



Science and  
Technology  
Facilities Council

Scientific Computing

# COMPUTING INSIGHT UK 2024

Catalysing Research


5 - 6 DECEMBER 2024

Manchester Central, UK

[www.ukri.org/CIUK](http://www.ukri.org/CIUK)

## CONTENTS

INTRODUCTION.....	3
CIUK 2024 PROGRAMME .....	5
CIUK 2024 KEYNOTE PRESENTATION .....	7
CIUK 2024 JACKY PALLAS MEMORIAL AWARD PRESENTATION .....	8
ABSTRACTS AND SPEAKER BIOS.....	10
Don't forget our networking events... ..	23
CIUK 2024 EXHIBITION .....	25
CIUK 2024 DAY ZERO.....	33
CIUK 2024 PARALLEL BREAKOUT SESSIONS.....	36
CIUK 2024 RESEARCH ZONE .....	41
CIUK 2024 STUDENT CLUSTER CHALLENGE .....	42
CIUK 2024 POSTER COMPETITION .....	47
We want your feedback... ..	51
See you in 2025... ..	51



## INTRODUCTION

We are delighted to welcome you all to [Computing Insight UK 2024](#).

The theme for this year's conference is "**Catalysing Research**", with sub-themes including "Sustainability", "AI Focus", "UKRI DRI and DSIT Update", "Technology Trends" and "Collaboration".


CIUK 2024 will include [an exhibition of the latest hardware and software releases](#) plus a full, [two day programme of presentations](#) and a series of [parallel breakout sessions](#), including the annual CoSeC Conference. There will be a [poster competition](#) plus the fourth instalment of the [CIUK Student Cluster Challenge](#). We will also present our annual [Jacky Pallas Memorial Award](#). Please take time to visit the [CIUK exhibition](#).

We would also like to invite you to join us on the evenings of Wednesday 4 and Thursday 5 December for our pre-CIUK networking event at the Gas Works Brew Bar, 5 Jack Rosenthal St, Manchester M15 4RA and the official CIUK networking event at Peaky Blinders Bar, 23 Peter St, Manchester M2 5QR (CIUK lanyard and badge required for entry).

We hope that you enjoy the conference.

All information about the conference can be found on the event website... [www.ukri.org/CIUK](http://www.ukri.org/CIUK).

You can also follow us on X [@CompInsightUK](#) (#CIUK2024) and [LinkedIn](#) for the latest live updates.



**Need help during the conference?  
Have a question about the event?**

The CIUK team will be happy to assist.

You can find them at the main reception desk in the foyer or on the exhibition floor... look for the yellow shirts!

You will also find an information screen at the entrance to the exhibition and next to the reception desk.

**SCIENTIFIC  
COMPUTING  
WORLD**

**Official CIUK 2024 Media Partner**

Scientific Computing World is the only global publication dedicated to the computing and information technology needs of scientists and engineers.

# BE A DATA DYNAMO

## The Hammerspace Global Data Platform

- Accelerates HPC/AI workloads on any storage
- Linearly Scalable Performance
- Standards-based Parallel File System
- Data Orchestration for Silos, Sites and Clouds



To learn more how Hammerspace eliminates data silos and makes data an instantly accessible resource for users, applications, and compute clusters, no matter where they are located visit:  
[www.hammerspace.com](http://www.hammerspace.com)

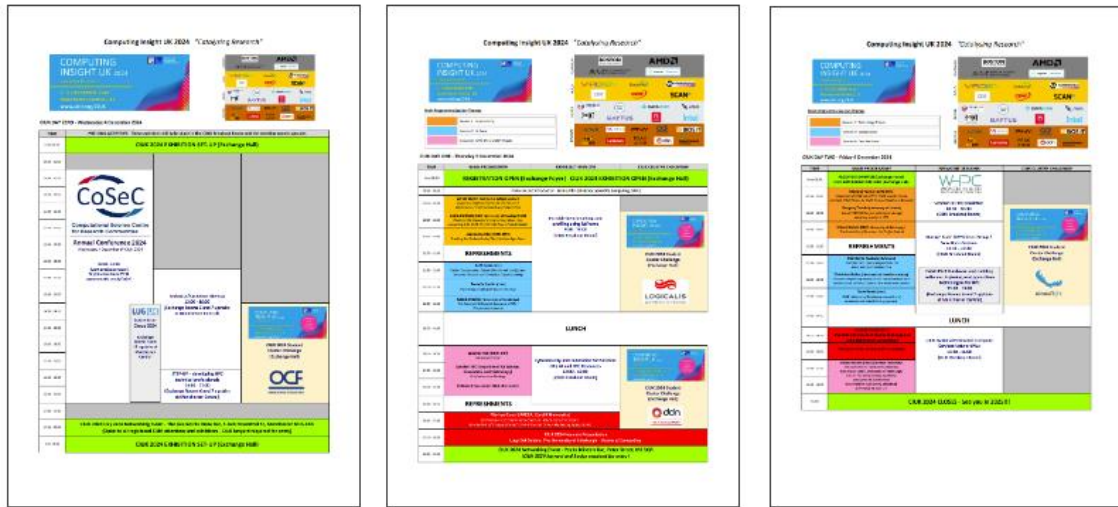
Look for your very own copy of Data Dynamo's latest adventure in your conference bag!



**HAMMERSPACE**

# CIUK 2024 PROGRAMME

The main [CIUK 2024 programme](#) will take place in the Auditorium. Look for the **BLUE** pull-up banners.



**Please click the image above to open the CIUK 2024 Programme**



**Accelerate Discovery:**

**Transform Unstructured Data Management for Faster, More Efficient Research.**

**Visit us at Booth 39/40** to discover how Quantum and its award-winning partner Vespertec can empower your research and IT infrastructure. From optimising unstructured data management to accelerating data center performance, our solutions are built to support faster, more efficient discoveries. Come chat with our team and learn how we can help overcome your IT challenges and reduce total cost of ownership with next-generation technology.



**vespertec**

**Quantum**

For information following CIUK,  
email [info@vespertec.com](mailto:info@vespertec.com) or  
[info.uk@quantum.com](mailto:info.uk@quantum.com)



Luigi Del Debbio  
The University of Edinburgh

### “Visions of Computing”

**ABSTRACT:** Computing has witnessed the most exciting growth over the last few decades. As we enter the era of Exascale computers, Computing is having a major impact on our lives across multiple disciplines. 'Visions of Computing' tries to capture some of the ideas, challenges and projects that are currently associated with Computing. Based on my personal research experience, but also on my work in defining the current European landscape in High-Performance Computing, I will present an overview of the opportunities, the complexity and the challenges that we are facing.

**BIO:** Luigi Del Debbio is a physicist based in Edinburgh, known for their work in theoretical physics, particularly in the areas of quantum field theory and lattice gauge theories. Their research often explores the fundamental aspects of particle physics and the underlying structures of the universe. Del Debbio is actively involved in teaching and mentoring early-career scientists in the field. Their work emphasizes both theoretical insights and computational methods, bridging the gap between abstract theory and practical applications.

Del Debbio has been associated with various European initiatives aimed at enhancing high-performance computing (HPC) capabilities. Their efforts have facilitated international collaborations, enabling researchers across Europe to tackle challenging scientific problems more efficiently. Through their engagement in HPC, Del Debbio not only promotes scientific innovation but also advocates for the importance of computational techniques in modern research, emphasizing their role in addressing fundamental questions in physics and other disciplines.



**Lisa Lampunio**  
Nuclear Engineering Group, Department of  
Mechanical Engineering, Imperial

### **“Advanced Modelling and Simulation for the Analysis of Novel Radiation Detection Technology and Thermal Fatigue Phenomena within the Nuclear Energy Sector”**

**ABSTRACT:** A key aspect of this scientific computing research is to address challenging problems within the nuclear energy sector, such as the development of multiphysics modelling and simulations (M&S) of thermal fatigue phenomena within nuclear power plants (NPPs) and the design of novel radiation detection technologies. Thermal fatigue phenomena are widespread within many areas of engineering and can limit the operational life of components, cause power outages, and require costly periodic inspections. Therefore, high-fidelity predictive models enable engineers to improve the reliability of engineering components and optimise maintenance scheduling. The development of novel radiation detection technologies is critical not only in the nuclear power sector but also for nuclear security and safeguard. Both these fields require the application of advanced high-performance computing (HPC) M&S along with uncertainty quantifications (UQ), surrogate modelling (SM), and machine learning (ML) algorithms. This research work aims to develop computationally efficient and improved M&S methods as well as advanced, cost-effective, radiation detector designs with benefits to both the engineering community and the nuclear energy sector.

**BIO:** I hold a Bachelor and Master’s degree in Nuclear Engineering from the University of Pisa in Italy and a PhD from the Nuclear Engineering Group of the Department of Mechanical Engineering at Imperial. My research work has been focused on developing high-fidelity multi-physics modelling and simulations and efficient data-driven machine learning surrogate models for critically important problems for nuclear power plant safety and control, such as thermal fatigue phenomena in T-junction pipes and design optimisation of neutron detectors within the reactor core.



KAYTUS



# Experience KAYTUS HPC & AI Solution

at #37




## Cornelis Networks CN5000 Omni-Path<sup>®</sup> Interconnect Solutions

Maximising performance and  
scalability for HPC & AI

Data Center Defining  
Next-Generation Interconnects

[www.cornelisnetworks.com](http://www.cornelisnetworks.com)



## ABSTRACTS AND SPEAKER BIOS

### Meet the speakers...

Carter Quinn  
University College London

**Creating a Staffing Pipeline for the Research Infrastructure Profession with Apprenticeships**

Thursday 5 December  
09:30 – 10:00

The CIUK Auditorium



**Abstract:** Within the research computing infrastructure community, staffing shortages are becoming more pronounced, producing top-heavy structures, where insufficient people are filling the entry level roles. This imbalance places increasing pressure in the short-term on teams and a future staffing crisis in the long-term. Partially responsible is the lack of a direct pipeline to create junior staff. In this presentation, we explore UCL's experience piloting an apprenticeship scheme in this field and how it may be applied

community wide.

