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

epcc news

Issue 92 AUTUMN 2022

From aerospace to biomedicine and Earth observation, we provide computational support for world-class research and innovation

Delivering UK supercomputing and data science excellence to the world

Image: GettyImages/Andrei Berezovskii

From our Director

Welcome to the Autumn 2022 issue of EPCC News. Over the past six months, the new ARCHER2 service has really got into its stride and is very busy supporting a huge variety of scientific and industrial research projects. We're really pleased that the system is being so heavily used – it's regularly almost completely full – but also that the ambition of our users has grown with the size of the system.

EPCC continuously bids for funding for new projects. In this issue we're particularly pleased to report that we've been successful in joining the Met Office Academic Partnership (MOAP). We've worked with the Met Office throughout our history, and MOAP opens a new chapter of joint activities. At the same time, our commercial activities continue to grow through new projects with a really broad range of organisations from some of the UK's largest, to some of Edinburgh's smallest companies.

Along with ARCHER2, the Edinburgh International Data Facility is gradually maturing. The data science cloud is now open to users and very shortly the GPU cluster will open as part of it. You can also read about our upgrade of the Cerebras CS-1 AI supercomputer to two CS-2 supercomputers, giving a total of 1.7 million AI cores. Along with our soon to arrive Graphcore Pod64 system, this gives us a unique AI capability in Edinburgh.

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

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.

Met Office Academic Partnership: advancing weather and climate science

The Met Office Academic Partnership (MOAP), a cluster of research excellence consisting of eight UK universities, is aimed at advancing weather and climate research. The University of Edinburgh, through EPCC, is delighted to have now joined this partnership.

MOAP brings together the Met Office and leading UK institutions in weather and climate science, while reaching researchers in other sectors including health, technology, artificial intelligence, and the social sciences.

The Partnership has invested in jointly funded Chairs at UK universities, with myself appointed to the Joint Chair for the University of Edinburgh. Each Chair will act as the leader for joint research programmes and related activities within their institution and across the Partnership.

EPCC will use its supercomputing and data science expertise to help the Met Office deliver its next-generation numerical weather modelling system, and exploit the data that results from it. Our experience with ARCHER2 and designing the EIDF gives us unique experience that will support our research with the Met Office.

A key ambition of the Partnership is to provide an outstanding environment to develop the science leaders of tomorrow in this very challenging area of research and delivery. The Met Office will sponsor undergraduate and PhD prizes, summer placements and internships and will seek to focus its existing studentship scheme on areas of strategic importance.

EPCC has a long track record of working with the Met Office, starting with the parallelisation of the Unified Model more than 30 years ago. Recent collaborations include:

• EPCC worked with the Met Office to replace its Large-Eddy Model (LEM), which was fundamentally limited in its scalability, with a new code, the Met Office NERC Cloud model (MONC). MONC has been demonstrated to scale to tens of thousands of cores, enabling scientific simulations that previously were not possible. Its modular nature has made it possible to explore GPUs and even FPGAs without the need to fundamentally restructure the code. As part of the MONC development, EPCC also developed a lightweight I/O server for asynchronous writing of diagnostics and prognostics and in-situ data analytics.

• The Numerical Atmospheric Modelling Environment (NAME) code is used to simulate a wide range of atmospheric dispersion phenomena. However its performance was limited because its parallelisation was shared-memory only. EPCC collaborated with the Met Office to introduce distributed memory parallelism using MPI to complement the shared-memory parallelisation, to enable both larger simulations and improved time to solution.

MOAP brings together UCL, University of Bristol, University of Exeter, University of Leeds, University of Oxford, University of Reading, University of Edinburgh and University of Birmingham.

Michèle Weiland, EPCC
m.weiland@epcc.ed.ac.uk

MOAP's aims

Draw together world-class expertise around a focused programme of joint research to tackle key challenges in weather and climate science and prediction.

Maximise return on the UK's investment in research and development in its leading institutions.

Combine our strengths to secure the UK's position in leading the world in weather forecasting and climate prediction.

Build a cluster of research excellence.

Provide an outstanding environment to develop the science leaders of tomorrow.

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

Giving business a competitive edge with HPC

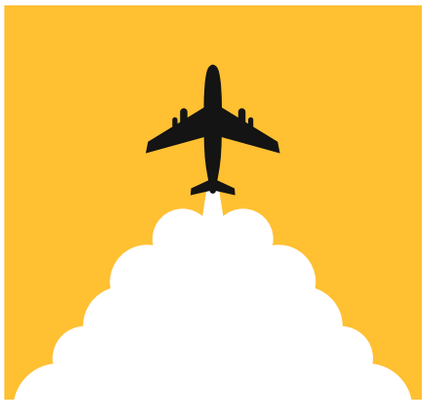

Quantum Base Alpha is working with EPCC to investigate the potential of quantum computing in the aerospace sector.

Image: Getty Images/Ihor Reshetniak

Since our foundation, EPCC has worked with commercial partners to bring them the competitive advantages of high performance computing.

Commercial activities at EPCC can be separated into cycle sales for on-demand access to our Supercomputing infrastructure and project-based income where EPCC staff work on collaborative or contract research projects with industrial customers. These can be short-term or long-term depending on the type of engagement and funding model.

The last quarter has seen a rise in new user accounts for on-demand access to both ARCHER2 and Cirrus from companies in the life sciences, manufacturing, and materials science sectors. These companies are leveraging the ability to parallelise large-scale problems across ARCHER2's 5,860 compute nodes and using the extensive GPU cluster available on Cirrus.

We have also seen an increase in applications for access to EPCC's new AI-specific infrastructure (including systems from Cerebras and Graphcore) and our team is looking forward to supporting new deep-learning proof-of-concept projects in the telecommunications and robotics industries.

Collaborative and contract research projects have continued to be successful at EPCC, with recent success stories including working with Jacobs on code optimisation for HPC environments, providing cloud computing and machine learning support to creative industries start-up Black Audio Goblin. We have also been supporting Quantum Base Alpha, a quantum computing SME, with its recent Innovate UK grant funding success that will research the effectiveness of quantum algorithms to optimise flight paths in the aerospace industry.

One of our HPC specialists recently worked with agri-tech SME and University spin-out Mercury to better optimise its data processing ingest pipeline using Python (see page 6). We have also started working with Orbital Microsystems (OMS) and the University of Edinburgh's School of Geosciences to provide a data ingest pipeline for OMS's microwave telemetry data from its GEMS satellite network. This work was funded by an EPSRC Impact Acceleration Account.

EPCC is providing significant resource and support to various strategic collaborations across the University including projects such as Data Slipstream with Geosciences (see p8), Natwest Group, and Abrdn with Edinburgh Futures Institute and Smart Data Foundry at the Bayes Centre.

These projects will utilise the unique capability of the Edinburgh International Data Facility, which includes the provisioning of data safe haven environments alongside world-class compute facilities to optimise research outcomes and develop innovative products and services.

