and go away with a good basic understanding of HPC then the Summer School will have been a success. Even if the students themselves don't choose HPC as a career, we hope that they will tell their friends about their experiences and perhaps encourage fellow undergraduates to learn more about the topic.

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

Education and training at EPCC:
https://www.epcc.ed.ac.uk/education-and-training

Improving access to professional skills development opportunities

The UNIVERSE-HPC project worked to widen access to careers in research software engineering and high performance computing. Its work will be continued by two new projects.

Research software engineers (RSEs) combine professional software engineering expertise with a good understanding of research practices. The role of RSEs in high performance computing (HPC) is critical to ensure sustainable and reproducible research outputs. However, there are still no established entry points or career progression routes for RSEs. Additionally, the increasingly diverse backgrounds of RSE and HPC professionals make training provision and professional development more difficult. Many professionals are forced to discover, develop and progress their skills on the job, which can be challenging and time consuming.

The UNIVERSE-HPC or 'Understanding and Nurturing an Integrated Vision for Education in RSE and HPC' project was funded under the ExCALIBUR programme with the aim of defining training frameworks for RSEs specialising in HPC. During the last three years the project has investigated the RSE skills landscape, outlined a curriculum for an MSc in RSE in HPC, explored the feasibility of the learning pathways approach for individualised professional development, developed and piloted course materials to make RSE and HPC skills more accessible, and led many community-building activities.

The underlying goal was to enable more people from a wide diversity of disciplines and backgrounds to obtain the skills and experience required for a successful RSE career. The project was very successful in engaging with the wider community and illustrating the need to make the training and education ecosystem more findable, accessible, interoperable and reusable (FAIR) to better address the specific requirements of professional skills development. The skills landscape needs to be easier to navigate for professionals to better align their learning objectives with available training opportunities or create more personalised learning pathways.

The work continues

The UNIVERSE-HPC project is coming to an end but the work will be continued across two follow-up projects. Over the next two years, DRIFT will address the training needs of research facilitators and research project managers. At the same time, the UKRI Digital Research Technical Professional Skills Network Plus project (CHARTED) will focus on making the ecosystem more FAIR. Across five work packages the project will:

- Map the skills and role landscapes and improve the FAIRness of training resources
- Develop tools that will make the ecosystem easier to navigate
- Create a Hub to facilitate professional development
- Support community building
- Administer the flexible fund for community-led initiatives.

The overall goal of the project is to better define skills and then connect skills to roles, roles to people, and people to people. We hope this project, over the next four years, will be a big step towards making RSE and HPC roles more widely accessible.

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

Read more about how UNIVERSE-HPC has worked to make a difference to the HPC community:

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

Providing career opportunities to young people and growing the UK's skills base

EPCC believes in the importance of encouraging young people and students to learn about our industry. To achieve this, we offer a range of opportunities from engagement activities at science fairs, through work experience for school children to internships for students.

Science festivals

To reach those in primary and early secondary education, we attend a series of science festivals across the UK, in particular the Big Bang Fair in Birmingham, New Scientist Live in London, and the Edinburgh Science Festival. All these events allow us to offer hands-on activities designed to explain the principles and benefits of supercomputing and large-scale data facilities in an accessible way, as well as encouraging young people to consider careers in STEM.

Work placements

For students in the senior phase of their school education, we offer short work placement opportunities, for example we have hosted students from the Nuffield Research Placement and In2STEM programmes during their summer break. Our aim is to provide opportunities for everyone, including those traditionally under-represented in higher education.

Paid internships

We are also involved in Employ.ed, a University of Edinburgh Careers Service programme that provides undergraduate students with paid internships during the summer months. The primary goal here is to help promote and grow the skills required to develop and operate complex supercomputing and data science systems, while helping young people find their way into careers in our field.

Graduate Apprenticeships

We have offered Graduate Apprentice posts at our Advanced Computing Facility for many years. These apprenticeships allow students to work with us in parallel while completing the final two years of their degree, with one day a week reserved for study leave in addition to gaining credit towards their degree for project work that is completed while working for us.

Post-graduate student placements

Finally we have had post-graduate students, who are paid to work on relevant software engineering projects, visit the centre for three to six months. Again this is aimed at increasing the availability of the skills required to work in our industry while helping young people gain relevant experience.

