Science and Technology Facilities Council

Scientific Computing

COMPUTING INSIGHT UK 2025

Computing Unites

4 – 5 DECEMBER 2025

Manchester Central, UK

www.ukri.org/CIUK

CONTENTS

INTRODUCTION

We are delighted to welcome you all to Computing Insight UK 2025.

The theme for this year's conference is "**Computing Unites",** with sub-themes including Sustainable Energy Efficient Computing, UKRI Digital Research Infrastructure, Scientific and Technical Community, Power of Computing and Industry Futures.

CIUK 2025 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 firth 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 3 and Thursday 4 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 Albert Hall, 27 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.

You can also follow us on 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 2025 Media Partner

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

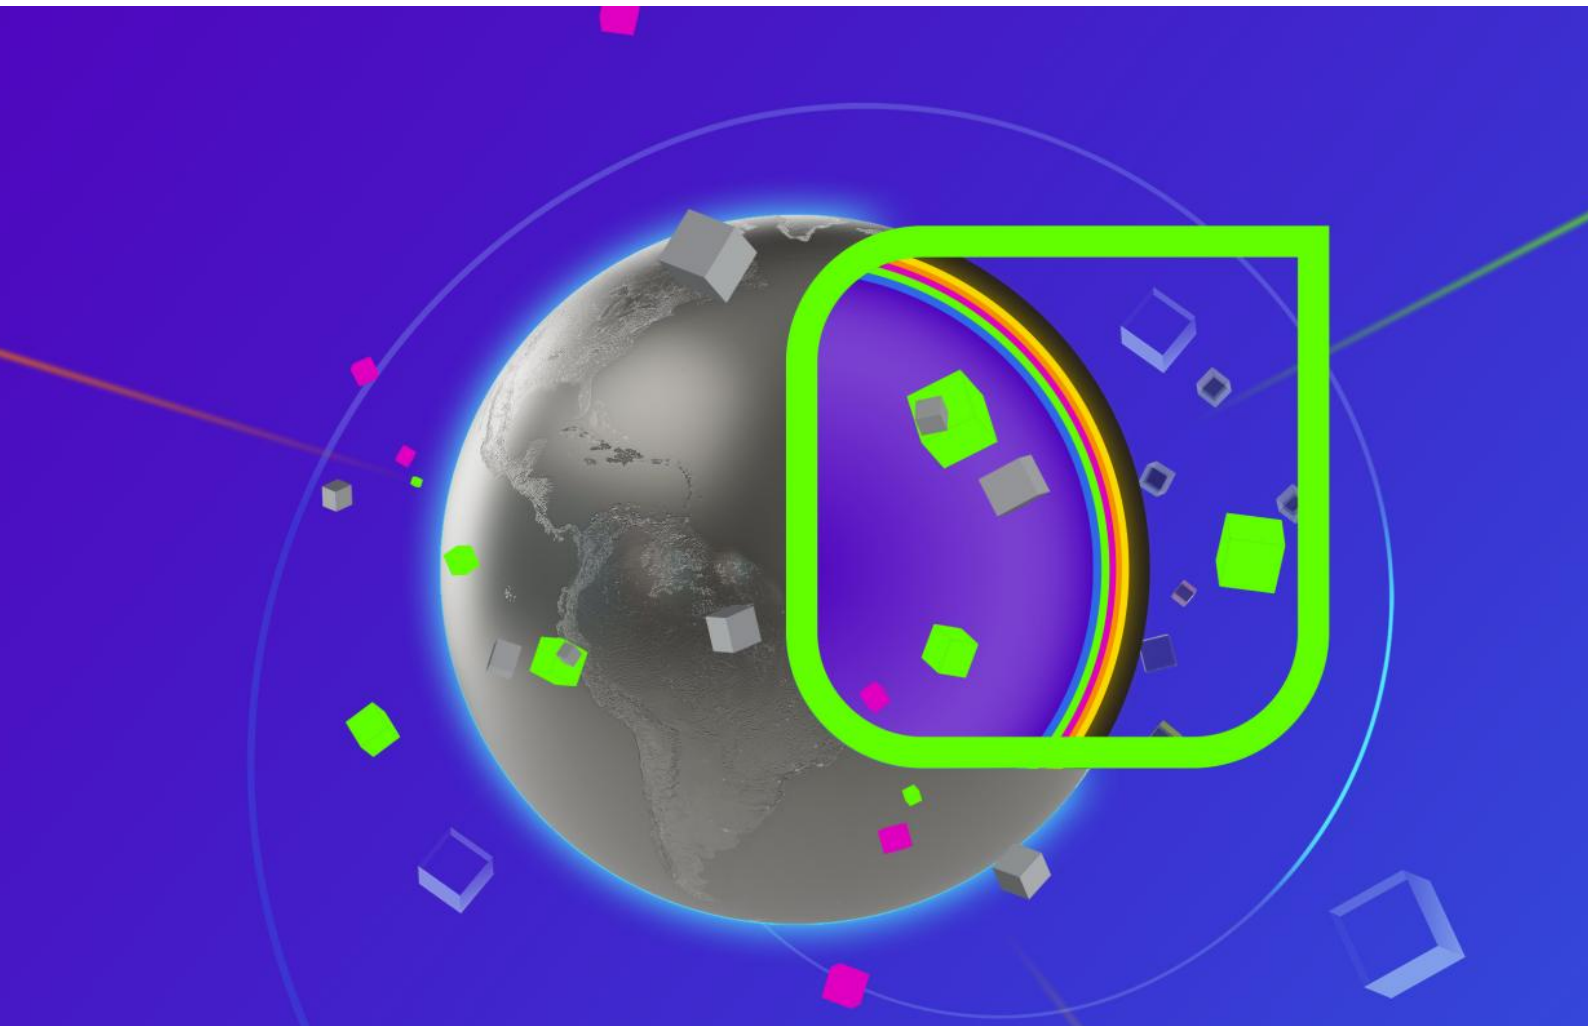

INDUSTRIAL CLOUD SOLUTIONS FOR AI AND HPC PRODUCTION

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

viridiengroup.com/hpc

SEE THINGS DIFFERENTLY

VIRIDIEN

CIUK 2025 PROGRAMME

The main CIUK 2025 programme will take place in the Auditorium. Look for the BLUE pull-up banners.