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

EPCC hosts and manages a remarkable collection of high-performance computing and data services.

These include UK national research systems such as ARCHER2, Cirrus and Tesseract, as well as smaller scale systems designed to explore new technologies.

We work with organisations of all sizes to unlock new frontiers of problem solving, prediction and data analysis.

We can help global enterprises, SMEs and start-ups to develop, manage and scale product development and manufacturing; to evaluate financial risks; and to generate new business insights.

See our website for more examples of how we work with business: www.epcc.ed.ac.uk/industry-solutions

To contact EPCC's business team, email Julien at: j.sindt@epcc.ed.ac.uk

Revolutionising recycling with AI

Image: Getty Images/Anna Kim

We've been providing computing and software support to start-up Danu Robotics, which aims to revolutionise the efficiency of the recycling sector with an AI-powered, vision-guided, robotic sorting system.

Danu Robotics is working with EPCC to develop a system which aims to automate waste sorting at recycling plants – also known as materials recovery facilities (MRFs) – with speed and accuracy, thereby increasing the quality of output and reducing costs.

Recycling waste is one of the easiest ways to reduce the use of limited resources and curb climate change, but often the materials that end up on the conveyors of the MRFs are not what recyclers want and many contaminants such as containers soiled with food waste have to be removed by hand.

The company's founder and CEO Xiaoyan Ma, who completed an MSc in HPC with Data Science at EPCC, wanted to come up with a solution after she found how relatively inefficient the current industrial recycling process is. She explains:

"It's estimated that the world generates three billion tonnes of domestic waste each year, but less than 10 per cent of it is recycled. I wanted to develop a system that would make this process easier for recyclers and more efficient and cost effective, and hopefully make it more attractive to recycle more materials."

The Danu Robotics solution is based on machine learning software that can visually identify recyclable and non-recyclable material and remove any items that should not be there. Before deploying the robot picking hardware, the company had to build up a waste image database to help the system identify contaminants.

Now that the initial system training is complete, Danu Robotics is working on the software which will direct the robotic sorting system to remove contaminants from a moving conveyor belt as efficiently and effectively as possible. For this part of the programme, the company called in EPCC for support.

EPCC initially worked with Xiaoyan to outline the system's architecture and this led to another project to train the AI part of the system to identify recyclable and contaminant items. EPCC's Cirrus supercomputer was employed to help process the data and train the software.

The project is now accelerating with two months of lab tests underway to integrate the software with the robotic hardware, and then a three-month trial of the prototype system at Glasgow City Council's recycling centre. Several large European recycling companies are interested in Xiaoyan's product.

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

"Our system is designed to be sustainable, flexible, affordable, scalable and future proof and we believe that our technology can help recycling companies recoup their investment within two years and double their profit within three or four years' time.

"But most importantly, our technology is a key enabler for efficient and effective recycling.

"The help of EPCC has been tremendous and definitely provided significant professional software engineering support to help drive the project forward."

Xiaoyan Ma,
CEO Danu Robotics

See our website for more examples of how we work with business: www.epcc.ed.ac.uk/industry-solutions

https://danurobotics.com

Satellite data and HPC combine to improve arable farming

A key objective of the EuroCC project (see page 12) has been the transfer of knowledge to industry. As part of this work, EPCC has helped a UK startup, Mercury Environmental Systems Ltd, to develop a service that uses satellite data and high-performance computing to monitor and forecast crop growth. This Mercury data service aims to aid more efficient and sustainable farming as well as broader ecological and environmental decision-making.

Mercury Environmental Systems is preparing to offer a commercial service that combines satellite observations with computational simulations in a novel way to provide farmers with continual crop growth monitoring and yield forecasting for fields under their management. This will enable farmers to plan ahead based on expected crop yields several months in advance of harvest.

Insights from the Mercury growth tracking and forecasting system should, thanks to the inclusion of nitrogen cycle modelling, also help facilitate more efficient use of fertiliser – an increasingly costly global commodity – by providing variable-rate nitrogen application maps to farmers. These maps can support better decisions relating to the optimal applications of nitrogen across a field while simultaneously reducing cost and environmental impacts. Moreover the inclusion of carbon cycle modelling supports environmental monitoring and carbon accounting.

At any stage of the crop growth cycle, satellite observations of crop vegetation are used to calibrate a growth model at the sub-field scale, which is then used to produce a forward forecast through simulation. EPCC has optimised execution of these runs on ARCHER2, the UK national HPC service, reducing their overall computational cost by a factor of 2.5. This is a significant saving in Mercury's projected operational costs.

We also implemented multi-node parallelisation so the time required to generate a forecast remains constant regardless of field size or the spatial resolution of the model. This has the identical financial cost to using only the cores on a single compute node, to which the software was previously restricted.

Finally, we helped Mercury design and implement in software an operational model underpinning its service that will enable it to efficiently automate batch executions of forecast workflows for all its customers. This was done with ARCHER2 as an initial target execution platform, but with sufficient generality and flexibility to allow execution on other HPC and

Arno Proeme, EPCC
a.proeme@epcc.ed.ac.uk
Andrew Revill, Mercury
ar@mercuryes.com

Image: Getty Images/Chris Hepburn

cloud computing resources, thereby enhancing Mercury's operational resilience and safeguarding against compute resource lock-in.

Mercury's forecasting system has the potential to have a significant impact on arable farmers and on broader agriculture, forestry, and emerging environmental markets in Scotland and worldwide. Innovative, data-driven forecasts will help farmers improve productivity and help reduce carbon emissions. Carbon forecasting and monitoring services will be a critical tool in planning and verifying land-based carbon sequestration projects at local to regional scales.

We are therefore pleased that by optimising and parallelising Mercury's software to make use of HPC, providing software architecting consultancy, and skills transfer with regards to parallel computing we have helped put the company in a better position to progress towards an operational launch date and to serve and grow its future customer base.

What our partner says

"We are very grateful for the expertise and guidance provided by EPCC, which was essential for developing a robust modelling framework that has the capacity to generate multiple outputs at the spatial and temporal scales required to deliver the Mercury data products. Comprehensive profiling and subsequent code development was carried out by EPCC to ensure the central model structure ran as efficiently as possible across multiple cores.

"The assistance and advice provided throughout the collaboration was crucial to overcoming issues related to interacting with HPC resources. Further support received from EPCC included safeguarding and future proofing the Mercury workflow when producing timely estimates of crop carbon and nitrogen fluxes along with yield forecasts."