**Bio:** Carter initially began a career in IT through self-employment during sixth form and moved into salaried positions for various employers. In 2021, he was exposed to HPC in research through a contact at UCL and interviewed for a position in UCL – ARC. Having been successful, about a year later, was invited to become involved in the department's pilot apprenticeship scheme due to his unique experience.

### Meet the speakers...

Sadie Bartholomew  
University of Reading / NCAS

**Reducing the impact of energy consumption from computing with CATS, The Climate Aware Task Scheduler**

Thursday 5 December  
10:00 – 10:30

The CIUK Auditorium



**Abstract:** As well as reducing computational energy consumption, it's important to consider its carbon cost. Performing compute when electricity has lower carbon intensity, due to more renewable generation, will reduce environmental impact. This idea led to the inception of CATS, the Climate-Aware Task Scheduler, at the Software Sustainability Institute (SSI) Collaborations Workshop 2023. CATS is a Python package that schedules tasks based on the estimated carbon intensity

of the electricity grid, using real-time data from the UK's National Grid ESO API. Our small team continued developing CATS after the workshop, with support from SSI, and released Version 1.0, designed for use with the 'at' command hence targeting smaller-scale tasks on local machines, this July. Major work is being finalised for Version 2.0, which will integrate with Slurm for HPC applications. Ethical HPC must promote efforts to minimise carbon impact and CATS can contribute to this through intelligent time shifting of jobs.

**Bio:** Sadie is a Computational Scientist working for the Computational Modelling Services (CMS) group within the National Centre for Atmospheric Science (NCAS) and the Dept. of Meteorology at the University of Reading. Her work largely involves developing, optimising and maintaining open-source tools that support research and collaboration in earth science and aligned domains. Prior to this role which she has held since 2020, she spent

two years as a Scientific Software Engineer at the Met Office, following a master's degree in (particle) physics. In 2022 she was awarded a fellowship by the Software Sustainability Institute.

## Meet the speakers...

Jessica Huntley  
UKRI-STFC

**Tracking the Carbon  
Cost of Optimization  
Algorithms**

Thursday 5 December  
10:30 – 11:00

The CIUK Auditorium



**Abstract:** One of the biggest global challenges is tackling climate change and building a sustainable future. Whilst software on its own doesn't emit CO<sub>2</sub>, understanding the cost of running it is important so that the way it is developed and used can be optimised for minimal carbon emissions. There are several carbon tracking tools available for estimating hardware power consumption, allowing users and developers to quantify the carbon cost of running a program and identify areas for improvement. FitBenchmarking

(<https://fitbenchmarking.github.io/>) is an open source Python package that allows users to compare the performance of different optimization algorithms and their implementations across various fitting problems, through metrics such as runtime and accuracy. In this talk, I will outline how a new carbon emissions metric has been added to FitBenchmarking, some of the challenges faced along the way, and the insight that reporting emissions offers developers and users.

**Bio:** Jessica works as a Mathematical Research Software Engineer in the Scientific Computing Department's Computational Mathematics Theme. Alongside her day-to-day role, she also leads the department's Energy Efficiency Team, a working group which aims to help staff reduce their carbon footprint.

## Meet the speakers...

Vassil Alexandrov and Adriano  
Agnello (Hartree Centre STFC)

**Harnessing AI to Address Grand  
Challenges in Scientific Innovation**

Thursday 5 December  
11:30 – 12:00

The CIUK Auditorium

**Abstract:** In this talk we will outline the Hartree Centre's engagement with industry, public sector and research communities in driving scientific innovation based on extreme scale computing and scalable AI solutions. In particular, we will discuss our partnerships and collaborations with UKAEA and US National Labs such as ORNL, LLNL, LBL and ANL to illustrate how we are working to address Grand Challenges with high societal impact. There are many interdisciplinary applications this

work can address, and we will showcase our work in AI to bypass simulation bottlenecks in magnetic-confinement fusion and materials discovery, and AI for effective elicitation of knowledge.

**Bio: Prof. Vassil Alexandrov (VA)**, Chief Science Officer, Hartree Centre – STFC. His focus is currently on the enhancement and implementation of the comprehensive Hartree Centre Research Strategy that enables to employ advanced HPC, Data Analytics, AI and quantum-classical methods and algorithms to address key industrial, societal and scientific challenges. He is also bringing key expertise in computational science, extreme scale computing, scalable fault-tolerant and resilient algorithms for advanced computer architectures and

Monte Carlo methods and algorithms, applied to diverse application areas and extreme scale (peta and exascale) architectures. He leads STFC's collaboration with UKAEA, based at Hartree Centre, and co-leads with Dr. Rob Akers (UKAEA) the UKAEA -Hartree Centre STFC Fusion Computing Lab. He led Hartree Centre's collaboration with US ECP as part of UK and US National Labs collaboration in exascale computing and now is leading Hartree Centre's collaboration with US national labs (LLNL and ORNL) in extreme scale computing and scalable AI.

**Bio: Dr Adriano Agnello (AA)** is a Principal AI Researcher at the Hartree Centre, supervising research projects in AI for Applied Sciences and Engineering, including a multi-year workstream in collaboration with the UKAEA and the Intelligent Observatory ISPF project with the South African Astronomical Observatory. Together with the AI Group Leader, he liaises with colleagues at US National Labs on common AI initiatives, including workflows for self-driving labs, simulation orchestration, and magnetic confinement fusion. Besides his work in AI (including contributions to the STFC AI/DR Strategy and ETP4HPC SRA6), Dr Agnello has led searches of rare cosmic milestones and worked on spectroscopic instrument design and recommissioning projects.



**Meet the speakers...**

Noel O'Connor

**Partnering to take LLM to the next level**

Thursday 5 December  
12:00 – 12:30

The CIUK Auditorium

**Abstract: Background:** Today, many industries manage large amounts of data and, as their requirements grow, the technology that supports them seeks to keep pace. Large language models (LLMs) such as ChatGPT step up to the challenge by providing a user-friendly way for organisations to interact with their data at the speed of business to automate processes, provide valuable data insights and create customized content based on their data. LLMs are a form of artificial intelligence (AI) trained on massive sets of text-based data to perform various tasks ranging from

detecting patterns to crafting original content.

Currently, LLMs are flourishing in the open with limitations and without efficiency and the gains of traditional open source software. Excitement about the potential of this technology is high, but it's met with equal measures of concern. Across industries, a consensus exists that data is a precious asset organisations should handle with care. As a result, worries about privacy, inaccurate data, biases and leaks of sensitive information are at the centre of the conversation.

**Objective:** Use the power of collaborative partnering to and end to end solution which will deliver open-source concepts to models and the tools for open evolution. In addition, enable community users and developers to create and merge contributions to LLMs, while increasing efficiency and performance.

**Methods:** CDW present a Red Hat offering, Instruct LAB where users can learn & experiment via limited desktop-scale training method on small datasets. which includes models and capabilities for model evolution and serving, packaged as immutable instances and tuned to specific architectures, with support, indemnification and lifecycle necessary for enterprise deployment. This is deployed onto Dell hardware, a world leader in technology that is trusted and compatible.

Companies are not comfortable using large language models due to the unknown license of the outputted text. RHEL AI includes indemnification and the model itself is licensed under the Apache-2.0 Open Source license.

**Results:** A fully strategic and holistic method for enabling learners and researchers to develop training of LLMs from a laptop. With CDW delivering InstructLab\* and the Granite models, we make the ability to add knowledge and skills to the language model accessible to everyone, allowing companies to add their specific knowledge and skills to align the model to their specific use case. The InstructLab provides a way for people to get started in a

cost effective way and scale out when needed under the guidance of CDW, reliability of Dell and the power of Nvidia .

**Conclusion:** Utilising the skills and expertise of complimentary technologies and services to deliver an innovative and ever evolving LLM with real world impact from day one.

**Bio:** Noel O'Connor is a Senior Principal Architect in Red Hat's EMEA Solutions Practice specializing in cloud-native application and integration architectures. Since 2008, he has worked with many of Red Hat's global enterprise customers in both EMEA and Asia. He's co-authored a number of books, including "DevOps with OpenShift" and "DevOps Culture and Practice with OpenShift."



**Meet the speakers...**

Simon McIntosh-Smith (University of Bristol)  
Priya Sharma (DSIT)  
Christopher Edsall (University of Cambridge)

**The National AI Research Resource (AIRR) – The journey to launch**

Thursday 5 December  
12:30 – 13:00

The CIUK Auditorium

**Abstract:** Join the Universities of Bristol and Cambridge, the Department of Science, Innovation and Technology (DSIT) and UK Research and Innovation (UKRI) to hear about the launch of the national AI Research Resource (AIRR). Find out about the journey so far, from construction and deployment of Dawn at Cambridge and Isambard-AI at Bristol to early access use cases and their experiences and lessons learnt, as well as the future outlook.

**Bio:**

**Professor Simon McIntosh-Smith** is the founder and Director of the Bristol Centre for Supercomputing, which runs the UK's Isambard-AI service. He began his career in industry as a microprocessor architect, first at Inmos and STMicro in the 1990s, before co-designing the world's first fully programmable GPU at Pixelfusion in 1999. In 2002 he co-founded ClearSpeed Technology where, as Director of Architecture and Applications, he co-developed the first modern many-core HPC accelerators. He previously founded the HPC Research Group in Bristol, where his research interests include advanced computer architectures and performance portability.

**Priya Sharma** is the Compute Delivery Lead at the Department for Science, Innovation and Technology (DSIT), overseeing the delivery of high-performance computing projects, particularly the AI Research Resource project in collaboration with the University of Bristol and the University of Cambridge. Since joining the Civil Service in 2016, Priya has successfully led and delivered complex digital infrastructure projects, including the UK Telecoms Lab (UKTL) and the UK Emergency Alerts Project.

**Chris Edsall** is the Head of Research Software Engineering at Research Computing Services, a co-director of the Institute of Computing for Climate Science, and a bye-fellow of Queens' College Cambridge. He leads several teams of research software engineers and HPC consultants with the goal of delivering better software enabling better research.

Chris studied physics at the University of Canterbury and then worked in several research institutions in New Zealand (NIWA) and the UK (National Oceanography Centre and the University of Bristol) administering their HPC systems and upskilling researchers in software engineering to make best use of the supercomputing facilities.