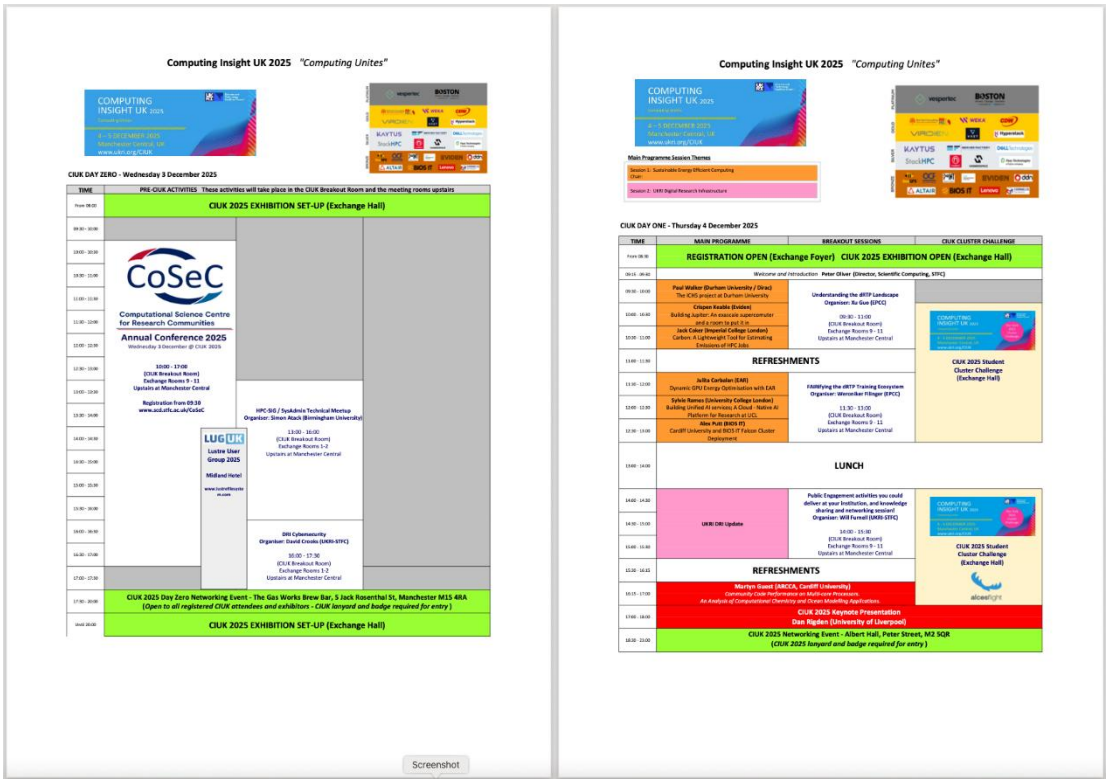

Please click the image above to open the CIUK 2025 Programme

CIUK 2025 KEYNOTE PRESENTATION

Prof Dan Rigden

University of Liverpool

The impact of AlphaFold on experimental and computational biology research (and researchers).

Abstract: AlphaFold is a ground-breaking deep learning-based method for prediction of the structure of proteins and biomolecular systems. Its development exemplifies how the patient acquisition of high-quality data, in this case archived by the PDB, can enable disruptive scientific innovation. The landscape of protein structure research, both experimental and computational, changed rapidly with the arrival of AlphaFold; some approaches and workflows have become largely obsolete, while other areas have seen more intense efforts and rapid progress, often using methods inspired by AlphaFold. This talk will try to present an overview and comparison of biology pre- and post-Alphafold, with an emphasis on the role of UK research communities and scientific computation.

Bio: Dan trained as a Biochemist in Oxford then dabbled in protein structure modelling in its infancy during a PhD in Edinburgh. Time in Leeds and Brasilia was spent doing protein X-ray crystallography alongside bioinformatics but increasingly gravitating towards the latter. He returned to the UK in 2003 to take up a position as Lecturer in Bioinformatics at the University of Liverpool and has been there ever since. Much of that time has been spent considering how bioinformatics innovations can be translated to benefit communities of experimental structural biologists, especially protein crystallographers through ties with CCP4. Latterly this has entailed adapting to the many opportunities and occasional challenges of the post-AlphaFold world.

CIUK 2025 JACKY PALLAS MEMORIAL AWARD PRESENTATION

Flaviano Della Pia

Yusuf Hamied Department of Chemistry - University of Cambridge

Towards modelling molecular crystals with the accuracy of diffusion Monte Carlo.

Abstract: Computational modelling plays a central role in molecular crystal discovery, fundamental to a wide range of applications, including pharmaceuticals and renewable energy. However, a reliable description of these systems requires both a high-accuracy description of the potential energy surface and a fully anharmonic quantum description of the nuclear motion. [1,2] In this talk, I will first show that reliable lattice energies can be obtained with quantum diffusion Monte Carlo. [3,4] Then, I will demonstrate the generation of machine learning interatomic potentials capable of describing molecular crystals at finite temperature and pressure with sub-chemical accuracy, using as few as ~ 200 data structures, an order of magnitude improvement over the current state-of-the-art. [5] Ourmodels successfully reproduce experiments for a diverse range of molecular crystals and open up the prospects of reliable modeling for drug discovery and beyond.

[1] G. J. O. Beran, Chem. Rev. 2016, 116, 9, 5567–5613 (2016)

[2] V. Kapil and E. A. Engel, Proc. Natl. Acad. Sci. U.S.A. 119 (6) e2111769119 (2022)

[3] F. Della Pia, A. Zen, D. Alfè, A. Michaelides, J. Chem. Phys. 157, 134701 (2022)

[4] F. Della Pia, A. Zen, D. Alfè, A. Michaelides, Phys. Rev. Lett. 133, 046401 (2024)

[5] F. Della Pia, B. X. Shi, V. Kapil, A. Zen, D. Alfè, A. Michaelides,Chem. Sci., 16, 11419-11433 (2025)

Bio: Flaviano Della Pia is from Naples, Italy, where he graduated in Condensed Matter Physics. He recently completed his PhD in Computational Chemistry at the University of Cambridge, where his work explored theoretical and computational approaches to understanding materials and molecular systems.

Accelerate Discovery: Transform Unstructured Data Management for Faster, More Efficient Research.

Visit us at Booth 25/26 to discover how Quantum, alongside datacentre infrastructure specialists Vespertec, can empower your research and IT infrastructure. From optimising unstructured data management to accelerating data centre 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.

For information following CIUK, email info@vespertec.com or info.uk@quantum.com

vespertec Quantum

CIUK 2025 Speakers

Abstract: The Immersion Cooling and Heat Storage (ICHS) project explores making waste HPC heat more useful through two approaches. First, we installed a single-phase immersion tank operating the COSMA5 cluster, gaining hands-on experience with technology offering lower embodied CO2 and PUE than conventional cooling. We welcome RTPs to visit our facility and offer travel funding. Second, we're investigating flooded abandoned mine workings beneath the University for inter-seasonal heat storage. Waste heat transfers into a lower seam while cooler water extracts from a higher seam, creating a thermal battery. Flow reverses during winter when heat is needed. We report facility experiences, modeling progress, and borehole drilling tender status.

Bio: Dr Paul Walker had a background in commercial and public sector IT before moving to HPC, and works on the DiRAC Memory Intensive HPC service, COSMA, at Durham University.

Abstract: This talk will peel back the covers on how Eviden built the JUPITER system. It will look at some of the problems we faced, and tell the story of what we achieved. From a standing start with no computer or data centre at the beginning of 2025, we now have a DC, the first exascale computer in Europe, a major HPC and AI power house already starting to produce science, and a leader on the top500. It's also a system which leads the world in its class on energy efficiency.