Andrew Revill,
Mercury Environmental Systems

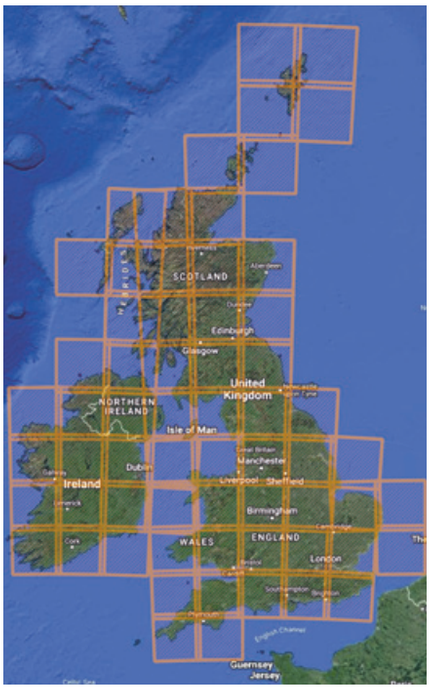

See our website for more examples of how we work with business: www.epcc.ed.ac.uk/industry-solutions

www.mercuryenvironmental.com

Data Slipstream: bringing together Earth-observation data, science, industry, and next-gen compute

The Data Slipstream project is building a system whereby the large, diverse and complex datasets vital to Earth Observation (EO) research at the University of Edinburgh and beyond can be brought together in one place.

Dave Mckay, EPCC
d.mckay@epcc.ed.ac.uk

Such data underpins research into geoscience and engineering challenges including climate change, land use change, agriculture, planning and infrastructure, landslip, forensic archaeology and anthropology.

A major difficulty is caused by the many data sources, APIs and protocols, licences, software, conventions, data formats and, of course, the sheer amount of data involved in EO research. The scientists who rely on this data are often forced to access it in a piecemeal fashion, limiting the scope of their research.

Another challenge is the merging of data. Fused EO data can be even more powerful than the sum of its parts; a flood warning system based on optical images of a flood plane would become vastly more predictive when those images are linked with weather and climate data, soil saturation data, and digital elevation models.

Building data infrastructure

But Data Slipstream is not just a collection of data. The project will design and build the infrastructure that brings the researchers to the data, and provide the tools, environments, and computing power required to develop algorithms and machine learning models, perform analysis, and produce results. And, with the expertise of the Data Slipstream team, develop services and products around that core science.

The project began in 2020 with £215K from the UK Space Agency (UKSA) National Space Innovation Programme to initialise the development of Data Slipstream on the Edinburgh International Data Facility (EIDF), with the overarching goal of facilitating climate change mitigation and adaptation from EO data.

Following on from the UKSA funding, further funds were sought from the UKRI Industrial Strategy Challenge Fund to deliver the PASTORAL (Pasture Optimisation for Resilience and Livelihoods) agri-tech service, which produces timely crop yield forecasts to aid in agricultural decision-making by combining satellite data, a novel modelling framework and weather forecast data, and real-time pasture productivity and carbon cycling information.

The role of Data Slipstream was to host the required data and provide the infrastructure for modelling to be conducted on ARCHER2. A new spin-out company, Mercury Environmental Systems Ltd, was formed and continues to work towards increasing productivity and sustainability whilst delivering net zero (see page 18).

Data Slipstream has been running as a collaboration between EPCC and the School of Geosciences since 2020, initially as a proof-of-concept project funded by the UK Space Agency (UKSA) in 2020.

It was an early adopter of the Edinburgh International Data Facility Data and will provide satellite and Earth Observation data ingest, processing, hosting and dissemination services.

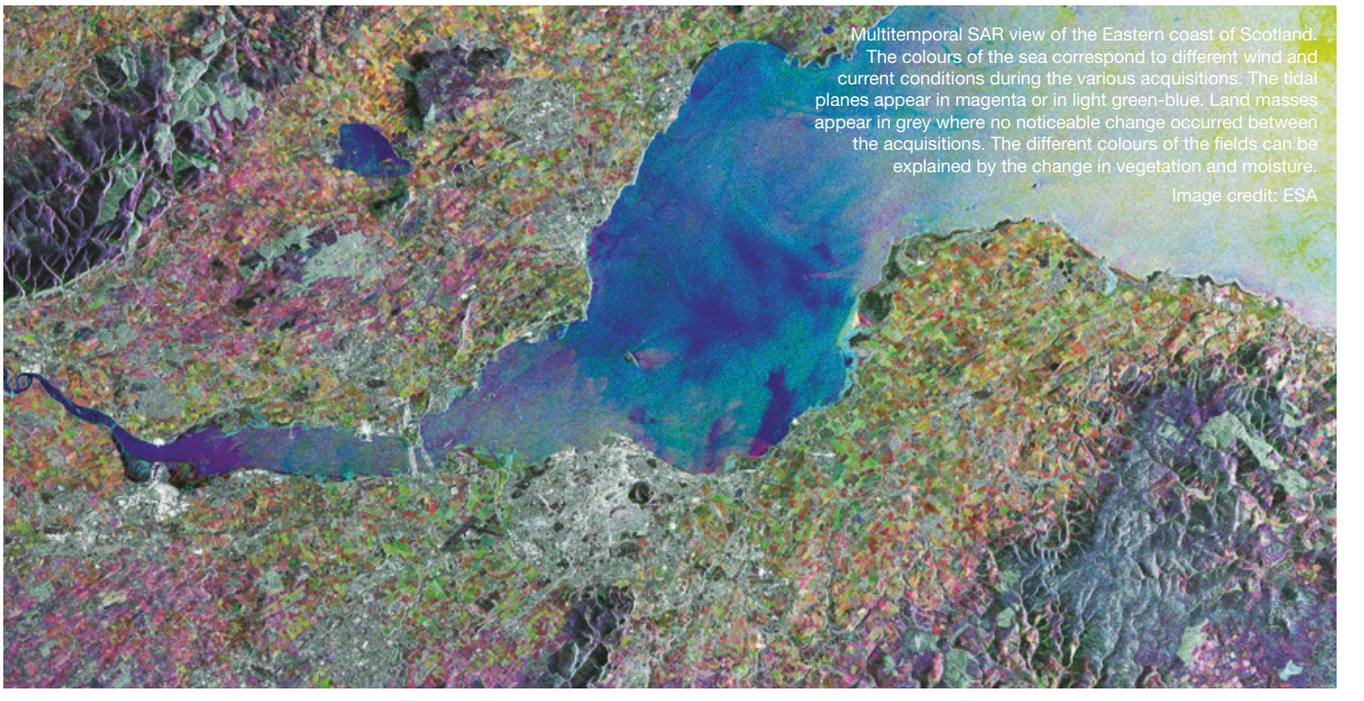

Multitemporal SAR view of the Eastern coast of Scotland. The colours of the sea correspond to different wind and current conditions during the various acquisitions. The tidal planes appear in magenta or in light green-blue. Land masses appear in grey where no noticeable change occurred between the acquisitions. The different colours of the fields can be explained by the change in vegetation and moisture.

Image credit: ESA

Shortly after this, UKSA SPRINT funding was won in partnership between EPCC, the School of Engineering, and Astrosat Ltd to develop a deep-learning platform for the prediction of soil moisture from data provided by the European Space Agency's Sentinel-1 satellite, with the COSMOS-UK network providing ground-truth data. The model, built using Tensorflow on EPCC's Cirrus supercomputer, was trained using data from 51 soil moisture measurement sites and thousands of images taken over the UK between 2014 and 2019.