## Meet the speakers...

Martyn Guest  
ARCAA, University of Cardiff

**Community Code  
Performance on Multi-core  
Processors. An Analysis of  
Computational Chemistry  
and Ocean Modelling  
Applications**

Thursday 5 December  
16:15 – 17:00  
The CIUK Auditorium



**Abstract:** This session will overview the measured performance of a number of popular community codes on a variety of HPC systems, with our established analysis based on both computational chemistry and ocean modelling applications. The former features codes from Molecular Dynamics (DL\_POLY, AMBER, and GROMACS), molecular electronic structure (GAMESS-UK) and Materials Science (VASP), while NEMO is the representative code from the ocean modelling community.

The variety of systems considered focus on both the Intel Emerald Rapids and recently released Granite Rapids processors, together with the AMD EPYC Genoa family of CPUs. Using the Intel Skylake Gold 6148 and AMD EPYC Rome 7502 as the baselines, a performance assessment is made across a variety of Granite Rapids (6980, 6972 and 6960) and Emerald Rapids (8562Y, 8568Y and 8592) SKUs, with system interconnects from both NVIDIA Networks and Cornelis Networks. Attention is also focused on systems featuring the AMD Genoa EPYC processors, including the Genoa 32-core 9354, the 48-core 9454 and 96-core 9654 and 9684X SKUs.

The benefits of the Intel® oneAPI Toolkit and Linaro Performance Reports are demonstrated throughout this analysis. To best capture a 'like for like' comparison amidst the extensive array of core densities, our analysis remains based on both a "node-by-node" and the more traditional "core-by-core" consideration.

**Bio:** Professor Martyn Guest has led a variety of high performance and distributed computing initiatives in the UK. He spent three years as Senior Chief Scientist and HPC Chemistry Group Leader at PNNL, before returning to the UK as Associate Director of Daresbury's Computational Science and Engineering Department. Martyn joined Cardiff University in April 2007 as their Director of Advanced Research Computing, retaining this position until February 2023. He is also Technical Director of the Supercomputing Wales programme and is co-I on the Isambard-2 and Isambard-3 systems at the GW4 Tier-2 HPC regional centre.


# HANDS-ON LEARNING

## WITH NVIDIA DLI TRAINING AT BOSTON LIMITED

Get practical experience with deep learning through Boston Limited's certified NVIDIA DLI Training. Led by expert instructors, our courses cover everything from foundational techniques to advanced applications, helping you and your team build essential skills for today's data-driven challenges.

**Scan the QR code** to see how Boston's training has made a difference for a prominent UK-based science and technology research organisation!

**JOIN US AND SEE WHAT YOU  
CAN ACHIEVE WITH NVIDIA DLI TRAINING**



## Meet the speakers...

Mayank Kumar  
UKRI-STFC

**Integration of SYCL with  
MPI in Multi-scale  
Universal Interface (MUI)  
library for Multi-Physics  
Coupling at Exascale**

Friday 6 December  
09:30 – 10:00

The CIUK Auditorium



**Abstract:** The Multiscale Universal Interface (MUI) is a versatile self-contained library designed to facilitate multi-physics and multi-scale simulations by providing an architecture that enables users to achieve seamless integration of solvers without requiring in-depth knowledge of their underlying topologies. This abstraction simplifies the coupling process, enabling researchers to focus on their scientific inquiries rather than the technical complexities of solver interaction. This work focuses on the improvements in portability

and performance achieved by leveraging Message Passing Interface (MPI) and SYCL to handle computationally expensive parts of the coupling algorithm.

**Bio:** Mayank is a computational scientist working at STFC Daresbury Laboratory, with a background in high-performance computing, computational fluid dynamics and code coupling. Currently working on the heterogeneous parallelization and optimization of codes used within the computational engineering group.

## Meet the speakers...

Gregory Tourte  
University of Oxford

**Use of WEKA FS for fast  
ephemeral storage  
powering academic HPC**

Friday 6 December  
10:00 – 10:30

The CIUK Auditorium



**Abstract:** As part of this year's University of Oxford Advanced Research Computing (ARC) storage infrastructure upgrade, we implemented a solution using a 275 TiB Weka filesystem for our fast scratch area in our HPC and HTC infrastructure, replacing the existing GPFS solution. This is an ephemeral data storage area with the primary requirement of extremely high-speed data read and write operations, supporting our ARC, HTC, and Jade clusters. In this talk we present an overview of our requirements, hardware and setup

design. We present benchmark results from our raw tests as well as the results from evaluating real world applications alongside our users in several disciplines, each of whom are working with applications where I/O is critical.

**Bio:** Gregory J. L. Tourte is a Research Systems Administrator in the Advanced Research Computing Team at the University of Oxford which provides and supports high performance computing facilities for researchers throughout the institution. Prior to joining Oxford, he worked as a Senior Research Associate at the University of Bristol in the Digital Health Department and School of Geographical Sciences where he started as a research software engineer and system administrator for the research group's supercomputer and used this opportunity to develop an understanding of the research data management required for the large quantity of data being generated. He holds an undergraduate Master's in Mechanical Engineering and a Master's in Computer Science from the University of Bath.



## Meet the speakers...

Oliver Brown  
EPCC, University of Edinburgh

### The Scalability of Quantum Air Traffic Control

Friday 6 December  
10:30 – 11:00

The CIUK Auditorium



**Abstract:** We describe a “top-down” approach for applying quantum computing to aviation; specifically, the problem of routing airplanes through an idealised airspace such that the carbon emissions associated with air travel are minimised. We follow a graph-based model, whereby airspace sectors are represented by vertices connected by weighted edges. The edge weights determine the cost of travelling through the airspace. The problem is naturally scalable in that one can adjust the number of vertices and edges, from a simple 3-

by-3 lattice digraph with 24 edges to an irregular digraph representing the UK airspace containing 80 edges. Finding the minimum cost path can of course be solved using classical techniques (e.g., Dijkstra, Bellman-Ford), but our concern here is to measure the scalability of a quantum-based approach, one that uses a Quadratic Unconstrained Binary Optimisation (QUBO) formalism suitable for a D-Wave quantum annealer. The QUBO formalism makes it possible to represent within a single matrix, the airspace network, the entry and exit sectors for an airplane requesting to enter the airspace, as well as the level and location of air traffic congestion at the time of the request.

**Bio:** Dr Oliver Thomson Brown is a Chancellor’s Fellow at EPCC and leads their Quantum Group. His research interests are all things quantum and HPC, including applications of quantum computing, classical emulation of quantum computers, and programming models for hybrid quantum-HPC. He is also a member of Edinburgh’s Quantum Software Lab.

## Meet the speakers...

Duncan McBain  
Codeplay Software

### Portable SYCL code using oneMKL on AMD, Intel and Nvidia GPUs

Friday 6 December  
11:30 – 12:00

The CIUK Auditorium



**Abstract:** In this session we will show you how to make your software run faster using oneMKL, an accelerated math library for AMD, Intel and Nvidia GPUs. oneMKL is based on the oneAPI specification and can be used to target multi-vendor and multi-architecture accelerators from a single code base. It is also now governed by the Unified Acceleration Foundation (UXL), an open governance body that is part of the Linux Foundation.

**Bio:** Duncan has worked with SYCL and GPUs at Codeplay Software after completing an MSc in high performance computing at the University of Edinburgh. Currently he is working with Intel's SYCL implementation, helping to deploy it on HPC systems across the world and providing support to users of Codeplay's plugins for oneAPI. He has a continued interest in how we can further leverage GPU hardware to accelerate the software that is trying to solve the biggest problems tackled by HPC systems today.



# Open Architecture, Enterprise Solutions

INNOVATIVE SOFTWARE  
AND UPDATES, FAST

A SAFE, SUPPORTED  
ENVIRONMENT

A STABLE LONG-TERM  
RESEARCH PLATFORM

Visit us on  
**Stand #36**



Linkedin: @CDW-UK | X: @CDW\_UK

[uk.cdw.com](http://uk.cdw.com)

## Meet the speakers...

Dimitrios Bellos  
The Rosalind Franklin Institute

**Flowcron: Improving access to HPC via a Function-as-a-Service which utilizes a Globus Flow and a cron service**

Friday 6 December  
12:00 – 12:30

The CIUK Auditorium



**Abstract:** Over the last 10 years, data generation rates in structural biology and material science have exploded, increasing computing power requirements for techniques such as cryo-EM and X-ray CT. Many UK institutes, including the Rosalind Franklin Institute, rely on HPCs, like the Baskerville Tier 2 HPC. Our research produces multi-terabyte datasets that need to be analysed quickly, thus requiring high computing resources and fast data transfers. Furthermore, computationally demanding AI algorithms are becoming more

popular to process such datasets. In the Franklin we encourage all our scientists to adopt the use of our HPC provider, Baskerville. However, making HPC use more convenient for them is a significant challenge. To tackle this challenge, we would like to showcase FlowCron, a Function-as-a-Service solution that utilises a Globus Flow and a cron service, to increase automation and minimise the number of steps necessary to transfer and process data on Baskerville.

**Bio:** Dimitrios is a researcher and software developer, in the Artificial Intelligence and Informatics (AI & I) theme in Rosalind Franklin Institute. He is a member of the Franklin's Advanced Research Computing (ARC) team which specialises in developing data and compute infrastructure for biological scientists and offers a centralised Research Software Engineering (RSE) capability through collaboration with the Franklin and external scientists to provide excellent software for research. Furthermore, he researches into DL/ML approaches for the processing or enhancing of Electron Tomography and X-ray Computed Tomography data. He studied in the school of Electrical and Computer Engineering in the Aristotle University of Thessaloniki in Greece before completing his PhDs in Computer Science at the University of Nottingham. His main research interests are regarding the development of AI solutions with focus on denoising, classification, segmentation, deep/machine learning, computer vision and image processing.