Bio: Crispin Keable, PhD - Senior solution architect, Eviden HPC & AI. I am a system architect for global HPC and AI solutions mapping customer needs to systems design, considering all aspects of technology choice and directions. In this role I was architect for the EuroHPC JUPITER system. I have had a long career in HPC, working for Cray Research, SGI, IBM, and Bull

Abstract: High-performance computing (HPC) plays a critical role in modern research, but it also consumes large amounts of energy and contributes to carbon emissions. As institutions seek to reduce their environmental footprint, there is a growing need for simple, transparent ways to make the environmental cost of computing visible to researchers and resource managers. This talk presents a lightweight tool for estimating the carbon emissions of HPC jobs. It requires only job resource usage from the scheduler and a minimal cluster configuration describing hardware power characteristics. From these inputs, energy consumption is estimated and converted to carbon emissions using region-specific intensity data. It does not depend on node-level monitoring or vendor-specific tools, making it straightforward to deploy across HPC systems. This non-intrusive tool helps researchers and support teams understand the impact of computing, encouraging more sustainable practices in computational research.

Bio: Jack is a computational physicist and research software engineer. Following a PhD researching semiconductor physics and its applications for renewable energy technology, his focus is now on understanding and mitigating the environmental impact of computing. He is currently employed by Imperial's Research Computing Service, on their HPC and RSE experience programme. Jack is funded by the ICICLE collaboration, a partnership between Imperial, Intel, and Lenovo, established to promote sustainable scientific computing.

Abstract: The sharp rise in GPU power consumption—up to 1100W per unit in the latest NVIDIA architectures—has made energy efficiency a critical challenge for HPC and AI data centres. Despite extensive research, no standard solution has been widely adopted due to hardware dependencies and power–performance trade-offs. Energy Aware Runtime (EAR) software system addresses this by providing transparent, vendor-agnostic monitoring and optimisation of workloads. EAR dynamically analyses applications at runtime, tuning CPU and GPU frequencies automatically while supporting system-wide power and thermal management. This paper presents the EAR architecture with a focus on GPU energy optimisation and shows results on real HPC and AI workloads. EAR is already deployed in major European supercomputing centres, proving its scalability and production readiness.

Bio: Raphaël Brochard is an entrepreneur. An entrepreneur with a passion for innovation and sustainability. Over the years, he has co-founded a couple startups across various industries, always driven by a commitment to impactful, forward-thinking solutions.

Raised in Paris, France, Raphaël moved to the United States at 18, where he earned a Bachelor of Science from Texas State University.

He began his career in the tech industry at Apple, working on Geographic Information Systems (GIS) for public space mapping. He then transitioned to Indeed, as an International Operations Analyst. After a decade in the U.S., he returned to France to launch his first company.

Raphaël's curiosity led him to start a new entrepreneurship adventure…

He joined forces with his father, Luigi Brochard, and colleague Julita Corbalan Gonzalez, to develop EAR: Energy Aware Runtime. Raphaël is now the General manager of EAR and focusing on growth.

We believe in sustainable data centers, and efficient HPC/AI computing. Gain +20% efficacy per watt. More performance per watt. How? With the new software standard for CPU and GPU: Energy Aware Runtime. EAR is a system software to monitor/analyse and optimise energy in data centers.

Raphaël continues to pave the way toward a more sustainable and efficient technological landscape.

Abstract: The Unified AI Services platform at UCL complements existing HPC services by providing a cloud-native, GPU-powered environment for AI/ML and Data Science. Built from scratch using Kubernetes on-premises, it integrates CNCF open-source tools and vendor solutions like the NVIDIA GPU Operator for performance. This talk shares our journey from initial setup to pilot users, highlighting lessons learned in deploying Kubeflow—a Kubernetes-based AI platform supporting the full AI lifecycle: data preparation, model development, training, hyperparameter tuning, and model serving—from both admin and user perspectives.

Bio: Sylvie Ramos is Principal Research Infrastructure Developer at the Advanced Research Computing Centre, where she leads the Unified AI Services team. She holds a Ph.D. in Physics Engineering and a master's in electrical and computer engineering, with academic training at former Pierre et Marie Curie University now part of the Sorbonne Université, University of Coimbra, and Instituto Superior Técnico. Her research background includes microwave diagnostics and plasma instability studies in nuclear fusion devices at the ASDEX Upgrade tokamak (Max Planck Institute for Plasma Physics, Germany) and other major European facilities such as JET (Culham, UK) and TJ-II (Madrid, Spain). Transitioning from experimental research to technical leadership, she has contributed extensively to large-scale data systems and research software development at Diamond Light Source and the Rosalind Franklin Institute. Her current focus is on AI infrastructure and advanced computational services for scientific research.

BOSTON
Servers | Storage | Solutions
A DIVISION OF SOURCECODE

NVIDIA Elite Partner

VISIT US AT BOOTH 7

TO RECEIVE 50% OFF NVIDIA DLI TRAINING

Get practical experience with deep learning through Boston Limited's certified NVIDIA Deep Learning Institute (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

Abstract: We would like the opportunity to present on the new Falcon Cluster deployment at Cardiff University. This deployment has consisted of a partnership of; The Arcca team at Cardiff University, BIOS IT, Supermicro, DDN and Clustervision. This project marks a new chapter for Cardiff Uni, utilising Supermicros Direct to Chip Liquid cooling to allow for large core count nodes resulting in a dense solution totalling 7,296 cores. Along with this we have provided H100 and L40s systems. We have also utilised Clustervision and their Trinity X cluster management software. This software offers robust 24/7 management and oversight for HPC&AI environments. Designed to be both dependable and SLA-compliant, this software enables you to concentrate fully on your research by monitoring and managing complex technologies. These include Linux, SLURM, CUDA, InfiniBand, Luster, and Open OnDemand, ensuring that your system runs smoothly and efficiently.

Bio: Religious Studies and Business studies graduate from Canterbury Christchurch Uni. Been at BIOS IT for 7 years with a background in Telco and SAAS.

Abstract: This session will provide an update on progress in developing UKRI's Digital Research Infrastructure (DRI) programme and delivery of the UK's Compute Roadmap. There will be opportunity for questions and discussion.

Bio: Richard Gunn is Programme Director for Digital Research Infrastructure at UK Research and Innovation (UKRI), where he leads on support for the ecosystem of large-scale compute and data infrastructures, software, networks, tooling and skilled DRI professionals that enables the UK's researchers and innovators.

With over a decade of experience at UKRI and EPSRC, Richard has held a range of senior roles focused on emerging technologies and research infrastructure. Prior to his time at UKRI, Richard held roles in intellectual property and public policy and completed a PhD in chemical biology.

Abstract: This session will provide an update on progress in developing UKRI's Digital Research Infrastructure (DRI) programme and delivery of the UK's Compute Roadmap. There will be opportunity for questions and discussion.

Bio: Tom is the AI Infrastructure Policy Lead at the Department for Science, Innovation and Technology (DSIT), where he works on the AI Research Resource (AIRR) programme. His work focuses on developing the UK's public network of supercomputers to support cutting-edge AI research in priority fields including healthcare, climate science, and public safety. Before joining the Civil Service, Tom spent a decade working on corporate responsibility and environmental, social, and governance (ESG) issues. He holds an MSc in Data and AI Ethics.