Evolving with EIDF

Data Slipstream was an early adopter of the EIDF. As the EIDF vision becomes fully realised, Data Slipstream is being improved and redesigned to exploit the infrastructure and systems the EIDF offers. Work is currently underway to migrate Data Slipstream to an OpenStack platform, where virtual machines can be spun up on demand and provide a full EO data analysis environment through Jupyter Lab, and CPU and GPU compute through Kubernetes. Together with increased, faster storage and improved automation and cataloguing, EO data will be placed at researchers' fingertips.

Data Slipstream now hosts UK-wide Sentinel-1 data, Sentinel-2 data covering all of Scotland and parts of Ghana, and many other datasets required in the processing of that data, such as digital elevation models of the Earth's surface. Data sources include the Alaska Satellite Facility, the Centre for Environmental Data Analysis, and the European Space Agency. Data acquisition is project led, meaning any new data acquired is directly linked to a use case.

Future developments

The latest Data Slipstream project, funded by the STFC Impact Accelerator Account, involves an international industrial collaboration with Orbital Micro Systems to receive low-latency passive microwave data recorded by their Global Environmental Monitoring (GEMS) CubeSat constellation.

Data from the demonstration instrument, launched in 2019, will be used to prepare the Data Slipstream infrastructure to receive raw satellite data to produce up-to-the-minute 3D precipitation, temperature, and moisture profiles of the Earth's atmosphere from instruments due for launch beginning in early 2023 from Shetland's SaxaVord UK Spaceport.

EIDF is at the heart of a plan to make Edinburgh the data capital of Europe. Through Data Slipstream, we aim also to become the space data capital of Europe.

Multispectral cloud-free composite image (showing combined RGB channels) produced by combining Sentinel-2 images from Summer 2019 produced using Data Slipstream by the School of Geosciences and EPCC staff on the EIDF.

Data Slipstream resources will be accessible through EIDF:

www.ed.ac.uk/edinburgh-international-data-facility

Building world-class computational services for data-driven innovation

The Edinburgh International Data Facility has entered an exciting phase with the launch of new services and hardware - and there's still more to come.

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

EIDF virtual desktops

In June 2022 we were delighted to launch the first version of our data science virtual desktop service, offering Linux virtual machines (VMs) through a "desktop in a browser" virtual desktop interface (VDI). This service has been designed for users who need more compute and storage than their local resources allow but don't need the full power of a high-performance computing (HPC) or dedicated machine learning system.

EIDF virtual desktops run on a computing cloud hosted entirely within the University of Edinburgh's Advanced Computing Facility. Desktop machines can be ordered with a range of pre-installed data science software including Python, R, and JupyterLab. The service also offers a full visual desktop rather than a bare Linux command line, enabling the use of graphical tools such as RStudio or Jupyter Notebooks.

EIDF computing services

The data science cloud is our latest EIDF computing service, joining the large-memory Ultra2 system and two Cerebras CS-2 AI supercomputers (see opposite).

Our plans for the rest of the year will complete the set, seeing the launch of a lightweight data science notebook service, a new Nvidia GPU cluster, and the acquisition of a Graphcore BowPod64 AI system. The GPU cluster and BowPod64 plus the CS-2 will give us an unrivalled array of leading-edge hardware for the hardest of machine learning challenges.

The combination of data science notebooks and desktops covers the smaller-scale, long tail of data science, while the Ultra2 system, in tandem with our traditional HPC services ARCHER2 and Cirrus, offers high-performance data analytics at scale.

But what about data?

The computational side of things, then, is in good shape, but EIDF will be nothing without data, of course. It currently sports some 50 petabytes of storage and to complete the overall picture we are working on the final stages of a general-purpose data ingest, processing, and archiving pipeline which will supply the rich and varied analytics-ready data lake at the heart of the Facility.

Our eventual aim is to make this data lake visible to any and all of the computational services – including the national HPC services – to enable our users to incorporate data of all types into their research.

It is expected that by the end of 2023 EIDF will provide a comprehensive, world-class set of computational services for data-driven research and innovation.

The Edinburgh International Data Facility brings together regional, national and international datasets to create new products, services, and research. It is funded by the UK and Scottish Governments under the Data Driven Innovation Programme of the Edinburgh and South-East Scotland City Region Deal.

www.ed.ac.uk/edinburgh-international-data-facility

The City Region Deal annual report includes a video showcasing EIDF's role in supporting COVID research: https://adobe.ly/3TvqR6n

The Cerebras system, which is housed at EPCC's Advanced Computing Facility.

Our AI capabilities extended with multi-million pound upgrade

Following the successful deployment of our Cerebras CS-1 system, the first to be launched in Europe, we have now replaced it with a CS-2.

The Edinburgh International Data Facility (EIDF) provides computing services designed for data science. This includes scalable notebooks, data science desktops and specialist compute platforms, and now two Cerebras CS-2 large-scale artificial intelligence (AI) engines.

This multi-million pound upgrade represents a significant expansion of our artificial learning (AI) capabilities as part of the EIDF programme.

With 850,000 AI optimised cores, the Cerebras CS-2 system boasts over twice as many cores and twice as much on-chip memory as the CS-1, and over a hundred times as many cores as the current class-leading GPUs. This allows EPCC to run industry-standard AI frameworks such as PyTorch and TensorFlow with models up to a million parameters.

With the option to cross-connect two CS-2 systems, we have the potential to use even larger models for tackling extremely large problems unthinkable on normal GPU clusters.

This also gives us the ability to include a CS-2 system in the next iteration of the National Safe Haven, giving researchers the ability to use cutting edge systems for Natural Language Processing within the strict confines required for access to controlled health data, for example to analyse, search and process digitised records at a scale previously not possible.

Safe haven services

EIDF provides a range of security environments for data analysis and archive services: standard; protected data access; and safe haven.

The Scottish National Safe Haven provides a governed, trusted research environment for accredited researchers to work on approved projects of public benefit using sensitive data.

The service is operated by EPCC as part of the EIDF infrastructure but is controlled by Public Health Scotland under the oversight of Scotland's Public Benefit and Privacy Panel.

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

Dynamic service

EIDF provides computing, data management, IoT and safe haven services for the Data-Driven Innovation (DDI) programme, and offers bespoke customisation and service co-creation for DDI partners, researchers and innovators.

It is dynamic and challenge-led, with services co-designed with our partners to solve real problems. EIDF's service catalogue will grow and change over the coming years in response to the needs of our stakeholders.

The CS-2 is generally available for DDI programme partners. For information on how to apply for access, email: eidf@epcc.ed.ac.uk

EuroCC@UK: a national centre of HPC competence

Some of the EPCC and Hartree EuroCC teams at ISC2022.