From school children to post-graduate students, we are looking to engage young people in our industry, provide opportunities and help grow the skills required within our field. See page 20 for current examples.

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

Discover and Learn is an outreach and public engagement website developed by the EPCC and ARCHER2 Outreach team:

https://discover.epcc.ed.ac.uk

New System Administration interns and Graduate Apprentices at EPCC

EPCC is committed to promoting and growing the skills required to develop and operate complex supercomputing and data science systems, and to help people find their way into careers in high performance computing (HPC) systems administration and research infrastructure engineering.

We offer twelve-month Systems Administration placements and two-year Graduate Apprenticeships through our long-standing partnership with Edinburgh Napier University. Here we introduce our latest cohort of placement students and Graduate Apprentices.

Kunjal Sancheti and Richard West started twelve-month placements with EPCC last summer, prior to completing the last two years of their degrees. Raul da Costa started as a Graduate Apprentice at the same time. We were also delighted to welcome back Connor Finlay as a Graduate Apprentice after he completed a twelve-month placement with us the previous year.

Above: left to right, EPCC interns and Graduate Apprentices Raul da Costa, Connor Finlay, Richard West and Kunjal Sancheti alongside the Tier-1 national supercomputer, ARCHER2, at EPCC's Advanced Computing Facility.

Zoe Payne, EPCC
z.payne@epcc.ed.ac.uk

The Advanced Computing Facility is the high performance computing (HPC) data centre of EPCC. It is one of Europe's largest research computing data centres and is acknowledged by the AI and HPC community as a global leading supercomputing site.

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

Below: left to right, Connor Finlay, Kunjal Sancheti, Raul da Costa and Richard West outside EPCC's Advanced Computing Facility.

Kunjal Sancheti
Systems Specialist intern

I'm in my third year at Edinburgh Napier University, studying BEng (Hons) Cybersecurity and Digital Forensics. My time with EPCC has given me the chance to work with cutting-edge technologies and gain hands-on experience in a fast-paced, innovative environment.

I primarily provide support for cloud services and shift-based support by addressing incoming tickets from users who are encountering issues with our systems. I also contribute to ongoing projects that aim to improve user experience and increase platform efficiency. Working on these projects has provided me with deeper insights into both technical challenges and the importance of user-centric design in system management.

Additionally, my involvement in Data Centre Support at the Advanced Computing Facility (ACF) has taught me a lot about physical security and its importance in protecting data. This has broadened my perspective on security, highlighting the importance of safeguarding infrastructure both digitally and physically.

What excites me most about my internship at EPCC is the opportunity to work on innovative projects alongside experts in the field. I'm grateful for the chance to grow my technical skills and contribute to meaningful improvements within the organisation.

Richard West
Systems Specialist intern

I am a third-year Software Engineering student at Edinburgh Napier University. I have been involved in hands-on support for the team, performing system administration work for the on-site supercomputers. This experience has provided me with valuable insight into machine maintenance and has allowed me to observe how a team collaborates to address a wide range of challenges.

I look forward to the rest of my placement and bettering my understanding of the systems by participating in project work. My role has fuelled my passion for high performance computing and its potential applications in scientific research.

Raul da Costa
Graduate Apprentice

I'm currently in my third year studying Cybersecurity at Edinburgh Napier University. During my time at EPCC, I've been learning about managing the large, powerful computer clusters that power the facility. It's been fascinating to visualise how these supercomputers work and the important role they play in scientific research. My day-to-day tasks involve system administration, monitoring, and troubleshooting any issues that come up. I'm also getting the chance to work on some other special projects and infrastructure tasks that help me to understand how everything operates.

Looking ahead, I'm excited about the opportunity to dive even deeper into the world of high performance computing, cluster management and automation. The team here has been so supportive, and I feel lucky to be gaining this hands-on experience as part of my studies.

Connor Finlay
Graduate Apprentice

I'm currently pursuing a degree in Cyber Security. My academic interests span penetration testing, networking, programming, and operating systems as I am deeply passionate about understanding and mitigating vulnerabilities in complex systems. This is an exceptional opportunity to apply the knowledge gained from my course.