Abstract: In a deviation from our usual format, this presentation looks to demonstrate the performance of a number of widely used GPU-enabled Molecular Dynamics (MD) community codes important in understanding biological systems and drug interactions. This field was an early and significant adopter of GPU acceleration, with the codes featured here – AMBER, GROMACS, LAMMPS and NAMD – widely used on both local and regional clusters.

Our objective is to demonstrate the performance improvements of the selected codes on GPUs compared, for example, to those seen on CPU systems, understanding the impact on performance related to the availability of various GPU features. Attention is focused here on systems hosting NVIDIA's V100, A100, L40S and H100 GPUs

Note that this work was in part undertaken under the auspices of the UKRI funded SHAREing project ("Skills Hub for Accelerated Research Environments Inspiring the Next Generation").

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-3 system at the GW4 Tier-2 HPC regional centre.

Martyn’s research interests cover the development and application of computational chemistry methods. He is lead author of the GAMESS-UK electronic structure program and has written or contributed to more than 260 journal articles.

Abstract: Limited standardization across HPC clusters creates challenges for administrators who must develop site-specific management solutions, for Research Software Engineers who waste time adapting and re-testing software, and for users who face a steep learning curve with each new environment. This talk introduces Charmed HPC, a new open source project from the Ubuntu community that addresses these issues by simplifying and standardizing the setup and management of HPC clusters. Charmed HPC is not "just another deployment tool", but rather a flexible toolbox for both production and test environments, enabling administrators and RSEs to easily manage complex tasks like scaling partitions or enabling the Slurm workload manager’s high-availability mode. This talk will also cover the open source governance model of Charmed HPC, highlighting how community and enterprise contributors are collaborating to add core features and will outline opportunities for new contributors to get involved with the project.

Bio: Jason is a young professional in the HPC industry passionate about developing the next generation of lean, mean, open source supercomputing machines. By day, he works at Canonical as one of their resident HPC engineers making Ubuntu better for supercomputing, and by night he leads the Ubuntu HPC community team as one of its “Not so Ancient Elders.” He is focused on addressing current and future challenges facing the HPC industry such as the convergence of cloud and HPC systems, supporting new heterogeneous architectures, and enabling advanced research software engineer workflows. In his free time, he likes to work on several of his open source projects, organise open source community conferences such as the Ubuntu Summit, and travel to new places to teach others about how great HPC and supercomputing is. Recently he’s started learning the Crystal programming language out of personal interest for its potential applications inside HPC environments. Why Crystal specifically? Well… after five years, he still hasn’t decided if he wants to write Rust or not.

Dominic is a Software Engineer at Canonical, where he contributes to the High-Performance Computing (HPC) team's mission of integrating HPC into the broader Ubuntu platform. He is currently focused on open source development of "charmed" Python operators for key HPC software to automate the lifecycle of clusters. His background includes previous roles in supporting HPC resources at multiple academic sites within the UK.