Since 2020 EPCC, together with the Hartree Centre, has been operating the UK National Competence Center of the EuroCC network.

The objectives of the EuroCC network are to increase the accessibility and availability of high-performance computing (HPC), high-performance data analytics (HPDA), and artificial intelligence (AI) in their respective countries. Also to provide a single "front door" into HPC research and development capabilities for both industrial customers and academic researchers.

The adoption of a strategy based on national competence centres (NCC) reflects the differing levels of maturity of these technologies across Europe, and it allows countries to focus on their own priorities. In addition to work done at the national level, the NCCs have served as conduits for exchanging knowledge between each other. This has helped to establish best practice and elevate the quality of research and innovation in HPC, HPDA, and AI to a common high standard across Europe.

HPC outreach

EPCC has been active in investigating state-of-the-art technologies, developing and delivering training, and general outreach. For example the UK NCC was well represented at the international ISC HPC conference earlier this year with EPCC hosting an exhibition stand. Our collaborations within EuroCC have included partnering with the Swedish NCC to investigate the use of CodeRefinery to develop accessible self-service HPC courses.

Knowledge transfer

A key objective of the programme has been the transfer of knowledge to industry, an area in which EPCC has a proven track record.

Through EuroCC we have helped local companies recognise the benefits of new technology, for example with Edinburgh-based Danu Robotics (see p5), which has developed an AI system to improve the efficiency of recycling through automation. EPCC has also given a presentation to other NCCs on the important issue of security requirements for HPC centres who wish to engage with industry.

EuroCC@UK workshop

With EuroCC set to continue into new phase, we will host a workshop in 2023 to showcase examples of how HPC and AI can benefit industry, and how companies can engage with the EuroCC network. The event will take place at the Bayes Centre in Edinburgh where EPCC is based. Further details of how to participate will appear on the EPCC website.

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 951732.

Lorna Smith and Mark Sawyer, EPCC
l.smith@epcc.ed.ac.uk
m.sawyer@epcc.ed.ac.uk

The National Competence Centres (NCCs) are the central points of contact for HPC and related technologies in their country. Their missions are to:

- Map HPC competencies and institutions in their country.
- Act as a gateway for industry and academia to HPC expertise and projects.
- Promote HPC training.

EURO

EuroCC
www.eurocc-access.eu

Hartree Centre
www.hartree.stfc.ac.uk

HPC Scalability Service for biomedical applications

Image courtesy CompBioMed.

Users of computational biomedical applications are increasingly encountering performance issues, perhaps code cannot run in an acceptable time, or large-scale sensitivity analysis is required by a regulatory authority. CompBioMed's free Scalability Service is designed to help.

EPCC is part of CompBioMed, a European Commission Centre of Excellence focused on the use and development of high performance computing (HPC) for biomedical applications. HPC can enhance industries in the healthcare sector including pharmaceuticals and medical device manufacturers.

CompBioMed now offers free support to organisations in their initial steps towards either parallelising their existing computational biomedical applications, or improving the scalability of those applications already parallelised, and thereafter deploy them on HPC platforms.

Its Scalability Service boosts the performance of biomedicine applications via a range of support routes, such as informal discussions about efficient use of parallel platforms in general and code reviews, to porting and profiling applications and suggesting improvements, or working closely with the client and adapting the source code on their behalf.

Please visit the website [1] for more information, where you'll find links to a group email of experts in both HPC and biomedical applications, and access to the public Slack channel #scalability, hosted by "In Silico World" Community of Practice [2], which provides a safe space to share scaling questions. The website also proffers both a detailed application form and a more informal webform version.

The Scalability Service includes a live, useful overview for programmers with ideas on improving the performance of parallel applications for high-end supercomputers [3].

In the field of biomedicine, many applications deal with sensitive data, and clients can be assured that great care is taken when adapting their applications and managing the associated data. The data policies cover data privacy, data security, and research data management. See link [4].

[1] compbiomed.eu/compbiomed-scalability-service

[2] insilico.world/community

[3] compbiomed.eu/rough-guide-to-preparing-software-for-exascale

[4] compbiomed.eu/compbiomed-data-policies

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

Personalised medicine is on the horizon now that computational simulations of human biology are reaching maturity in the biomedical domain, rendering predictive models of health and disease increasingly relevant to clinical practice.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 675451 (phase 1) and grant agreement No. 823712 (phase 2).

www.compbiomed.eu/compbiomed-scalability-service/

Building a common ecosystem for DSLs

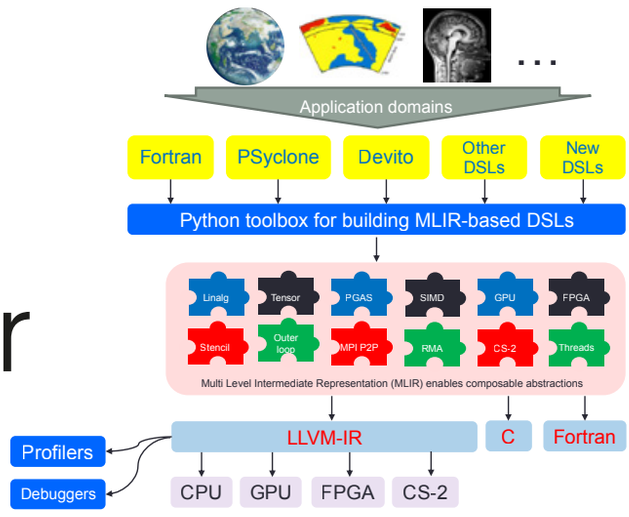

As we enter the era of exascale supercomputing, a major challenge is how to program these next-generation machines. We are often limited to languages designed for sequential work, which are time-consuming to employ.

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

Domain Specific Languages (DSLs) potentially offer a way to address this issue. Programmers can work at much higher levels of abstraction and, as a rich source of information is available to the compiler, it can effectively make the tricky, low-level decisions that commonly take time. Also, moving from one architecture to another is much easier because the compiler has more information with which to adapt the code for different hardware.

However a major drawback with DSLs is that they tend to be heavily siloed, where one DSL shares very little or no infrastructure with any other. This is disadvantageous for numerous reasons, such as the high overhead required to actually write a DSL in the first place, and a significant risk for potential users as there are then uncertainties around long-term maintenance of the stack and whether the DSL will still be around in five years time!

In the xDSL project we are looking to address this by leveraging a common ecosystem for DSLs which builds upon the ubiquitous MLIR and LLVM. MLIR, first developed by Google, is a key enabling technology as it provides composable dialects and transformations which can be shared between different compilers. Leveraging these existing building blocks to develop a compiler significantly lowers the barrier to entry, and commonly they generate LLVM-IR which feeds into LLVM. LLVM provides a rich set of backends which target many different hardware architectures and consequently enables code to run on these with minimal changes.