## Meet the speakers...

Dave Bond, Andy Herdman and Neil Gasper  
AWE National Security Technologies

**AWE – Centres of Excellence, an outline of academic engagement**

Friday 6 December  
12:30 – 13:00

The CIUK Auditorium

**Abstract:** AWE is pursuing a centre of excellence programme, and the Scientific Computing Centre of Excellence, and Quantum Centre for Nuclear Defence and Security are two such centres out of the ten that will be established over the next few years covering all areas of science at AWE. The vision for a Scientific Computing Centre of Excellence (ScCoE) and Quantum Centre for Nuclear Defence and Security (QCND), is to create a focus of expertise in those technical areas of interest identified in AWE's scientific community. Bringing

together subject matter experts, external expertise, and academic institutions within the UK. To achieve these goals there is the need to engage with internal subject matter experts as well as the wider academic community via PhD sponsorship. This talk will give an overview to this scheme and will focus on the scope of the centres, how AWE intends to engage with the wider academic community and research opportunities these present.

**Bio: Dave Bond** - Dave is the AWE business lead for delivering the centre of excellence scheme at AWE Nuclear Security Technologies. Dave's background is in scientific computing with an emphasis on high-speed distributed storage systems; a discipline where he has spent most of his career. Dave has previously run an academic partners Horizon 2020 bid surrounding high-speed data placement aware object stores.

**Dr Andy Herdman** - Andy is the Principal Technical Authority for Scientific Computing at AWE Nuclear Security Technologies. He has worked in the HPC field for over 25 years. Roles have included Technical Team Leader, responsible for code porting and optimisation to then emerging Massively Parallel Processor supercomputers, Group Leader of multi-disciplined teams, covering applied computer science, advanced networking, storage, engineering codes, data visualisation, and the interception of emerging technologies.

**Dr Neil Gaspar** - is a Principal Applied Physicist working at AWE National Security Technologies. He lectured at the University of Exeter researching into the microstructure of materials then joined AWE in 2008. He was awarded chartered status as a Physicist and Mathematician and worked in Nuclear Threat Reduction building strategic academic partnerships. More recently Neil started the Quantum Technology group at AWE to investigate how new technologies can be adopted into the nuclear defence and security mission.

**Meet the speakers...**

Alastair Basden  
DiRAC/ Durham University

**The Durham HPC  
Hardware Laboratory**

Friday 6 December  
15:00 – 16:00

The CIUK Auditorium



**Abstract:** The Durham HPC Hardware Lab is hosted by the DiRAC COSMA HPC facility and provides UK researchers with access to cutting edge technologies and facilities, to allow testing and benchmarking of codes, software migration to new hardware, and study of new paradigms. This lab has grown in scope over the past few years, funded by ExCALIBUR H&ES, DiRAC, Durham and various UKRI grants. Of particular interest to many users is access to new GPU systems, novel networking topologies, composable

infrastructure, and access to BlueField DPUs. This talk presents the Hardware Lab, including information about how it can be accessed.

**Bio:** Alastair manages the DiRAC Memory Intensive HPC service, COSMA, at Durham University.

**Meet the speakers...**

Nick Brown  
EPCC at University of  
Edinburgh

**RISC-V: The most exciting  
hardware innovation of  
recent times**

Friday 6 December  
15:00 – 16:00

The CIUK Auditorium



**Abstract:** Powering over 20 billion devices since it was first released around a decade ago, the open, community driven, RISC-V Instruction Set Architecture (ISA) has enjoyed phenomenal growth especially in embedded computing. However, whilst RISC-V is yet to become widespread in HPC recent developments such as the high-core count SG2042 RISC-V CPU and several RISC-V accelerator cards have the potential to change that. From my perspective as leading the RISC-V HPC Special Interest Group (SIG) I will

introduce RISC-V, explore the potential impact it could have on HPC and describe how people can get involved and be a part of potentially the most exciting hardware revolution of our lives.

**Bio:** Dr Nick Brown is a Senior Research Fellow at EPCC, the University of Edinburgh. His main interest is in the role that novel hardware can play in future supercomputers and is specifically motivated by the grand-challenge of how we can ensure scientific programmers are able to effectively exploit such technologies without extensive hardware/architecture expertise. Combining novel algorithmic techniques for new hardware, programming language & library design, and compilers, he has over 80 peer reviewed publications and has worked on a number of large scale parallel codes. He is chair of the RISC-V HPC SIG, leads EPCC's RISC-V testbed.



**Meet the speakers...**

Tina Friedrich  
University of Oxford

**A first look at JADE 2.5**

Friday 6 December  
15:00 – 16:00

The CIUK Auditorium

**Abstract:** Acting on behalf of the JADE consortium, the University of Oxford has recently purchased and provisioned the JADE 2.5 system - a technology pilot to investigate the suitability of AMD Instinct Accelerators for Artificial Intelligence and Machine Learning (AI/ML) research. The JADE

2.5 system is operated by the University of Oxford's IT Services Advanced Research Computing (ARC) team; in this talk, we give an overview of the system hardware and

setup, as well as first experiences with the system.

**Bio:** Tina Friedrich is a Senior HPC Systems Administrator in the Advanced Research Computing (ARC) team at the University of Oxford, which provides access to High Performance Computing resources, support, and advice to researchers within the University of Oxford. Before joining Oxford, they worked at Diamond Light Source, managing HPC systems and contributing to various systems administration responsibilities. They hold a Master's in Physics from the University of Heidelberg.

## Delivering a secure future

### JOIN US AT CIUK 2024

**We'll be exhibiting on booth 43-44, and presenting on Friday 6 December at 12:30pm.**

You'll get to hear about AWE's new and future Centres of Excellence. Each centre provides strategic research and key services to our programme for both PhD and postdoctoral research projects. AWE is growing its science and engineering portfolio and we invite UK universities to be a part of this exciting opportunity.

Over the next few years we are establishing up to 10 centres to address key research areas, and to build a UK leading capability that will improve nuclear skills and expertise, and promote further collaboration with interest groups.

The Scientific Computing Centre of Excellence (ScCoE) and Quantum Centre for Nuclear Defence and Security (QCNDs) are two new CoEs, joining AWE's existing and established academic partnerships that support our mission.

#### **Find out more at CIUK 2024**

Our experts will be giving an overview of the CoEs, with insights into how AWE intends to engage with the wider academic community on many fronts that will deliver mutual benefit.




## PLATINUM SPONSORS

Computing Insight UK  
**2024**

**Find us at  
booth: 43-44**



@AWE\_plc



### Don't forget our networking events...

This year we have two networking events that are open to all registered CIUK delegates and exhibitors. Come along and join in the discussion with other CIUK attendees. Food and drink will be served at both events.



On Wednesday 4 December we will host a Pre-CIUK Networking Event at the [Gas Works Brew Bar](#), 5 Jack Rosenthal Street, M15 4RA. The event will run from 5:30pm until 11pm with food and drink available. The venue is a short walk from the Manchester Central Convention Complex.

**A CIUK lanyard and badge is required for entry**, and these can be collected from the registration desk at the convention centre between 9am and 5pm on Wednesday 4 December.



On Thursday 5 December the official CIUK 2024 Networking Event will take place at [Peaky Blinders Bar](#), 23 Peter Street, M2 5QR. The event will open at 6:30pm and run until 11pm. Food will be available from 6:30pm. **Your CIUK badge and lanyard is required for entry.**

We hope to see everyone there as we continue the conference discussions, initiate new partnerships and collaborations and celebrate another successful Computing Insight UK Conference.

# Data & Storage Management



PoINT products allow efficient usage of storage systems and help to reduce costs and issues caused by data growth. The software solutions fulfil compliance and archiving requirements and provide independence from storage technologies and vendors.

**S3-to-Tape**

**On-Prem Glacier**

**File Archiving  
and HSM**

**Backup and  
Replication**

**point.de**



**SOLVING COMPLEX TECHNOLOGY PROBLEMS, SUSTAINABLY.**

## **ASK US ABOUT**

**HPC CLUSTERS**

**SLURM**

**SupremeRAID™**

**VIRTUALISATION**

**COMPUTE SERVERS**

**STORAGE SERVERS**

**WORKSTATIONS**



[business.novatech.co.uk/hpc-cluster](https://business.novatech.co.uk/hpc-cluster)



[education@novatech.co.uk](mailto:education@novatech.co.uk)




Stand 1



## CIUK 2024 EXHIBITION

The [CIUK 2024 Exhibition](#) will take place in the Exchange Hall. The exhibition will include the exhibition stands, the poster competition, the cluster challenge and the research zone.

PLATINUM	 	
	 	
GOLD	  	
	  	
SILVER	   	
	   	
BRONZE	    	
	    	



# ADVANCING AI FROM CLOUD TO EDGE TO ENDPOINTS

## AMD ADVANCES END-TO-END AI INFRASTRUCTURE

### AMD AI STRATEGIC PILLARS

Broad portfolio of training & inference compute engines from CPUs, GPUs, & adaptive computing

Open & proven software capabilities help ensure leading AI frameworks, libraries, and models are fully enabled for AMD hardware

Expanding AI ecosystem with deep co-innovation with cloud providers, OEMs, software developers & AI leaders






HPC | CLOUD | ENTERPRISE | EMBEDDED | PC

## THE INDUSTRY'S BROADEST AI COMPUTE PORTFOLIO

### AMD AI COMPUTE PORTFOLIO

AMD Delivers the Broadest Technology Portfolio to the Data Center



Cloud & Enterprise	Cloud	Edge and Embedded	Endpoint	Gaming AI
AMD Instinct™ MI300X Accelerators	AMD EPYC™ Processors	AMD Versal™ AI Edge Series	AMD Ryzen™ 8040 Series	AMD Radeon™ RX/PRO 7000 Series
Data Center Training & Inference	CPU AI Leadership	AI + Sensor Embedded Inference	AI inference for Windows PCs	Gaming AI

### STRONG OPEN AI SOFTWARE AND ECOSYSTEM MOMENTUM



**Hugging Face**

62,000+ Models Running Nightly Fully Integrated Optimum Library

**PyTorch**

From 'port-to' to 'develop-on' with latest platforms













Expanding Open Source Contributions & Footprint


The information contained herein is for informational purposes only, and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of non-infringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. GD-18











Use or mention of third-party marks, logos, products, services, or solutions herein is for informational purposes only and no endorsement by AMD is intended or implied. GD-83






© 2024 Advanced Micro Devices, Inc. all rights reserved. AMD, the AMD arrow, EPYC, Instinct, Infinity Fabric, Radeon, ROCm, Ryzen, and combinations thereof, are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective owners.