At EPCC I work on the infrastructure of the Edinburgh International Data Facility (EIDF), where my responsibilities include supporting the managed edge that controls access to EIDF services, maintaining hardware for OpenStack, and supporting the GPU service. I also use Ansible extensively for automation and configuration management, which has been an excellent opportunity to strengthen my technical abilities in this area. In addition, I contribute to the Safe Haven Services by taking shifts, helping to ensure the infrastructure supporting sensitive data remains secure and operational. These experiences have allowed me to work with advanced technologies and gain invaluable exposure to real-world systems and processes.

I am thoroughly enjoying the opportunity to expand my knowledge of Linux, a critical component in cybersecurity and infrastructure management, while also honing my skills with Ansible. Working with cutting-edge technology and innovative applications has been both challenging and rewarding, providing me with the chance to apply my academic knowledge in practical, impactful ways.

This experience is an incredible step forward in my career ambitions, and I am excited to contribute to meaningful projects and collaborate with experts in the field.

acf

Software Sustainability Institute update

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

The Software Sustainability Institute (SSI) was the first organisation in the world dedicated to improving software in research. It remains committed to its founding 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.

Research Software Maintenance Fund

The SSI has launched the Research Software Maintenance Fund, a new funding initiative offering £4.8 million to support existing research software. The fund aims to improve how research software is maintained and to reduce technical debt, ensuring essential tools remain available to the research community.

Many researchers rely on software but the people and projects developing it may find it hard to secure resources to maintain or update it. This funding is designed to help sustain key research software, ensuring it remains reliable and accessible.

The Research Software Maintenance Fund will hold up to three rounds of funding and conduct an evaluation of the programme to gain a better understanding of how research software maintenance and development can be effectively supported. This initiative directly benefits researchers and organisations relying on research software, the Research Software Engineers (RSEs) responsible for developing critical research tools, collaborative research projects, and open source software research communities.

Funding will be available for individuals and teams who currently maintain existing research software or are responsible for multiple software tools. It is not intended for developing new software or software that does not already have an established community of users.

Find out more about eligibility and how to apply at: https://www.software.ac.uk/research-software-maintenance-fund

Fellowship Programme

The SSI Fellowship programme offers a unique financial support package, networking opportunities, and professional guidance. Each of our 217 Fellows is an important ambassador for good practice in research software. The Fellowship is an opportunity to join an engaging, supportive, and inspiring community. We believe that "Once a Fellow, always a Fellow", and continue to support and engage with our Fellows beyond their inaugural 15-month period.

The SSI's selection process for inducting new Fellows has proven remarkably adept at identifying emerging leaders over the last decade. Following selection, we award each successful candidate a package of flexible financial support worth £4,000 to spend towards sustaining their activities related to research software.

The next round of applications will open in August 2025. All information is available on our website at: https://www.software.ac.uk/programmes/fellowship-programme

Research Software Camp

The SSI runs free online Research Software Camps once a year over the course of two weeks. Each Camp focuses on introducing basic research software skills and good practices, thus starting discussions among various research communities. Researchers from all career stages are welcome to join in the live and offline discussions that will take place throughout the Camp.

The event is filled with diverse online events held at different times each week, accompanied by a carefully curated collection of online resources, from engaging discussion forums to informative blog posts and helpful guides.

The next Research Software Camp will take place in November 2025. Find out more at: https://www.software.ac.uk/training/research-software-camps

Collaborations Workshop 2025

The SSI's flagship event, Collaborations Workshop (CW), took place as a hybrid event at the University of Stirling in May. CW brings together researchers, developers, innovators, managers, funders, publishers, policy makers, leaders and educators to explore best practices and the future of research software.

This year's theme centred around future proofing research software and how we evolve together as a diverse community. Over three days, delegates took part in discussion sessions, collaborative ideas sessions, and hack day sessions, to explore together and create together.

Among the key highlights were the panel discussion on the importance of equity, diversity and inclusion (EDI) for the future of research software, and the keynotes on EDI, the future of research software, and AI.