In xDSL we are developing a Python toolbox to lower the barrier to entry to MLIR and LLVM which, until this point, have required a fairly significant learning curve. We are also providing a variety of high performance computing-focused dialects and transformations that DSL developers can use as the foundation for their own abstractions. We are using the popular Devito and PSyclone DSLs as demonstrators as well as working with other groups to help them integrate their technology with xDSL.

xDSL is a collaboration between The School of Informatics at the University of Edinburgh, Imperial College, and EPCC. It is funded by ExCALIBUR. ExCALIBUR is a £45.7m UK research programme that aims to deliver the next generation of high-performance simulation software for the highest priority fields in UK research. EPCC has funding from this programme to explore a range of topics.

We have a growing and thriving open source community around xDSL on GitHub. See our website to find out more: www.xdsl.dev

We will also have a research poster at SC22. Please drop by booth 2022 if you will be attending the conference.

EPCC joins RISC-V International

CPU Instruction Set Architectures (ISAs) are typically proprietary, which limits the number of implementations and how they can be adapted. By contrast, RISC-V is an open ISA: developed by the community, anybody can use its specification to freely develop a CPU implementation.

Not only does this encourage a collaborative effort to develop a solid and mature ISA, but it also results in many RISC-V CPU implementations and a rich software eco-system. The world's largest companies support RISC-V and it is already common in embedded systems. It has been forecast that there will be over 60 billion RISC-V-based devices in existence by 2025.

EPCC hosts the ExCALIBUR H&ES RISC-V testbed which aims to provide scientific computing users with access to this technology and explore its role in HPC. While it may seem far fetched that RISC-V will displace x86, such as the AMD Rome we have in ARCHER2, it could certainly complement such technologies in the future, for example numerous accelerators have been developed which are driven by RISC-V CPUs.

One of the most interesting aspects of working with this technology is that, while several physical RISC-V CPUs are publicly available, the most cutting-edge designs are typically provided as soft-cores. These are software descriptions of the CPU which can be compiled and used to configure reconfigurable hardware, making this appear and function like the CPU. They can be very advanced, for instance including the ability to run an OS such as Linux. Using Xilinx's latest-generation Versal FPGA to host our soft-cores, we are especially interested in this part of the project as it means we can experiment with a diverse set of latest-generation RISC-V CPUs and provide them to users too.

RISC-V International

As a part of the testbed we have joined RISC-V International as a community member. RISC-V International is the body that oversees RISC-V and has over 200 members ranging from the world's largest companies to SMEs and academic institutions.

There are many activities on-going within RISC-V International to further develop and promote the technology, and a major benefit for us in joining is that we can now help influence such activities and the community more widely as we undertake the RISC-V testbed project.

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

RISC-V International comprises a large member organisation building the first open, collaborative community of software and hardware innovators. RISC-V combines a modular technical approach with an open, royalty-free ISA. As a non-profit, RISC-V does not maintain any commercial interest in products or services. As an open standard, anyone may leverage RISC-V as a building block in their open or proprietary solutions and services.

For more details on the testbed see riscv.epcc.ed.ac.uk

RISC-V website https://riscv.org

https://excalibur.ac.uk/themes/hardware-and-enabling-software/

Ensuring continuity of service at the ACF

EPCC's Advanced Computing Facility (ACF) delivers a world-class environment to support the many computing and data services which we provide. This article takes a behind the scenes look at some of the activities the ACF team undertakes to provide the stable services our users expect.

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

Computer Room 3 (CR3) of the ACF has recently been the focus of a great deal of activity, with all of the main power distribution units' (PDUs) supply cabling and sub-floor power supply cables being replaced. The ARCHER2 cooling distribution units (CDUs) and pipework have also been being cleaned internally by HPE.

The need for the PDU cabling upgrade arose from regulatory 5-year inspection and testing which showed that, although adequate for the original ARCHER machine, the electrical infrastructure was being pushed to its operating limit when ARCHER2 was at maximum capability.

The simplest way to accomplish this work would have been to switch off ARCHER2 for two months while the cables were replaced. As this was obviously not possible, the system power supplies had to be maintained throughout the works to ensure ARCHER2 remained operational. In total we had to install around 3km of large diameter power cables without disrupting any services.

Working with our electrical contractor and HPE we devised a way to connect enough 125 amp 3-Phase temporary cables to power ARCHER2 compute cabinets while we isolated and worked on each PDU in turn to install and connect the new cabling. Although difficult to programme and implement, the supply cables were all installed and connected with zero disruption to users of ARCHER2.

ARCHER2's cooling and management cabinet supplies, which are connected to UPS power, were an additional consideration. To keep these systems operational, each of the UPS-fed PDUs were back-fed from one of the other UPS PDUs to ensure these more critical circuits were not interrupted during the works. This was not a simple task, but they have now been fully rewired to meet present and future needs.

All these works were completed with no disruption to ARCHER2 service provision thanks to great coordination and liaison between our electrical contractor, HPE, and the ACF site team.

Another electrical issue which came to light at this time was that the 125 Amp 3-Phase plug/socket connections to the ARCHER2 compute cabinets were susceptible to overheating under heavy load, a situation which has also been seen in other locations worldwide.

To address this we decided to not only replace the plugs with hard, bolted connections (all 69 of them), but to fully rewire all the sub-floor cabling for ultimate safety of supply. This again presented the same

Above: Some of the removed 125 amp 3-Phase plugs.

Below: 800 amp 3-Phase PDU with new supply cables at top left.

Hot and cold cooling connections to the rear of ARCHER2 compute cabinets.

challenge: to complete the works without disrupting the service.

Working with the contractor, HPE, and ACF site team we devised a system to provide and connect temporary supplies to the compute cabinets with “hot swaps” of these critical connections, which allowed the existing cables to be removed, replaced, and the cabinets connected to the new cabling without any down-time. This approach proved to be acceptable to all stakeholders and was completed again with no loss of ARCHER2 service, which was a great achievement for all concerned. The electrical contractor’s staff deserve everyone’s thanks for working in very cramped underfloor conditions for several months.

While these electrical works were progressing in CR3, HPE identified an issue with ARCHER2’s cooling distribution units (CDUs) and their connected cooling pipework to the compute cabinets. This issue has also been seen in similar systems around the world. It was therefore decided to clean the insides of all this pipework.

To minimise disruption to the ARCHER2 service, it was agreed that four compute cabinets (a sixth of the total system) would be taken out of service at a time for remedial work, thus allowing work to proceed while ensuring the system remains operational, albeit at a reduced capability. Users may have noticed parts of the system being taken our of service and returned a few days later.

This work has been carried out by HPE specialists flown in from USA and Europe, supported by the site-based HPE team of Martin, Felipe, and Greg.

The works are now generally finished with only some final cleansing ongoing, which is expected to be completed in the coming weeks. The ARCHER2 service will then be back to full capability including the ability to consider new High Performance Linpack benchmark runs to verify its maximum performance now we have more experience with the system.