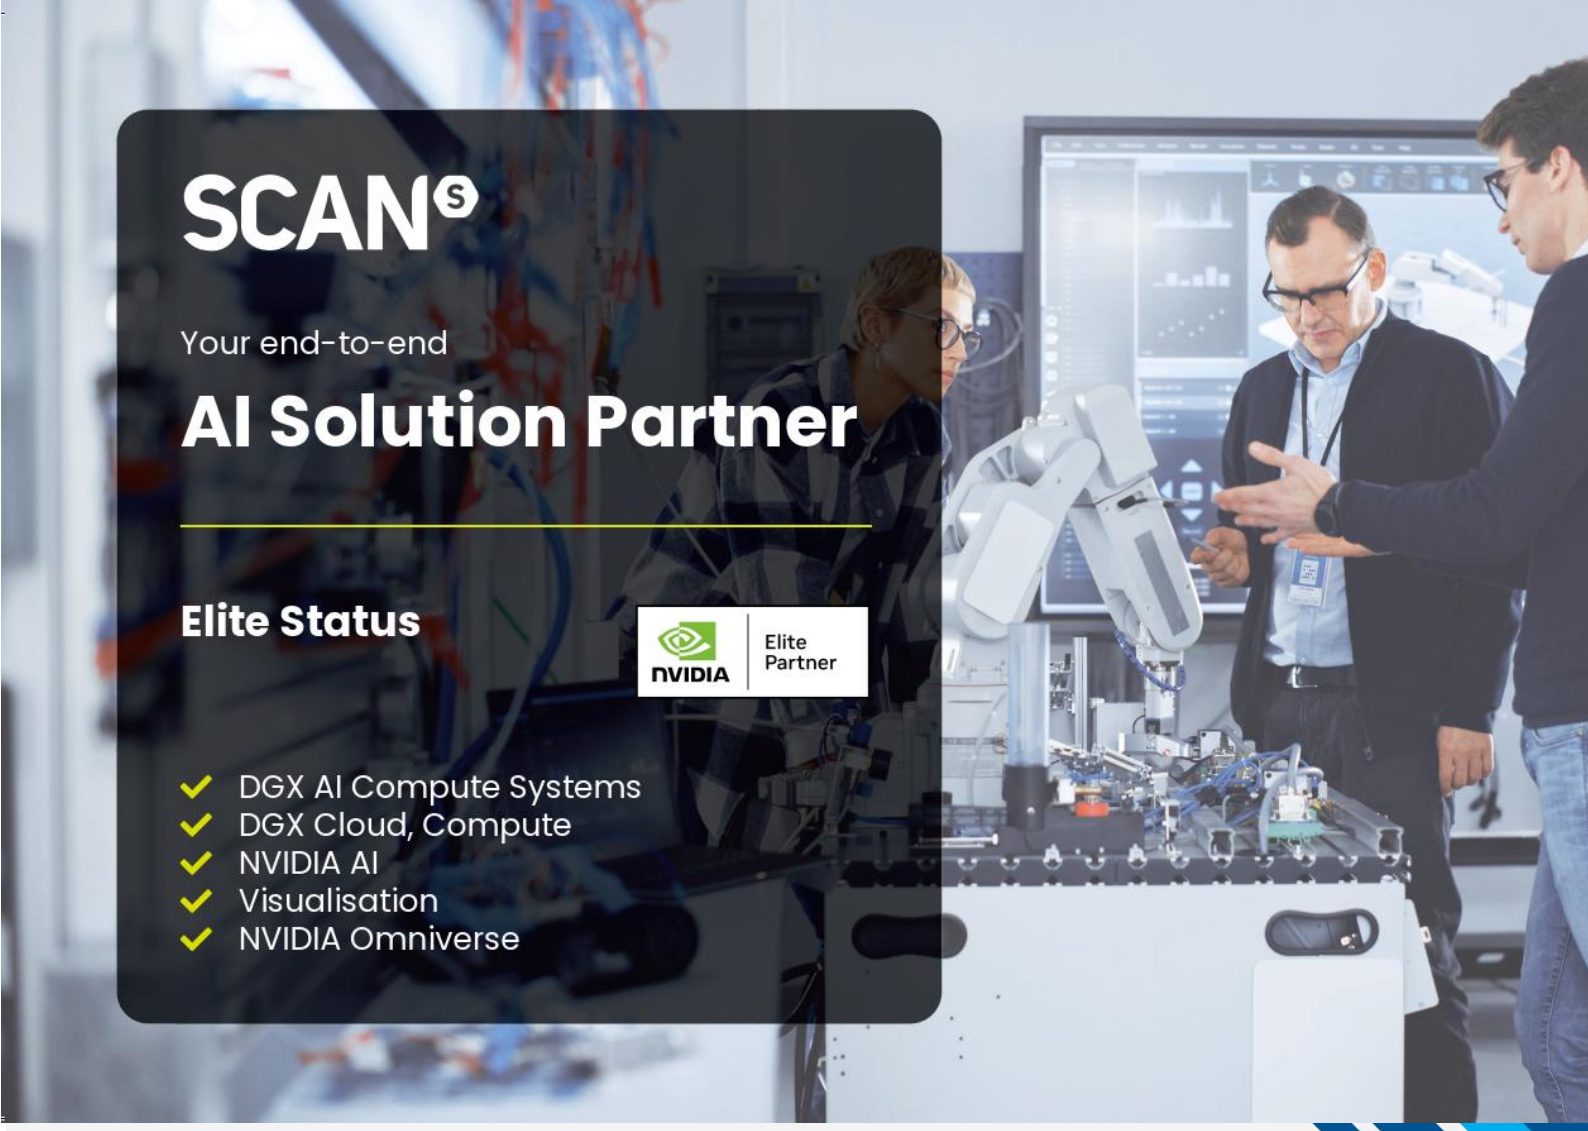

SCAN

Your end-to-end

AI Solution Partner

Elite Status

NVIDIA Elite Partner

- DGX AI Compute Systems
- DGX Cloud, Compute
- NVIDIA AI
- Visualisation
- NVIDIA Omniverse

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

Abstract: High-performance computing enables ever-larger materials simulations, but efficiency remains essential for both performance and sustainability. This session will focus on scalable implementations of the Wang–Landau flat-histogram method, widely used for calculating phase diagrams in atomistic and lattice models. We benchmark multiple strategies for distributing the workload—energy-domain decomposition, replica exchange, and multi-walker approaches—to determine where additional computational investment yields real benefit. Applying these methods to an AlTiCrMo alloy model, we show that adaptive, non-uniform energy windows deliver superior scaling and reduce wasted compute effort. The results emphasise how thoughtful algorithm design can enhance scientific output while conserving energy and computational resources.

Bio: Hubert J. Naguszewski is a final-year PhD student in the interdisciplinary EPSRC Centre for Doctoral Training in Modelling of Heterogeneous Systems at the University of Warwick, based in the Department of Physics. His research focuses on scalable algorithms for high-performance computing, with applications in computational materials science, particularly in alloy thermodynamics. He is also developing machine learning approaches to predict low-dimensional descriptors, such as the committor, from simulations, with the goal of generating interpretable insights for new reaction coordinates. His work emphasizes creating efficient, reliable, and understandable simulations while minimizing wasted computational resources.

Abstract: In this talk, I will discuss the use of the JAX library for solving problems in scientific computing. After a brief discussion of the main features of the library, I will focus on how these features enable researchers to develop more efficient implementations of their scientific models that can run on both CPUs and GPUs and that can interface with machine-learning models seamlessly. I will show examples of open-source JAX-based solvers for a few problems across computational science and engineering and share personal experience of developing software for applications in fluid mechanics and plasma physics. Finally, I will compare JAX to other solutions and discuss their relative limitations and strengths.

Bio: Dr Ubaid Ali Qadri has a background in computational modelling for fluid mechanics and leads the Multi-Fidelity Design and Twinning team at the Hartree Centre, STFC. He has more than 10 years

experience in developing and using adjoint techniques to understand sensitivities and optimal design parameters in multi-physics applications. He has worked across academia and industry on projects related to the aerospace, automotive and fusion sectors. At the Hartree Centre, Ubaid leads efforts in the application of novel simulation and design methodologies and high performance computing to industrial challenges.

Abstract: As university research computing services grow and expand in compute power and user-base, it becomes increasingly important to engage the users more proactively and ensure that the services are working for them and with them. With a little over a year in the role, the learning curve has been steep and has required the support and encouragement of a great team within Imperial and further afield. This talk will outline some of the strategies that have been used to improve the community around Research Computing at Imperial, the work that has been done to engage with the community on specific projects and the challenges that have been faced during the process.

Bio: Emily has been working as the Research Computing Services (RCS) Engagement Lead at Imperial College London for 18 months. Before this, she worked for 2 years as programme manager for the Centre of Injury Studies within Department of Bioengineering at Imperial and for 5 years as a project manager for the Computational Biomedicine Centre of Excellence at UCL. Her journey to these roles has been anything but straightforward, with a research and educational background in Chemistry, forays into Biology and a brief exploration of Drama straight out of school.

Overview of the NetworkPlus Programme – Friday 5 December, 12:00 - 13:00

Overview of the NetworkPlus Programme - Sophie Janacek - UKRI

Lightning Sessions

- CAKE: Nick Brown (EPCC)
- CaSDaR Louise Saul (University of Southampton)
- CHARTED : Weronika Filinger (EPCC)
- DisCouRSE: Jonathan Cooper (University College London)
- MInDS : Hajnal Farkas (Warwick University)
- NetDRIVE : Alex Owen (Queen Mary University of London)
- NFCSPlus : Jon Hays (Queen Mary University of London)
- SCALE-UP: Michael Bearpark <m.bearpark@imperial.ac.uk>

Followed by a Panel Discussion

Abstract: As Europe accelerates towards a future defined by data-intensive research and disruptive technologies, the need for sustainable, high-performance computing has never been greater. FSAS Technologies is proud to introduce Monaka – Fujitsu's next-generation Arm-based CPU – designed to deliver exceptional energy efficiency, scalability, and performance for HPC and AI workloads.

In this session, I will explore how Monaka provides a sustainable compute foundation for universities, research centres, and enterprises tackling climate change, healthcare, and digital inclusion challenges. Beyond traditional HPC, Monaka underpins the evolution toward quantum computing, enabling hybrid architectures where classical and quantum systems can collaborate seamlessly. This convergence will be key to solving the most complex problems of our time while keeping sustainability at the heart of innovation.

Join us to understand how FSAS and Fujitsu are shaping a responsible, energy-efficient pathway to exascale and beyond, bridging today's HPC with tomorrow's quantum future

Bio: Chris has been working with HPC systems since joining Altair Engineering in 2008. In 2013 he moved to Fujitsu, and worked on a number of systems, culminating as the HPC Lead for the Defence and National Security Chief Digital Officer.

Abstract: Founded in 2017 by our CSO, Dr Peter Leek, Oxford Quantum Circuits (OQC) has quickly become a prominent leader in the global quantum computing revolution. Our leadership is defined not only by an innovative approach to building performant, scalable, and high-quality superconducting qubits but also by providing Europe's first Quantum-Compute-as-a-service system. OQC has successfully transitioned quantum computing from the laboratory setting into data centres worldwide. This talk will outline how our current and future (4th and 5th) generation systems are designed to maintain our leadership position as we enter the 'logical era' of quantum computing. Furthermore, I will share our vision for the future, detailing how we plan to strategically organize our technology and teams to deliver the fault-tolerant quantum computers necessary to help solve the world's most significant challenges.