	<p><b>Web:</b> <a href="http://www.alces-flight.com">www.alces-flight.com</a>  <b>X:</b> <a href="https://twitter.com/alcesflight">@alcesflight</a></p>
	<p><b>Web:</b> <a href="http://www.altair.com">http://www.altair.com</a></p>
	<p><b>Web:</b> <a href="http://www.amd.com">http://www.amd.com</a>  <b>X:</b> <a href="https://twitter.com/AMD">@AMD</a></p>
	<p><b>Web:</b> <a href="http://www.atempo.com">www.atempo.com</a>  <b>X:</b> <a href="https://twitter.com/Atempo">@Atempo</a></p>
	<p><b>Web:</b> <a href="http://www.awe.co.uk">www.awe.co.uk</a>  <b>X:</b> <a href="https://twitter.com/AWE_plc">@AWE_plc</a></p>
	<p><b>Web:</b> <a href="https://thinkparq.com/">https://thinkparq.com/</a>  <b>X:</b> <a href="https://twitter.com/beegfs">@beegfs</a></p>
	<p><b>Web:</b> <a href="https://bios-it.com/">https://bios-it.com/</a>  <b>X:</b> <a href="https://twitter.com/Bios_it">@Bios_it</a></p>
	<p><b>Web:</b> <a href="http://www.boston.co.uk">www.boston.co.uk</a>  <b>X:</b> <a href="https://twitter.com/bostonlimited">@bostonlimited</a></p>
	<p><b>Web:</b> <a href="http://www.cornelisnetworks.com">www.cornelisnetworks.com</a>  <b>X:</b> <a href="https://twitter.com/CornelisHPC">@CornelisHPC</a></p>
	<p><b>Web:</b> <a href="https://www.uk.cdw.com/">https://www.uk.cdw.com/</a>  <b>X:</b> <a href="https://twitter.com/CDW_UK">@CDW_UK</a></p>

	<p><b>Web:</b> <a href="http://www.datacore.com">www.datacore.com</a>  <b>X:</b> <a href="https://twitter.com/DataCore">@DataCore</a></p>
	<p><b>Web:</b> <a href="http://www.datavita.co.uk">www.datavita.co.uk</a></p>
	<p><b>Web:</b> <a href="http://www.ddn.com">www.ddn.com</a>  <b>X:</b> <a href="https://twitter.com/DDNStorage">@DDNStorage</a></p>
	<p><b>Web:</b> <a href="http://www.define-technology.com">www.define-technology.com</a>  <b>X:</b> <a href="https://twitter.com/definetechnology">@definetechnology Ltd</a></p>
	<p><b>Web:</b> <a href="http://www.digitalrealty.com">www.digitalrealty.com</a>  <b>X:</b> <a href="https://twitter.com/digitalrealty">@digitalrealty</a></p>
	<p><b>Web:</b> <a href="http://www.dell.com/en-uk">http://www.dell.com/en-uk</a>  <b>X:</b> <a href="https://twitter.com/Dell">@Dell</a></p> <p><b>Web:</b> <a href="http://www.intel.com">www.intel.com</a>  <b>X:</b> <a href="https://twitter.com/IntelBusiness">@IntelBusiness</a></p>
	<p><b>Web:</b> <a href="http://www.doit-now.tech">www.doit-now.tech</a>  <b>X:</b> <a href="https://twitter.com/DoItNow_Tech">@DoItNow_Tech</a></p>
	<p><b>Web:</b> <a href="http://www.eviden.com">www.eviden.com</a>  <b>X:</b> <a href="https://twitter.com/evidenlive">@evidenlive</a></p>
	<p><b>Web:</b> <a href="http://www.groq.com">www.groq.com</a>  <b>X:</b> <a href="https://twitter.com/groqinc">@groqinc</a></p>

 <p><b>HAMMERSPACE</b> DATA IN MOTION</p>	<p><b>Web:</b> <a href="http://www.hammerspace.com">www.hammerspace.com</a></p>
 <p><b>Hewlett Packard Enterprise</b></p>	<p><b>Web:</b> <a href="http://www.hpe.com">www.hpe.com</a> <b>X:</b> <a href="https://twitter.com/HPE_UK">@HPE_UK</a></p>
 <p><b>IBM</b></p>	<p><b>Web:</b> <a href="https://www.ibm.com/uk-en/storage">https://www.ibm.com/uk-en/storage</a></p>
 <p><b>KAYTUS</b></p>	<p><b>Web:</b> <a href="https://www.kaytus.com/">https://www.kaytus.com/</a> <b>X:</b> <a href="https://twitter.com/KAYTUS">@KAYTUS</a></p>
 <p><b>KAO DATA</b></p>	<p><b>Web:</b> <a href="http://www.kaodata.com">www.kaodata.com</a> <b>X:</b> <a href="https://twitter.com/KaoDataDC">@KaoDataDC</a></p>
 <p><b>Lenovo</b></p>	<p><b>Web:</b> <a href="http://www.lenovo.com">http://www.lenovo.com</a> <b>X:</b> <a href="https://twitter.com/Lenovo">@Lenovo</a></p>
 <p><b>LOGICALIS</b> Architects of Change</p>	<p><b>Web:</b> <a href="https://www.uki.logicalis.com/">https://www.uki.logicalis.com/</a> <b>X:</b> <a href="https://twitter.com/LogicalisUK">@LogicalisUK</a></p>
 <p><b>NAG</b></p>	<p><b>Web:</b> <a href="http://www.nag.com">www.nag.com</a> <b>X:</b> <a href="https://twitter.com/NAGTalk">@NAGTalk</a></p>
 <p><b>novatech</b></p>	<p><b>Web:</b> <a href="https://business.novatech.co.uk/">https://business.novatech.co.uk/</a></p>
 <p><b>OCF</b></p>	<p><b>Web:</b> <a href="http://www.ocf.co.uk">www.ocf.co.uk</a> <b>X:</b> <a href="https://twitter.com/ocf_hpc">@ocf_hpc</a></p>

	<p><b>Web:</b> <a href="https://par-tec.com/">https://par-tec.com/</a>  <b>X:</b> <a href="#">@ParTec_AG</a></p>
	<p><b>Web:</b> <a href="http://www.peakaiio.com">www.peakaiio.com</a></p>
	<p><b>Web:</b> <a href="https://www.pny.com/">https://www.pny.com/</a>  <b>X:</b> <a href="#">@PNYTechnologies</a></p>
	<p><b>Web:</b> <a href="https://www.point.de/">https://www.point.de/</a></p>
	<p><b>Web:</b> <a href="http://www.qumulo.com">www.qumulo.com</a>  <b>X:</b> <a href="#">@qumulo</a></p>
	<p><b>Web:</b> <a href="http://www.quobyte.com">www.quobyte.com</a>  <b>X:</b> <a href="#">@Quobyte</a></p>
	<p><b>Web:</b> <a href="http://www.scan.co.uk/business/scan-it">www.scan.co.uk/business/scan-it</a>  <b>X:</b> <a href="#">@ScanComputers</a></p>
	<p><b>Web:</b> <a href="http://www.stackhpc.com">www.stackhpc.com</a>  <b>X:</b> <a href="#">@stackhpc</a></p>
	<p><b>Web:</b> <a href="https://starfishstorage.com/">https://starfishstorage.com/</a>  <b>X:</b> <a href="#">@StarfishStorage</a></p>

	<p><b>Web:</b> <a href="https://www.storj.io/">https://www.storj.io/</a>  <b>X:</b> <a href="#">@storj</a></p>
	<p><b>Web:</b> <a href="http://www.vspl.co.uk">www.vspl.co.uk</a>  <b>X:</b> <a href="#">@vspl_solutions</a></p>
	<p><b>Web:</b> <a href="http://www.viridiengroup.com">www.viridiengroup.com</a>  <b>X:</b> <a href="#">@weareviridien</a></p>
	<p><b>Web:</b> <a href="http://www.vespertec.com">www.vespertec.com</a>  <b>X:</b> <a href="#">@Vespertec</a></p> <p><b>Web:</b> <a href="https://www.quantum.com/">https://www.quantum.com/</a>  <b>X:</b> <a href="#">@QuantumCorp</a></p>
	<p><b>Web:</b> <a href="https://www.vdura.com/">https://www.vdura.com/</a>  <b>X:</b> <a href="#">@VDURAInc</a></p>
	<p><b>Web:</b> <a href="https://www.weka.io">https://www.weka.io</a>  <b>X:</b> <a href="#">@wekaio</a></p>
	<p><b>Web:</b> <a href="http://WWW.WDC.COM">WWW.WDC.COM</a>  <b>X:</b> <a href="#">@westerndigital</a></p>



# INDUSTRIAL CLOUD SOLUTIONS FOR AI AND HPC PRODUCTION

At Viridien, we leverage >70 years of industrial computing expertise to drive scientific discovery, advanced simulations and transformative insights across industries.

[viridiengroup.com/hpc](https://viridiengroup.com/hpc)

SEE THINGS DIFFERENTLY

 VIRIDIEN



## CIUK 2024 DAY ZERO

Following the success of [CIUK DAY ZERO](#) in 2023, we will again be offering attendees at CIUK 2024 the opportunity to attend extra activities for no extra cost!

On Wednesday 4 December there will be a number of different activities taking place that will be available to attendees of CIUK 2024. Some will be included with your CIUK 2024 registration, others will require a separate registration. Full details of all the sessions can be found below.

Wednesday 4 December	<b>CoSeC Annual Conference 2024</b>	10:00 - 17:00	CIUK Breakout Room
Wednesday 4 December	<b>Technical/SysAdmin Meetup</b>	13:00 - 16:00	Exchange Room 6 and 7 (Upstairs at Manchester Central)
Wednesday 4 December	<b>The Lustre User Group UK 2023</b>	14:00 - 17:00	Exchange Rooms 9 and 10 (Upstairs at Manchester Central)
Wednesday 4 December	<b>STEP-UP - developing HPC technical professionals</b>	16:00 – 17:00	Exchange Room 6 and 7 (Upstairs at Manchester Central)



**The CoSeC Annual Conference 2023** will take place on Wednesday 6 December in the CIUK Breakout Room. It will run from 10:00 to 17:00 and will include updates on everything that has taken place within the [CoSeC](#) programme over the last twelve months, as well as a forward look to what the future holds for CoSeC.