While all of this has been ongoing, we’ve also swapped out our Cerebras CS-1 for a new CS-2 (see p11), racked an additional CS-2 for future use, ordered several more HPE ARCS cabinets for near future installation, started planning for a potential expansion of the DiRAC Tursa system, upgraded our site management network, continued planning and design for Exascale, and all of the other daily tasks to ensure the ACF functions smoothly, as well as looking forward to whatever the future throws at us.

The Advanced Computing Facility is EPCC’s high performance computing data centre. It is home to the ARCHER2 national supercomputing service, the Edinburgh International Data Facility, and other systems of national importance. Each of the ACF’s four computer rooms hosts specific HPC and storage equipment and is supported by associated plant rooms which provide dedicated power and cooling infrastructure for each room.

Read more about the ACF on our website: www.epcc.ed.ac.uk/hpc-services/advanced-computing-facility

Net Zero: investigating the energy use of ARCHER2

UKRI is committed to net zero emissions by 2040. As part of this, it has recently funded the UKRI Net Zero Digital Research Infrastructure Scoping Project to investigate how UKRI can achieve net zero emissions in the digital research infrastructure (DRI) that it funds.

Digital Research Infrastructure (DRI) funded by UKRI covers a huge variety of things, from researchers' laptops and the networking that connects research institutions across the UK together, up to large scale digital facilities. It also includes ARCHER2, the UK's National Supercomputing Service based at, and operated by, EPCC.

As part of EPCC's contribution to the Net Zero Scoping Project we are preparing a case study on ARCHER2, investigating the current greenhouse gas emissions arising from the energy required to operate the service (often referred to as Scope 2 emissions). The study will look at everything from the type of energy used by the service (electricity is procured through renewable sources which are REGO certified) through to the energy used by different software and research communities and hence their impact on Scope 2 emissions.

Looking at Advanced Computing Facility (ACF), EPCC's data centre where ARCHER2 is housed, we utilise free-cooling as much as possible; this is the cooling of systems solely using outside air via water pumped to fans on our roof. Our systems will automatically use this whenever outside air temperatures are sufficiently low. We would typically expect to make use of free-cooling where the outside air temperature is 12 degrees centigrade or less, which is quite often in Scotland!

This summer we also experienced extreme heat events, and we have been preparing robust plans that will allow the temperature of the cooling water to be controlled while maximising workload and minimising the impact on our users.

A key area of research is how the energy use of individual software and research areas impacts on total energy consumption (and emissions) and whether strategies to influence service operation and user behaviour to reduce emissions are possible. Our energy usage and emissions investigations around software and research areas is closely linked with the work of one of the Net Zero sandpit projects, HPC-JEEP.

The full report is due by the end of 2022.

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

Net Zero sandpit projects

EPCC is involved in four Net Zero sandpit projects:

- HPC-JEEP (High Performance Computing: Job Efficiency and Energy Usage, Monitoring & Reporting)
- ENERGETIC (ENergy-aware hEteRoGenEous compuTIng at sCale)
- ARINZRIT (Applying Responsible Innovation to the Net Zero Research Infrastructure transformation)
- IRISCAST (IRIS Carbon Audit SnapshoT).

UKRI Net Zero Digital Research Infrastructure Scoping Project

https://net-zero-dri.ceda.ac.uk

The EPCC Outreach team gets ready for the 2022 Big Bang Fair.

EPCC outreach is back!

During COVID, our ability to deliver outreach events to the community was limited, although we did undertake a number of virtual events and enjoyed seeing and engaging in the novel formats used to deliver activities online. The relaxation in travel restrictions has allowed us to attend in-person activities again, with events across the country.

We returned with a bang, to the Big Bang Fair (BBF). This is a favourite event for us, bringing tens of thousands of children to the NEC in Birmingham to enjoy interactive activities and gain career inspiration. Our outreach objectives match well to BBF, as we want to highlight the importance of supercomputing to society and to showcase the benefits of a career in computational science.

Our first large-scale event in two years, it was a huge success for us, with children attending from all over the country. We did face a few challenges, not least last-minute staff drop-outs due to COVID infections and a rail strike. I am grateful to everyone for keeping their positivity despite the many curve balls thrown our way.

An interesting change from pre-COVID events has been our focus on sustainability and reducing the carbon footprint of our outreach activities. From travel to reusing materials and minimising waste, this is now an integral part of our planning process.

Our most recent big event was New Scientist Live in London, a large three-day event for the general public. Our next major activity will be hosting a drop-in session at the 2023 Edinburgh International Science Festival.

We are in the process of developing new and enhanced activities to demonstrate the principles of supercomputing. These include a revitalised Wee Archie (our suitcase-sized supercomputer), a Minecraft Advanced Computing Facility (inspired by the facility where EPCC hosts its computing systems), duck scratch games, and ways to learn about binary numbers.

Finally, we have introduced our new mascot for ARCHER2. Our puffin (see right) now sits alongside ARCHER's red squirrel.

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

EPCC's outreach activities are part of the ARCHER2 service. See the website for more details: www.archer2.ac.uk/community/outreach/

Make your Python code 10,000 times faster with parallel numpy!

Photo by @TyneSightPhoto

That was the title of the half-day workshop I ran at RSECon2022, The Sixth Annual Conference for Research Software Engineering.

The workshop came about because I realised that there was something of a gap in EPCC's introductory training material, between our 'Supercomputing' Futurelearn MOOC and our parallel programming courses in MPI such as those run for ARCHER2.

The MOOC is purely conceptual with no programming, whereas the programming courses require knowledge of C, C++, or Fortran. What was missing was a course where people could actually do some parallel programming using the language that today most people learn as their first: Python (despite knowing almost nothing about Python myself!).

We therefore devised a workshop to address this. The main goal was not to promote Python but rather to give people practical experience of parallel programming so they could understand its full potential, which might then motivate them to learn a compiled language. However, it was also an opportunity to demonstrate how you can achieve reasonable performance in Python using relatively straightforward programming techniques. To make the training as accessible as possible I also decided to use Jupyter notebooks.

When my workshop was accepted for RSECon22 I realised I had set myself something of a challenge. Of the three areas I was going to cover – parallelism, Python, and Jupyter notebooks – I only really knew about one of them!

By coincidence, one of the founders of EPCC contacted us out of the blue asking for advice: he was running a workshop in parallel Python at PyData in London and, although very experienced in Python, was worried that his parallel programming might be a little rusty. After meeting Nick Radcliffe for a very informative chat I felt a lot more comfortable on the Python front (and I hope Nick gained confidence in his already excellent parallel material).

Around 60 people attended the workshop with almost 50 registrations for login accounts on ARCHER2, which we used for the final practical session. Certainly people seemed to enjoy the session.

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

My test example was a simple cellular automaton model of traffic flow. When I chose the name of the workshop, the 10,000-fold improvement was just an aspiration but we did achieve it on the day. About a factor of 40 came from using numpy arrays (bringing the performance to about a third of that of C or Fortran), with the remaining 250-fold increase from running on many hundreds of CPU-cores of ARCHER2 using the mpi4py interface to MPI.