Bio: Lamar serves as the Director of Software Engineering at Oxford Quantum Circuits (OQC). He earned his PhD (Engineering) from the University of Sussex, where his research focused on developing ultra-fast, single-photon sensitive camera systems for applications in the biological sciences. His

professional background includes experience across both research organizations and industry, with notable positions at UKRI and GE Aerospace before joining OQC. Lamar is committed to applying engineering best-practice to solve complex problems and advance scientific inquiry.

KAYTUS CIUK

Experience KAYTUS End-to-End AI & Liquid Cooling Solutions

at #39

Fsas Technologies
a Fujitsu company

CIUK
Stand 45

The Future of HPC

MONAKA
Next-Gen Arm-based processor

Fsas Technologies Europe is a leading provider of advanced data centre solutions, HPC, AI and Quantum. This combines our world-class people with future-proof technology to enable customers and partners in Europe to harness the power of infrastructure for sustainable digital transformation. As a Fujitsu company, Fsas Technologies is dedicated using its expertise to realise the Fujitsu Group's purpose of bringing trust to society through innovation and making the world more sustainable.

Artificial Intelligence
On-prem AI: Secure and Trusted

Hybrid IT
Optimised performance, cost, security and workload portability

Data Resilience
Protect business data, compliance and secure growth

uSCALE Consumption IT
Combining cloud agility and on-prem control

Technology Innovation
For a sustainable and more efficient future

Contact:
education@fujitsu.com

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 3 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 3 December.

Thank you to the sponsors of the Pre-CIUK Networking Event.

On Thursday 4 December the official CIUK 2025 Networking Event will take place at Albert Hall, 27 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.

NVIDIA Partner

Hyperstack

ACCELERATE YOUR RESEARCH WITH ON-DEMAND GPUS

INSTANT ACCESS TO HIGH-PERFORMANCE NVIDIA GPUS FOR HPC, AI, AND DATA-INTENSIVE WORKLOADS.

- Works with SLURM, Docker, and on-demand Kubernetes
- Built for CUDA, PyTorch, and TensorFlow
- Autoscaling for dynamic HPC and AI workloads
- Pay-as-you-go GPUs — no queues, no contracts

Visit Booth 3 for FREE GPU credits
Start running your next experiment today!

WWW.HYPERSTACK.CLOUD

POWERED BY N3XGEN CLOUD

CIUK 2025 EXHIBITION

The CIUK 2025 Exhibition will take place in the Exchange Hall. The exhibition will include the exhibition stands, the cluster challenge and the research zone.

	**Web:** www.alces-flight.com
	**Web:** http://www.altair.com
	**Web:** https://thinkparq.com/
	**Web:** https://bios-it.com/
	**Web:** www.boston.co.uk
 	**Web:** https://www.uk.cdw.com/
	**Web:** www.cornelisnetworks.com
	**Web:** www.computacenter.com
	**Web:** www.ddn.com

	**Web:** https://deepgreen.energy/
	**Web:** www.define-technology.com
	**Web:** www.digitalrealty.com
	**Web:**
	**Web:** http://www.dell.com/en-uk **Web:** www.intel.com
	**Web:** www.eviden.com
	**Web:** https://www.fsastech.com/en-eu/index.html
	**Web:** www.hammerspace.com
	**Web:** www.hpe.com

Hyperstack	**Web:** https://www.hyperstack.cloud/
IBM	**Web:** https://www.ibm.com/uk-en/storage
KAYTUS	**Web:** https://www.kaytus.com/
Lenovo	**Web:** http://www.lenovo.com
mavin POWERCUBE® CONTAINERISED MODULAR DATACENTRES + CRITICAL INFRASTRUCTURE SOLUTIONS	**Web:** https://www.mavin.global/
novatech	**Web:** https://business.novatech.co.uk/
OCF	**Web:** www.ocf.co.uk
PEAK AIO	**Web:** https://www.peakaio.com/
point software & systems	**Web:** https://www.point.de/

	**Web:** www.qumulo.com
	**Web:** www.quobyte.com
	**Web:** www.redoakconsulting.co.uk
	**Web:** www.serverfactory.com
	**Web:** https://www.scan.co.uk/shop/ai-and-robotics
	**Web:** https://www.scalecomputing.com/
	**Web:** https://www.seagate.com/gb/en/
	**Web:** https://www.spectralogic.com
	**Web:** www.stackhpc.com
	**Web:** https://starfishstorage.com/

	**Web:** https://tenstorrent.com/
	**Web:** www.vastdata.com
	**Web:** www.vspl.co.uk
	**Web:** www.viridiengroup.com
	**Web:** www.vespertec.com **Web:** https://www.quantum.com/
	**Web:** https://www.vdura.com/
	**Web:** https://www.westerndigital.com/
	**Web:** https://www.weka.io

WEKA

Powering the Future of HPC and AI

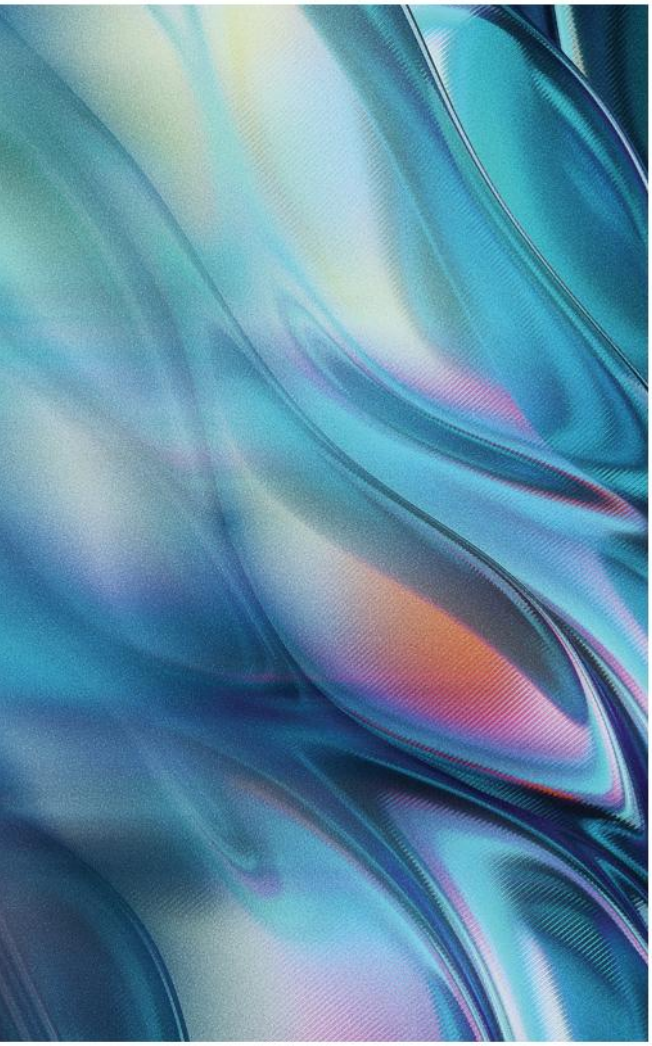

Meet WEKA at CIUK 2025 @ Booth 33

At WEKA, we're redefining how the world's most demanding HPC and AI workloads move, manage, and perform.