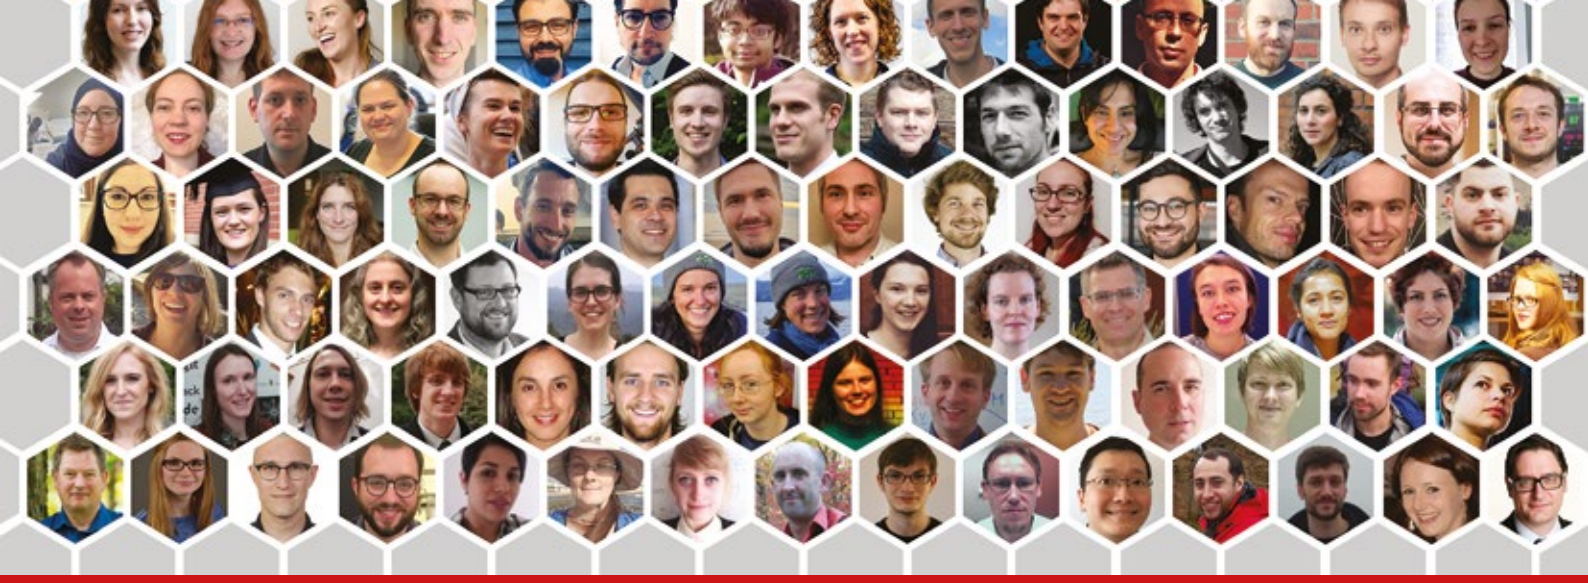

Find out more about CW25 at https://www.software.ac.uk/cw25 and keep an eye on our website to receive all news about the next Collaborations Workshop.

Research Software Practices in the Social Sciences

The recently-completed SSI Research Software Practices in the Social Sciences programme was aimed at social sciences researchers who want to learn new software skills relevant to their field of study. Its purpose was to raise awareness about the importance of publishing software in social sciences and to encourage researchers to follow best practices.

With outreach activities including talks, panel discussions, workshops, recordings, guides, and blog posts, it provided clear guidance on using research software effectively in the social sciences, ensuring that its benefits go beyond technical considerations and support.

Find out more about this initiative: https://www.software.ac.uk/programmes/research-software-practices-social-sciences

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

epcc

Study HPC with us

We offer Masters' degrees in High Performance Computing (HPC); HPC with Data Science; and Imaging, Vision, and HPC.

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 the EIDF GPU Service which supports data processing and AI workloads.

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.

This year we also launched a new on-campus MSc programme in Imaging, Vision, and High Performance Computing in collaboration with Heriot-Watt University, UK.

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 27th by QS World University Rankings 2025.

"Overall, my experience with EPCC has been excellent, especially thanks to my supervisor who has been incredibly dedicated. I am grateful to have had the opportunity to know such a great mentor and friend." **Yongkang Qiu** *On-campus MSc in HPC with Data Science, 2023-24*

"The programme was incredibly rewarding. I now feel well-prepared for a future career in HPC." **Petter Sandås** *On-campus MSc in HPC, 2023-24*

www.epcc.ed.ac.uk/msc