More details about the conference can be found on their [website](#).

**PROGRAMME now available to view and download.**



### **Technical/SysAdmin Meetup**

**Wednesday 4 December - 13:00-16:00 - Exchange Room 1**

A meeting for technical staff and systems administrators to discuss the latest topics in the field and share information and ideas. This will be an unconference-style event, where the technical staff are able to steer the discussion to a variety of topics of interest. This meeting has been very popular for the last two years, and we expect a similar level of interest this year. If you have any topics of particular interest you feel strongly about, [feel free to register them in advance](#).



**The Lustre User Group UK 2023** will be hosted alongside CIUK in Manchester on 4th December 2024. This event brings together the UK Lustre Community, with keynote speakers from Lustre open source contributors, engineers and partners, sharing their experience and technology roadmap updates.

The User Group Meeting will run from 12:30 to 17:00 and will take place in Exchange Rooms 9 and 10 (*upstairs at Manchester Central*).

**Registration is FREE via the LUG UK website.**

*Please note this registration is separate from the general CIUK 2024 registration. If you are planning to attend any of the CIUK 2024 activities you must also register for [CIUK 2024](#).*



### **STEP-UP - developing HPC technical professionals**

**Wednesday 4 December - 16:00-17:00 - Exchange Room 1**

The STEP-UP project aims to change the landscape for digital Research Technical Professionals (dRTPs) within the London and South East of England region, while sharing developments, approaches and findings with the wider community, across the UK and beyond. We class dRTPs as anyone undertaking work related to software, data and computing infrastructure / High Performance Computing (HPC) within the research community. In this session we'll focus on the HPC-related aspects of STEP-UP's activities. How

did you become an HPC technical professional? How do you progress as an HPC technical professional? What's missing to support the next generation of HPC experts? Share your journeys and learn about the STEP-UP project and how we're looking to support the HPC community. The session will begin with a couple of short introductory presentations to set the scene. However, this will be an interactive session where we would like to hear from you about the routes you have taken your journeys into your roles, your learning journeys and your perspectives on career opportunities within the HPC community. The information shared in this session will help us in our identification of career and learning pathways for HPC professionals' careers and the skills and career professional development requirements. As part of its activities over the next few years, the STEP-UP project will also be setting up a programme to support short-term secondments to help technical professionals develop their specialist expertise and share skills, and creating a mentoring programme. We want to know how you think you could benefit from such opportunities and what you'd like to see within these schemes? What can we offer you to help further develop the HPC community, address challenges around skills shortages and recruitment, and provide more opportunities for technical engagement across institutions. We'll close with a summary, thoughts around next steps and an opportunity for participants to provide their contact details to engage further with us or receive updates on the work of STEP-UP.

Stack  
HPC

# The Convergence of HPC, AI and Cloud

Delivering:

- ◆ Simple self-service platforms for HPC and AI
- ◆ Secure environments for trusted research
- ◆ Open Cloud Infrastructure for Research Computing
- ◆ The performance of HPC with the flexibility of Cloud

StackHPC.com



## ENABLING YOU STORAGE TO DO MORE

-  Ensure **high availability** of data
-  Modernize your infrastructure
-  Unlock the benefits of HCI
-  Protect data against ransomware
-  Preserve data on cost-effective active archive

datacore.com

## CIUK 2024 PARALLEL BREAKOUT SESSIONS

The [CIUK 2024 Breakout Sessions](#) will run alongside the main CIUK programme in Charter Room 4 and Exchange room 1. Look for the **PINK** pull-up banners.

### Technical/SysAdmin Meetup

Wednesday 4 December - 13:00-16:00 - Exchange Room 1

A meeting for technical staff and systems administrators to discuss the latest topics in the field and share information and ideas. This will be an unconference-style event, where the technical staff are able to steer the discussion to a variety of topics of interest. This meeting has been very popular for the last two years, and we expect a similar level of interest this year. If you have any topics of particular interest you feel strongly about, feel free to register them in advance.

### STEP-UP - developing HPC technical professionals

Wednesday 4 December - 16:00-17:00 - Exchange Room 1

The STEP-UP project aims to change the landscape for digital Research Technical Professionals (dRTPs) within the London and South East of England region, while sharing developments, approaches and findings with the wider community, across the UK and beyond. We class dRTPs as anyone undertaking work related to software, data and computing infrastructure / High Performance Computing (HPC) within the research community. In this session we'll focus on the HPC-related aspects of STEP-UP's activities. How did you become an HPC technical professional? How do you progress as an HPC technical professional? What's missing to support the next generation of HPC experts? Share your journeys and learn about the STEP-UP project and how we're looking to support the HPC community. The session will begin with a couple of short introductory presentations to set the scene. However, this will be an interactive session where we would like to hear from you about the routes you have taken your journeys into your roles, your learning journeys and your perspectives on career opportunities within the HPC community. The information shared in this session will help us in our identification of career and learning pathways for HPC professionals' careers and the skills and career professional development requirements. As part of its activities over the next few years, the STEP-UP project will also be setting up a programme to support short-term secondments to help technical professionals develop their specialist expertise and share skills, and creating a mentoring programme. We want to know how you think you could benefit from such opportunities and what you'd like to see within these schemes? What can we offer you to help further develop the HPC community, address challenges around skills shortages and recruitment, and provide more opportunities for technical engagement across institutions. We'll close with a summary, thoughts around next steps and an opportunity for participants to provide their contact details to engage further with us or receive updates on the work of STEP-UP.

### Portable benchmarking and profiling using ReFrame

Thursday 5 December - 09:30-13:00 - CIUK Breakout Room

Testing and benchmarking scientific applications on HPC systems still often relies on manual effort and specialist knowledge. This makes studying application performance across HPC systems time consuming and error prone. In this session we show how building, running and measuring the performance of applications can be automated to develop benchmarks across HPC systems using ReFrame. ReFrame is widely used for regression testing by HPC services from small group clusters up

to the largest supercomputers in the world (including the UK national supercomputer, ARCHER2). Together with Spack and lightweight profiling tools, we have used ReFrame to build an automated benchmarking pipeline for the ExCALIBUR Hardware and Enabling Software program for cross-platform benchmarking. The workshop will introduce using ReFrame to automate benchmarking and collecting performance data. Participants will learn to write a simple ReFrame test, run it on a HPC platform, measure its performance and compare against achievable peak performance of the system. Following the session, attendees should have enough knowledge to start using ReFrame in their day-to-day activities. The workshop is aimed at attendees who have an understanding of basic HPC concepts and are interested in the performance of scientific applications. This includes both system administrators who run benchmarks to monitor the performance of a system, as well as application developers who are interested in performance benchmarks in their development workflow. Attendees will need to bring a laptop with an SSH client installed to access a remote facility to develop and run the ReFrame tests. A working knowledge of Python is required to develop tests in ReFrame. We will provide a repository with setup instructions and examples that will be communicated to attendees before the session starts and during the session. The workshop will be taught by the developers from the ExCALIBUR benchmarking project, and ARCHER2 service staff who have experience of using ReFrame in production. We will provide access to ARCHER2 and example applications to work on. The workshop is a mix of short presentations, carpentries-style live coding, and hands on sessions with helpers.

[Cybersecurity and Federation for National DRI, AI and HPC Resources](#)  
[Thursday 5 December - 14:00-16:00 - CIUK Breakout Room](#)

The threat faced by the UK research and education community from cyberattack is persistent, with well publicised incidents against members of our community both nationally and internationally. Digital Research Infrastructures involving these organisations form a critical part of the supply chain of the UK and thus we must evolve a collective, collaborative approach to our defence in the face of this threat. Organisers of this session include representatives of the DRI Cybersecurity community including JISC and national HPC and AI Research Resource providers. With the HPC focus of CIUK, we seek to bring together a community of practice to share ongoing activities and engage with the wider CIUK participants. In parallel, the research communities have ever-increasing needs for securely sharing data and accessing powerful computing resources within a federated ecosystem; the impact of the risks involved with these requirements must be carefully assessed so that we can provide secure, assured access to our national resources. A tentative outline of the session:

- Welcome and introduction
- Overview from the organisers (DRI Cybersecurity, JISC, and AIRR)
- Cybersecurity topics
- AAI topics
- Open discussion
- Close and next steps/events.

Women in HPC Breakfast: Organiser Cristin Merrit (Alces Flight)  
Friday 6 December - 08:00-10:30 - CIUK Breakout Room



CIUK is thrilled to host the 3rd Annual [Women in HPC](#) Breakfast on **Friday 6 December!** Join us starting at 8:00 AM for breakfast, with presentations kicking off at 8:30 AM. This year's event features the exciting finale of *Move the Needle* - a 12-month project on equity, diversity, inclusion, and accessibility (EDIA).

*Move the Needle* brought together nine participants from the UK and US to set impactful EDIA goals, with the project evolving along the way as the HPC/AI community added their own contributions. Three of those team members will take the stage and share their experiences, with an overview presented by the team project manager. Following the presentation we will host an open forum with attendees.

This breakfast is included with your conference pass. To see the full agenda, [click here](#), and to learn more about Move the Needle, [click here](#).

Storage Scale (GPFS) User Group / New Users Session  
Friday 6 December - 10:30-13:30 - CIUK Breakout Room

GPFS User Group CIUK 2024 Planned to be 2 hours, but with a new User session extending it to three hours. Continuing on from user groups at CIUK from a number of years, we want to run another Storage Scale User group at CIUK 2024. The user group will combine user talks, Vendor presentations and IBM engineer Talks around the user of and future directions of Storage Scale. The user group brings together representation from media, academia, finance, research, automotive, defence and pharmaceutical industries. We would also like to continue with the new user sessions that have been successfully run at previous user group meetings including at the event hosted at IBM in June 2024. These sessions have been well received and are helping to build expertise for the future to fill the admin roles required to run these systems. The User group continues to be run by users for users, and is not run by IBM, but we work closely with IBM. The Spectrum Scale User Group aims to:

- Bring together users of Spectrum Scale and Spectrum Scale with Spectrum Protect ILM (TSM HSM) into a collective environment
- Represent the needs of the Spectrum Scale User Group members
- Liaise with IBM and our User Group members to improve Spectrum Scale and Spectrum Scale with Spectrum Protect ILM
- Provide a "localised" community for support and knowledge sharing
- Raise awareness of Spectrum Scale as a capable data management platform I can provide a more detailed description if needed.