We're at CIUK 2025 to connect with the high-performance computing community, share innovation stories, and demonstrate how our purpose-built storage solution, NeuralMesh™ by WEKA eliminates infrastructure constraints and delivers the ultra-high throughput and microsecond latency required to maximize the efficiency of exascale HPC environments, and accelerates insights and outcomes. NeuralMesh enables researchers and enterprises to:

- Remove performance bottlenecks for **metadata-intensive workloads**
- Simplify management with **a zero-tuning architecture**
- Increase ROI with **over 300GB/sec to a single client** (single mount point)
- **Liberate metadata-intensive workloads** with fully distributed metadata management

From genomics to AI model training and inference, WEKA's unique approach delivers unmatched speed, scalability, and freedom giving scientists and engineers more time to focus on discovery, not data.

Visit Booth 33 to experience how WEKA powers breakthrough performance in HPC and AI environments worldwide.

Win Big with WEKA

Join us at Booth 33 and enter our Prize Draw for a chance to win a Mercedes-AMG F1 Lego Set — because at WEKA, we deliver more than just speed — it's innovation that keeps you laps ahead.

Stop by, chat with the team, and discover why the future of data performance runs on WEKA. Prize winner will be announced on the booth at 1 pm on Friday 5th December.

CIUK 2025 DAY ZERO

We will again be offering attendees at CIUK 2025 the opportunity to attend extra activities for no extra cost!

On Wednesday 3 December there will be a number of different activities taking place that will be available to attendees of CIUK 2025. Some will be included with your CIUK 2025 registration, others will require a separate registration. Full details of all the sessions can be found below.

Wednesday 3 December	**CoSeC Annual Conference 2025**	10:00 - 17:00	Exchange Rooms 8 – 10 (Upstairs at Manchester Central)
Wednesday 3 December	**Technical/SysAdmin Meetup**	13:00 - 16:00	Exchange Room 1 (Upstairs at Manchester Central)
Wednesday 3 December	**The Lustre User Group UK 2023**	14:00 - 17:00	Midland Hotel
Wednesday 3 December	**DRI Cybersecurity**	16:00 – 17:30	Exchange Room 1 (Upstairs at Manchester Central)

The CoSeC Annual Conference 2025 will take place on Wednesday 3 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 3 December - 13:00-16:00 - Exchange Room 1

This year at CIUK there will again be a meetup for technical staff to talk, learn & discuss important topics on Day Zero (Wednesday) from 1300-1600. This is for all technical people wherever you work, e.g. vendors, private companies, academic institutions, contractors and more. We would love to see as wide a representation as possible this year. After the success last year we will again be running this session as an unconference. If you can't arrive for the start then please just come along and join in when you can!

We would like to include a few short, focused technical talks (5 to 15 minutes each). If you are willing to give a talk please fill in the form: https://hpc-sig.org.uk/event/sysadmin-technical-meetup-facilitated-by-hpc-sig/

The Lustre User Group UK 2025 will be hosted alongside CIUK in Manchester on 3rd 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 13:00 to 17:00 and will take place in the Midland Hotel (*across from Manchester Central*).

Registration is FREE via the LUG UK website.

Please note this registration is separate from the general CIUK 2025 registration. If you are planning to attend any of the CIUK 2025 activities you must also register for CIUK 2025.

DRI Cybersecurity
Wednesday 3 December - 16:00-17:30 - Exchange Room 1

A meeting to discuss topics relevant to cybersecurity across the UK research and education landscape. Key areas would include:

- Foundational aspects of cybersecurity governance and risk management
- Operational Security and Incident Management
- Skills and Training
- building the relationship between operational security and the rich cybersecurity research community in the UK
- Common, coordinated standards and approaches that would benefit participation in the DRI community while providing essential assurance at a funding level.

Red Oak Consulting

Expert advice, exceptional delivery

Discover The Possible with Red Oak Consulting

Red Oak Consulting delivers expert solutions in High-Performance Computing, AI infrastructure, and Cloud Computing. With over 20 years' experience, we empower academia, research, and industry to innovate, adapt, and achieve sustainable results. Our comprehensive consultancy and managed services help clients overcome complex challenges and realise new possibilities

Helping to Unlock the Future of Research

- We collaborate with universities, research institutions, and industry leaders to design, build, and support cutting-edge IT and AI infrastructure,empowering you to meet today's challenges and seize tomorrow's opportunities.
- Our experts provide trusted guidance on managing and migrating research workloads in cloud and hybrid environments, always prioritising sustainability and compliance.
- Through comprehensive system assessments, we help you boost performance, minimise environmental impact, and drive research productivity beyond conventional limits.

Our Expertise

Specialist IT Consultancy: Decades of experience in HPC, AI infrastructure, and advanced cloud technology.

Research Infrastructure: Designing and optimising resilient, future-proof platforms to drive groundbreaking discoveries.

Strategic Planning: Forward-thinking guidance for IT strategy, business case investment, and building research environments that last.

Training & Support: Tailored training and ongoing managed services to maximise institutional capability and long-term impact.

Where computing unites -
accelerating collaboration, discovery and innovation

CIUK 2025 PARALLEL BREAKOUT SESSIONS

The CIUK 2025 Breakout Sessions will run alongside the main CIUK programme in the below rooms. Look for the PINK pull-up banners.

Understanding the dRTP Landscape

Thursday 4th December - 09:30-11:00 - Exchange Room 8-10

Organisers: Xu Guo and Weronika Filinger (EPCC, The University of Edinburgh - CHARTED), Jeremy Cohen (Imperial College London - SCALE-UP), Jonathan Cooper (University College London - DisCouRSE) and Adrian Harwood (The University of Manchester - DisCouRSE)

Description: This breakout is the third session in the series of events organised collaboratively by three digital Research Technical Professional (dRTP) skills Network Plus projects - CHARTED, SCALE-UP and DisCouRSE. Our goal is to engage with the community to understand, capture and map out the landscape of skills, roles and training content to improve access to dRTP roles in the UK. During this interactive session, we will build on the activities and data previously collected to better link skills to both training content and roles. We are hoping to focus on the different dRTP roles across different sectors to understand career progression paths.

Anyone interested and invested in training, professional skill development and workforce development is invited to attend and contribute to this session. The results will be disseminated to the community to help create a more sustainable and resilient DRI workforce.

FAIRifying the dRTP Training Ecosystem

Thursday 4th December - 11:30-13:00 - Exchange Room 8-10

Organisers: Weronika Filinger, Steven Carlysle-Davies (EPCC, The University of Edinburgh), Jeremy Cohen (Imperial College London) and Steve Crouch (The University of Southampton).