Workshop material:
https://github.com/EPCCed/PythonHPC#readme

Nick Radcliffe's workshop:
https://london2022.pydata.org/cfp/talk/HH7SUG

Nick's slides:
https://stochasticsolutions.com/pdf/mpi-pydata-london-2022.pdf

Futurelearn MOOC on Supercomputing:
www.futurelearn.com/courses/supercomputing

Understanding RSE training pathways in HPC

Photo by @TyneSightPhoto

Research Software Engineers (RSEs) are essential to research in high performance computing and computational science. While significant efforts have been made to raise awareness of the RSE role and improve career development, training pathways still require better understanding.

The UNIVERSE-HPC project is working to enable people from a wide diversity of disciplines and backgrounds to have well-defined paths for obtaining the skills and experience required for a successful RSE career.

A Research Software Engineer combines professional software engineering expertise with an intimate understanding of research. The bridging of these two practices, combined with the importance of software in modern research, places RSEs at a critical point in the research pipeline. By supporting and ensuring the development of sustainable research outputs that generate reproducible results they have a key role in a research team but, despite their importance, most RSEs are self-taught.

This gap in professional education and the lack of cohort networks for RSEs specialising in HPC is a major challenge. There is specialist training available, including MSc programmes in HPC and scientific computing, but no established education route for those wishing to pursue a career as an RSE.

To address this, UNIVERSE-HPC will develop a clear curriculum framework that will describe the skills required by RSEs. It will also develop and pilot Masters-level courses for those seeking to join the profession.

In addition, the project will establish the key professional support activities that will ensure a vibrant community, enhance knowledge exchange, and improve engagement with industry. To support these efforts, UNIVERSE-HPC includes a programme of community-building and diversity, equality, and inclusion efforts.

To ensure the sustainable and healthy growth of the RSE community, it is important to increase the transfer of skills and talent between academia and industry, improve the facilitation of professional networks and peer-support, and encourage diversity and inclusion to attract RSEs from outside the established RSE community.

The project will engage with the wider HPC community to trace current learning pathways and understand what the community thinks are the most appropriate routes for gaining specialist RSE skills. The input and information captured through these and follow-up conversations will inform and guide the work being undertaken in UNIVERSE-HPC as the project progresses.

The UNIVERSE-HPC project is led by EPCC, and includes Imperial College London, the University of Oxford, and the University of Southampton.

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

What training should RSEs undertake, at what career stage? What materials exist and which are missing? How do we navigate this largely undefined landscape?

In short: how does one train to become an RSE?

We invite everyone in the RSE community to share their perspectives with us by emailing: info@software.ac.uk

Co-founded by EPCC, The Software Sustainability Institute cultivates better, more sustainable, research software to enable world-class research: www.software.ac.uk

Some staff and students of the 2021-2022 MSc programmes at EPCC.

MSc 2021–22: finding the new normal

Being ‘close to normality’ for this year’s MSc programmes in High Performance Computing (HPC) and HPC with Data Science may, in some ways, have been a stranger experience than the COVID-hit year of 2020-21.

In 2020-21 we knew from the start that it would not be a normal year: Semester 1 teaching had been planned for online-first with a possibility to add more hybrid activity as the year progressed if the COVID situation improved. However with the spread of the first variants and the ‘cancelled Christmas’ of that winter, such hopes did not come to pass. The small cohort adapted well to that year’s challenges and many of them performed extremely well, especially considering the circumstances.

Academic Year 2021-22’s challenges were far more nuanced as we tried to find a ‘new normal’. Staff and students were back on campus, with over 90% of teaching in Semester 1 being scheduled on-campus and Semester 2 returning to an entirely pre-COVID schedule, with the exception of exams remaining online. The differences were far more subtle with attendances impacted by isolation and similar requirements, and students finding discussion-based activity less familiar than pre-COVID. Things which in 2019/20 seemed second nature needed time to bed in again and this was a learning experience for both staff and students.

Despite the difficulties this presented, many of our students achieved excellent outcomes and are leaving the MSc to further their careers in a variety of areas, including working at EPCC, PhD study within the UK and abroad, positions at ECMWF, financial institutions, and software companies. We look forward to seeing their careers flourish and possible future collaborations within the vibrant wider HPC and associated communities.

This year also saw the second intake of our online part-time MSc programmes and the first group of students from cohort 1 complete their taught components. We look forward to seeing this group embark on their dissertations and our first graduating class of online students in November 2023.

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

EPCC runs two established Masters programmes, which are available both on-campus and online:

- High Performance Computing
- High Performance Computing with Data Science.

We also offer the opportunity to study for a PhD in areas related to HPC, computational data science, and software sustainability.

Some of our 2022/23 intake.

Postgraduate programmes update

The start of a new academic year each September, especially so soon after end of the previous one, is an incredibly busy, but exciting, time for staff involved with our postgraduate programmes at EPCC.

The Postgraduate Programmes team was delighted to welcome James 00Richards as our new Administrator, to support Jemma Auns in her new role as Coordinator and, in the University's new Student Support Model, Student Adviser.

Alongside James, we have also welcomed our new MSc cohorts for both on-campus and online programmes. Due to a mixture of the cost-of-living crisis, accommodation issues in Edinburgh, and a lack of available English Language Testing internationally, the intake is smaller than 2021/22, but still a good size and staff are already remarking on the enthusiasm of the new cohorts.

Curriculum-wise, the major change to the on-campus MSc is the development of a new 20-credit course, Practical Software Development, for EPCC students. This course combines the core competencies of the Programming Skills and Software Development courses, while also providing the programmes with a bespoke course for our students only. The original courses remain, but now have greatly increased capacity to take students from external programmes.

Online, 2022-23 has seen the launch of the final set of optional courses to bring our full curriculum to the online programmes, while also seeing increased cooperation with the Edinburgh Futures Institute and Data Upskilling for some of our core and introductory courses.

2022 will also see the first ever graduation from EPCC's specific PhD programme. We have had multiple doctoral graduations in the past through the School of Physics & Astronomy, or Pervasive Parallelism Centre for Doctoral Training, but our EPCC colleague Dr Maurice Jamieson will be the first on our specific programme.

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

EPCC is a leading provider of HPC and data science education and training in Europe.

Find out more: www.epcc.ed.ac.uk/education-and-training

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 praSctical 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 the Times Higher Education World University Rankings 2022, and 16th by QS World University Rankings 2022.

"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
MSc in HPC with Data Science graduate, 2021
Now CEO, Danu Robotics

"The exposure to world-class computing facilities and the National HPC Service coupled with lectures taught by highly experienced tutors has opened a world of new possibilities for me, including skills that I am now further honing at EPCC."

Daniyal Arshad
MSc in HPC with Data Science graduate, 2022
Now Applications Developer, EPCC

www.epcc.ed.ac.uk/msc