# Redefining Data

## Centres and Hosting

Scotland has all the attributes to become one of Europe's emerging data centre hubs and **DataVita offers organisations the opportunity to optimise their data centre strategies for the future.**

### Why host with DataVita in Scotland?

#### Hosting Capacity

We are **capable of hosting 130kW + per rack**, and can host the latest generation of servers and GPUs from all the hardware vendors.

#### Cost

DataVita is **typically more cost effective than any HPC capable data centre**. We also offer a remote hands service at no additional cost, so you never need to come to site.

#### Low Carbon

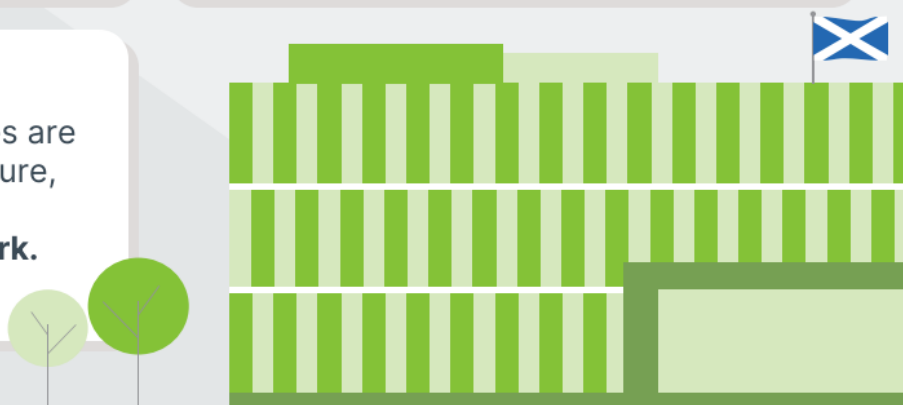
We can **reduce your carbon emissions by up to 11 times**, Scotland has the lowest carbon intensity in the UK.

#### Cooling

We offer comprehensive cooling solutions to meet all customer needs. Air cooled and **direct to chip liquid cooling** available.

#### Connectivity

Connectivity is key. All our services are designed to offer you reliable, secure, and high-speed solutions, such as **connectivity to the JANET network**.



[Visit our Website](#)

[Follow us on LinkedIn](#)

## ExCALIBUR Hardware and enabling software: Exploring next generation technologies for HPC

Friday 6 December - 11:00-13:00 - Exchange Room 1

There are a wealth of new hardware technologies being made available for HPC. Driven by the boom in AI, vendors promise that these can deliver improved performance and significant energy savings, but a key question is how closely this matches reality and how realistic these are for HPC workloads. This is the challenge that the ExCALIBUR H&ES programme has been looking to explore since 2019 by setting up a range of novel hardware testbeds across the UK. H&ES has provided not just the hardware itself, but also addressed the observation that a key challenge in leveraging such technologies is in the software ecosystem. In this session we will reflect on the successes of the H&ES programme, describe some of the key testbeds that audience members can gain access to, and explore future opportunities to build upon the outcomes of H&ES with the CIUK community.

## UKRI National Federated Compute Services NetworkPlus

Friday 6 December - 14:00-16:00 - CIUK Breakout Room

Introduction to the community of the recently DRI funded National Federated Compute Services NetworkPlus. The project team is led by Jon Hays (QMUL) with co-leads Sadaf Alam (University of Bristol), Adrian Hines (JASMIN, STFC), Josephine Beech-Brandt (EPCC, University of Edinburgh) and Vasil Alexandrov (Hartree, STFC). We would like an opportunity to engage with the community and present the timeline for the workshops and calls which are being planned and discuss flexible funding for projects. We can provide more details if required.



## CIUK 2024 RESEARCH ZONE



The [CIUK 2024 Research Zone](#) will allow all EPSRC Tier-2 Centres, and other sites with significant size computing facilities, the opportunity to join the CIUK exhibition and update attendees on their progress. You can visit them in the exhibition hall to find out about their systems, discuss potential projects and discover how to get access.

**DiRAC**

**JASMIN**

**UKRI** Science and Technology Facilities Council  
Scientific Computing

**archer2**

**hpc**  
MIDLANDS+

**B**

**GW4**  
Isambard

**CSD<sup>3</sup>**  
Cambridge Service for Data Driven Discovery

**cirrus**  
Powered by epcc

**MATERIALS AND MOLECULAR MODELLING HUB**

**NI-HPC**  
NORTHERN IRELAND HIGH-PERFORMANCE COMPUTING

**EXCALIBUR 10**

**N8 CIR**  
Computationally Intensive Research  
**BEDE**

**UKRI** Science and Technology Facilities Council  
Hartree Centre

**HPC SIG**  
The High Performance Computing Special Interest Group for UK Academia

**UK Atomic Energy Authority**

**iris**

## CIUK 2024 STUDENT CLUSTER CHALLENGE

Following the success of our Student Cluster Challenges in 2023, the competition will return in 2024 for its fifth and biggest edition. As in 2023 this years' [Student Cluster Competition](#) will be a hybrid competition with a series of online challenges leading up to the conference followed by some in-person challenges during the conference itself.

Meet the teams...

### Team: (u)CLI University College London



Rozenn Raffaut

Yuliang Huang

Tom Bickley



Zak Morgan

Shahid Khalid

Qi Li

### Team: Brunels University of Bristol



Vladimir Dementyev

Lucas Oliveri

Hanzhong Qiu



Dhillon Thurairatnam

Daniel Adigamov

Thomas Bale

### Team: Cardiff Dragons Cardiff University



David Dan

Dean Jukes

Arnav Jain



Jiatong Zhu

Ronan Osborne

### Team: ClusDur Durham University



Charles Dubois-Veltman

Elizabeth Holland

Keiran Barnes



Sam Ezeh

Edward Jackson

Jeremy Mariani

**Team: clusterphobia**  
Manchester Metropolitan University  
and University of Liverpool



Thankgod Egbe



Richard Omole



Majid Kamali



Ayomide Popoola



Tawakalit Agboola



Ali Naeimi

**Team: CompuDur**  
Durham University



Nam Le



Michal Pluta



Oliver Turner



Toby Davis



Jane Manning



Matt Richardson

**Team: Decarb**  
Warwick University and  
Imperial College London



William Jones



Vinz Kakilala



Alex Do



Haaris Farooq



Viraj Ganatra



Aalaap Dey

**Team: Imperial Bytes v2**  
Imperial College London



Harvey Gould



Giacomo Rossetti



Stephanie Wong



Freddie Nicholson



Bella Hurrell



Rose Newkey-Burke


**Team: Isambards**  
University of Bristol



Lawrence Douglas



Oliver Marks



Aidan Ewart



Aleksandra Walusiak



Irek Swierczynski



Ajinkya Bhalerao

**Team: Most Significant Bit**  
The Rosalind Franklin Institute



Chloe Cheng



Victoria Hann



Gabryel  
Mason-Williams



Natalie Nord



Aparajita Karmakar



Melissa Mitchell

Team: PlyCompSoc  
University of Plymouth



Gabriel Marangoci



Joseph Troughton



Mohammad Khazaei



Nikodem Drabik



Thomas Cogzell

Team: return -1  
University of Birmingham




Tahashin Muaz



Lydia Barucka



Ayush Acharjya



Fearn Townsend



Ivan Korniiichuk




Vivan Waghela

Team: TED  
University of Birmingham



Aaron Abraham



Shay Patel



Siddarth Shringarpure



Adwit Mukherji



Daniel Ebbutt



Victor Andraous

Team: TODO: Refactor Later  
King's College London



Alex Rubio Blobaum



Mark Chester  
Jude Emmanuel



Adam Moujar Bakhti



Vaibhvkumar Patel



Harshit Sharma



Chenhao Liu

Team: YorKluster  
University of York



Luke Arnold




Indigo Edwards



Kit Brown



Aoife Oliver



Juniper Dunn



Emma Fare

You can find the Cluster Challenge teams at the entrance to the exhibition hall, by the information point. Please feel free to pay them a visit as they complete their challenges during the conference.

You can also follow the competition on Twitter [#CIUK2024\\_SCC](#).

The winning team will be announced before the final presentation session on Friday 6 December with the champions going forward to represent CIUK at the ISC'24 Cluster Challenge in Germany next summer.

As always, we could not host the cluster challenge without the support of our cluster challenge partners and we thank them for their time, efforts and the opportunity to access their systems...



SCAN<sup>®</sup>

# Your **Partner** for End-to-End Technology Solutions

## Meet the Scan team at CIUK

- ✓ Discover how we can enhance your HPC environment
- ✓ Explore solutions designed to fit your unique goals
- ✓ Connect with our experts to discuss your enterprise needs
- ✓ Hear more about our own cloud from individual cards, through to GPU compute clusters of DGX B200's, GH200 and HPC infrastructure



Scan are proud to exhibit alongside some of our leading partners NVIDIA, PNY, Peak:AIO, DDN, AMD, Intel, and supported by Crunchyard and SUSE.




Together, we provide the technologies and expertise needed to ensure your success.

TECH  
HEART  
& SOUL

## CIUK 2024 POSTER COMPETITION

10 finalists have been chosen for the [CIUK 2024 Student Poster Competition](#). Their posters can be viewed on the exhibition floor and also [online](#). The winner will be announced prior to the final presentation session on Friday 8 December.

<p><b>Kiran Jonathan</b> <i>UKRI – STFC</i></p> 	<p><b>Accelerating Research with HPC: A Collaborative Study of Space Charge Compensation in H- Ion Sources</b></p> <p>Particle accelerators such as the ISIS Neutron and Muon Source (ISIS) have been crucial in probing our understanding of fundamental physics and have a wide range of applications across academia and industry. One challenge faced in the design of such accelerators is the space charge effects exhibited in the low-energy beam transport (LEBT) region, where charge repulsion can cause a defocussing of the beam, reducing the overall transmission and energy efficiency. These effects can be reduced via a process known as space charge compensation (SCC), through the addition of factors such as a neutral background gas. This project is a collaborative effort between the Scientific Computing Department (SCD) and ISIS, which aims to use a combination of real-world beam diagnostics and high-fidelity particle-in-cell simulations to develop a deeper understanding of the SCC process and its key parameters, ultimately informing design decisions for current and future iterations of the ISIS H- ion source</p>
<p><b>Gokmen Kilic</b> <i>Durham University</i></p>	<p><b>Visualising HPC Node Performance via Swift ReFrame Benchmark and Grafana</b></p> <p>High-Performance Computing systems are fundamental enablers of complex computational workloads across scientific, engineering, and industrial domains. Accurately measuring and visualizing the performance of HPC nodes is analytical for optimizing resource allocation, identifying performance bottlenecks, and improving overall system efficiency. This work presents an approach to monitor and visualize HPC node performance using the Swift ReFrame Benchmark framework in conjunction with the Grafana data visualization platform. The ReFrame Benchmarking Framework is a highly flexible and portable tool designed for writing regression tests for HPC systems. It focuses on scalability, ease of use, and seamless integration with various HPC environments. This work demonstrates ReFrame's capabilities to design customizable performance benchmarks that gather key metrics such as CPU utilization, memory bandwidth, and network latency across multiple HPC nodes. The benchmark results produced by ReFrame</p>

**Jinjiang Li**  
*University of Manchester*



**MetaWAAM: Real-Time Digital Twin Architecture for Direct Energy Deposition-Arc Manufacturing**

Digital twins (DT) are increasingly used in manufacturing for real-time monitoring, process optimization, and fault prediction, but challenges remain in achieving seamless physical-virtual integration and real-time system optimization. This study leverages NVIDIA Omniverse™ to develop a DT architecture for a Direct Energy Deposition-Arc (DED-Arc) manufacturing unit, enabling real-time monitoring of robotic arm movements and welding states. The architecture integrates edge computing and finite element analysis (FEA) of laser-scanned 3D models, allowing on-site data processing and efficient transmission to high-performance computing (HPC) systems for rapid simulation. This scalable DT system provides immediate feedback to the physical manufacturing process, supporting high-quality additive manufacturing. As the first study to achieve real-time DT monitoring in DED-Arc manufacturing, it offers a foundation for future advancements in simulation and process optimization.

**Terence Lobo**  
*UKRI - STFC*



**Accelerating High-Order Finite Difference Methods for Turbulence and Combustion Using CUDA Fortran**

High-performance computing (HPC) has become a cornerstone in advancing scientific and engineering research, particularly in fields requiring intensive numerical computations. This project focuses on accelerating an existing CPU-only high-order finite difference (FDM) solver for turbulence and combustion simulations by leveraging the computational power of Fortran and GPU programming.

Fortran is widely used in the scientific community due to its efficiency in handling numerical tasks. However, the increasing complexity and scale of simulations necessitate further optimisation to meet the demands of modern research. This project explores the integration of NVIDIA's HPC SDK (CUDA) with Fortran to enhance computational performance. By utilising CUDA Fortran, this project aims to parallelise all the computational tasks, thereby significantly accelerating numerical computations.


The primary objective is to demonstrate the feasibility and effectiveness of CUDA Fortran in parallelising tasks to achieve significant acceleration in numerical computations. Key results include a detailed comparison of performance metrics between traditional CPU-based computations and the optimised GPU-accelerated approach, using a single A100 GPU for comparison. The findings highlight the potential of CUDA Fortran to transform HPC applications, achieving up to 20x acceleration, which includes IO operations, making it a viable solution for large-scale simulations in turbulence and combustion research.

This poster will present the methodology, implementation details, and performance outcomes, providing insights into the practical benefits and challenges of using CUDA to adopt GPU programming for a high-order compressible flow solver. The results underscore the importance of leveraging modern HPC tools to push the boundaries of computational research



**Miao Zeyuan**

*University of Manchester*



**Thermal Modelling in Neutronics Using Physics-Informed Neural Networks: Leveraging HPC for Surrogate Model Training**

Conducting thermal analyses in neutronic heating environments, such as nuclear reactors, is essential for component design and optimization. However, traditional finite element methods (FEM) have limitations, particularly in their inability to transfer information between similar analyses, which leads to redundant computations during design exploration. Additionally, FEM requires complex meshing, especially for intricate geometries, with each simulation often requiring a unique mesh design. This need for distinct meshes complicates data mapping between simulations, reducing both accuracy and computational efficiency.

This study explores the application of Physics-Informed Neural Networks (PINNs) as a mesh-free, efficient alternative for thermal modeling in neutronics. Unlike FEM, PINNs do not require mesh generation and allow for knowledge transfer across simulations by encoding learned information within network weights, enabling more efficient design exploration through transfer learning. Building a PINN surrogate model benefits from utilizing substantial computational resources, as network training can be easily scaled using parallel computing on GPUs. In this work, we leverage the powerful computational capabilities of the Isambard HPC center to accelerate the training process.

**Aadya Mudgil**

*University of Warwick*



**On Optimal Resource Allocation Algorithms**

Highly parallelisable HPC jobs are creating unprecedented demands on compute resources. Heuristic-based algorithms are widely used to distribute servers among such jobs but lack theoretical guarantees and often underperform. This research presents a simple, non-pre-emptive allocation algorithm with provable performance guarantees, optimally addressing the trade-off between server allocation and job speed-up. Unlike prior methods, our approach eliminates the need for advance job characteristic knowledge while achieving the same optimal performance. Extensive simulations validate its effectiveness across diverse traffic conditions and heterogeneous workloads, mirroring real-world scenarios. This work provides theoretical insights and practical solutions for resource allocation in HPC/Cloud computing systems.




**Maitrayee Singh**

*UKRI – STFC*



**MD-based spectroscopic simulations in Py-ChemShell**

Spectroscopic technologies, such as UV-vis and IR, are essential in molecular and materials science for uncovering molecular structures and interactions. However, interpreting experimental spectra accurately requires simulations that account for molecular dynamics. This poster introduces a new workflow in Py-ChemShell for MD-based spectroscopic simulations, designed to enhance spectral analysis by integrating molecular conformational fluctuations. Through the addition of a "Spectroscopy" task class, this workflow supports UV-vis spectral simulations on molecular snapshots from MD trajectories, offering a more realistic depiction of molecular behaviour over time. Key features include data averaging and convolution techniques for refining spectral

	<p>outputs. Preliminary benchmarking shows that these simulations improve alignment with experimental spectra, making Py-ChemShell a more powerful tool for researchers.</p>
<p><b>Raska Soemantoro</b> <i>The University of Manchester</i></p> 	<p><b>A case study of real-time collaborative design in FreeCAD and NVIDIA Omniverse</b></p> <p>FreeCAD, a robust open-source CAD tool, is popular among small organizations and users seeking an efficient, bloat-free design platform, with many applications integrating it as a CAD backend. However, it lacks collaborative design features essential in modern engineering. This work introduces a FreeCAD connector for NVIDIA Omniverse, enabling real-time collaboration, version control, and photo-realistic rendering. Demonstrated through a case study of engineers designing a sensor for a stellarator fusion plant, the connector enhances FreeCAD workflows by connecting to Omniverse's Nucleus—a secure, central storage platform compatible with third-party tools. This integration offers FreeCAD users a collaborative, non-proprietary CAD solution, expanding design capabilities for engineers and creators within Omniverse.</p>
<p><b>Yazhmozhi Vasuki Murugesan</b> <i>University of Dundee</i></p> 	<p><b>Tamil Spell Checker</b></p> <p>The poster showcases the development of a spell checker for Tamil. This research examines how different language models such as LSTM, BiLSTM, mBERT, mT5, XLM-R and Llama perform with spelling error detection and correction of a highly inflectional language like Tamil. A systematic literature review of spell checkers for the major written Dravidian languages – Tamil, Telugu, Kannada, and Malayalam - led to the development of two Tamil language resources – "TamilCorp," a balanced corpus of Tamil written texts containing approximately 1.7 billion tokens across 16 genres and TamilSpell, first-of-a-kind Tamil spelling error corpus with ~10 million entries to support the spell checker. The TamilContextSpell, an open-source toolkit that features both a command-line interface and a graphical user interface, will be published along with the performance metrics of all models tested. Dictionary-based hashing techniques and cross-lingual language models will be used to handle non-word errors.</p>
<p><b>Jools Wills</b> <i>UKRI – STFC</i></p> 	<p><b>Linking data sources with CCP4 Cloud: A fully online setup for structure determination</b></p> <p>Over the last decade, Macromolecular Crystallography (MX) has seen a prominent shift towards online solutions in both experimental and computational parts of the structure determination process. An online approach may start with sending crystals to a synchrotron, followed by a remote diffraction experiment with subsequent data processing and structure solution using a suitable online software setup, and finishing with the PDB deposition. We have developed DataLink, a file/database storage service with a web based API capable of acquiring raw image data from various sources, including on-line x-ray diffraction image repositories and potentially any MX data source such as a synchrotron or a diffractometer. We also now have incorporated Globus into CCP4 Cloud allowing users to send data from any facility that has a Globus Endpoint</p>



## We want your feedback...

Help us to shape next year's event by providing your thoughts and ideas. We welcome all feedback – good or bad – and we are specifically interested in ideas for the main theme of the conference in 2024.

See you in 2025...

The banner features a blue background with a decorative pattern of yellow and white diagonal lines. On the right side, there is a large, colorful graphic consisting of overlapping, curved lines in shades of blue, purple, red, and orange. A prominent green circle is overlaid on this graphic, containing the text 'SAVE THE DATE'. The UKRI logo is positioned in the top right corner, with the text 'Science and Technology Facilities Council' to its right. The main text is centered on the left side of the banner.

**COMPUTING  
INSIGHT UK 2025**

**4-5 December 2025**  
**Manchester Central, UK**

**SEE YOU IN 2025**  
[www.ukri.org/CIUK](http://www.ukri.org/CIUK)

**UKRI** Science and  
Technology  
Facilities Council

**SAVE  
THE  
DATE**