Description: Despite the multitude of training opportunities and resources available, many digital Research Technical professionals struggle to identify their learning needs and to find time to gain or progress desired skills. To better facilitate professional skill development, we need to make the training ecosystem easier to navigate by making the training resources more FAIR - Findable, Accessible, Interoperable and Reusable. This session will interactively build on previous efforts to adapt the FAIR principles to the dRTP training resources and develop good practices for FAIRifying training resources. Together we will explore what it means for the training resources to be FAIR, metric for evaluating and practical tips for improving FAIRness.

The outcomes of this session will be incorporated into a set of resources that will be shared with the community through the dRTP skills NetworkPlus project CHARTED (https://drtp-skills.ac.uk/), which will also offer opportunities to apply for funding to improve the FAIRness of the existing content. Anyone involved or interested in the development or provision of training content is warmly invited to attend!

Public Engagement activities you could deliver at your institution, and knowledge sharing and networking session!

Thursday 4th December - 14:00-15:30 - Exchange Room 8-10

Public Engagement activities you could deliver at your institution, and knowledge sharing and networking session!

Public engagement (PE) is an important part of our work. By delivering PE activities to children and adults in schools, festivals and at our workplaces, we can inspire the next (and past!) generations about how genuinely interesting, exciting and fun it is to work on high performance computing (HPC) systems and write applications for them. We will be sharing a small number of activities you could deliver, and also facilitating peer support, knowledge exchange, time to share experiences and have your questions about public engagement answered!

Attendees don't need any experience in delivering PE activities to attend this session - and if they do, we hope they will share ideas, hints, tips and tricks of what works for them, and take inspiration from the session.

Women in HPC Breakfast

Friday 5th December – 08:30-11:30 - Exchange Room 8-10

CIUK is thrilled to host the 4th annual Women in HPC Breakfast on **Friday 5 December**! Join us starting at 8:00 AM for breakfast, with presentations kicking off at 8:30 AM. We are bringing the HPC community together for:

- Inspiring early-career tech talks
- Introductions to the UK's WHPC Chapter & Affiliate community
- Networking with peers and colleagues
- A relaxed start to the day with breakfast provided

Please be aware that you will need a conference badge to attend, so make sure to register for CIUK. To see the full agenda, click here.

GPFS User Group 2025

Friday 5th December - 11:30-13:30 - Exchange Room 8-10

Continuing on from user groups at CIUK for a number of years, we want to run another Storage Scale User group at CIUK 2025. 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 beings 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 May 2025.

From ExCALIBUR to DRI, linking the UK's compute programmes

Friday 5th December - 14:00-16:00 - Exchange Room 8-10

ExCALIBUR, which ran between 2019 and 2025, aimed to get the UK's HPC software ready for exascale computing. In many ways, the Digital Research Infrastructure (DRI) programme follows on from ExCALIBUR but also takes things in a somewhat different direction with emphasis on different communities. This session aims to bring these two programmes together, to share the outcomes and successes of ExCALIBUR with DRI projects to ensure that the work done in the one programme can be leveraged, and is not reinvented, by the other. Linking these different groups will also help form and solidify networks between what is going on in the UK, and we also aim to influence those in the DRI programme to fund ExCALIBUR follow on style activities by taking some of the key outcomes and "pulling them through" into the DRI programme. We plan to have a series of talks from those involved in both programmes about their work, and will invite representatives from UKRI to be present at the session.

CIUK 2025 RESEARCH ZONE

The CIUK 2025 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.

VAST

Architecture Matters

The Data Foundation for Deep Research

visit us at
vastdata.com

CIUK 2025 STUDENT CLUSTER CHALLENGE

A NEW LOOK CLUSTER CHALLENGE FOR CIUK 2025... over recent years we have seen massive growth in the CIUK cluster challenge. From 3 teams in our first competition to 16 teams last year the challenge is proving to be very popular amongst university students. We have now become victims of our own success, and the previous format of the competition is no longer sustainable. So, we are making some positive changes...

- The CIUK2025 Cluster Challenge will see the introduction of a qualifying round. We are partnering with three Universities - The University of Edinburgh, UCL and Durham University. These Universities will be hosting a series of online challenges for the teams. Team members will also be required to complete the ARCHER2 driving test before commencing the challenges.
- At the conclusion of these online challenges the top 8 teams will qualify to attend CIUK2025 at Manchester Central in December, with all travel and accommodation costs covered by CIUK. Once at CIUK2025 the qualified teams will complete a series of in-person challenges working with our partner companies Alces Flight and OCF as in previous years to crown the overall CIUK2025 Cluster Challenge Champions 🏆
- Our champion team will then go forward to represent CIUK in the online version of the ISC'26 cluster challenge competition in June 2026.

This year's competition saw a record 17 teams join the qualifying round. Our 8 finalists can be found in the exhibition hall completing their in-person challenges. Please feel free to pay them a visit as they complete their challenges during the conference. You can also follow the competition on LinkedIn #CIUK2025_SCC.

The winning team will be announced before the final presentation session on Friday 5 December with the champions going forward to represent CIUK at the ISC'26 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…

Team: Beavers
London School of Economics


Rajib Ahmed


Zahin Ayman


Rachel Van Wijk


Mateusz Mikolajczak


Daria Godorozha


Aditya Sheth

Team: Imperial Bytes
Imperial College London


Scott Yapp


Maryam Fetanat


Marci Talbot


Freddie Nicholson

Team: ClusDur
Durham University


Carlos Valdemiro


Michal Pluta


Charles Dubois-Veltman


Laura Kraft


Keane Wong


Matthew Richardson

Team: Beariables
University of Birmingham


Harry Vicaradge


Sereya Mee


Jamie Cline


Ewan Hallet


Adam Wilson


Arthur Baverstock

Team: Hex
Queen Mary University of London


Rayaan Sheikh


Hardik Mathur


Oscar Gabriel Moreno Veliz


Mohamed Amiin Mohamed


Min Thant Kyaw


Abdulaziz Saud Albarrak

Team: Isambards
University of Bristol


Vladimir Dementyev


Tom Bale


Rowan Johns


Mobeen Mohammed


Kyrian Salas


Asher Bloom

Team: Clusty Boys
University of Exeter


Daniel Spicer


Wiktor Wiejak


Finn Van Monfort


Alex Morgan

Team: CodeSoc
University of Warwick


Sacha Cooper


Nok-Hang


Nevin Ahluwalia


Mayukha Pochana


Luke Turnyanskiy


Femi Segun

Open, Scalable Architecture

Empowering research with flexible, future-ready infrastructure.

Enterprise-Class Solutions

Delivering performance, reliability, and simplicity for demanding workloads.

Innovative Software, Rapid Updates
Stay ahead with cutting edge technology and seamless enhancements.

Secure, Supported Environment Confidence and compliance built into every solution.

Stable, Long-Term Research Platform
Designed to accelerate discovery and sustain growth for years to come.

CIUK 2025 POSTER COMPETITION

24 finalists have been chosen for the CIUK 2025 Student Poster Competition. Their posters can be viewed on the exhibition floor. The winner will be announced prior to the final presentation session on Friday 5 December.

James Connell *UKRI - STFC*	**Pointing the Way: Azimuth for the STFC Cloud** **Abstract:** Azimuth, developed by StackHPC, is a self-service portal that simplifies creating and managing virtual machine clusters and cloud resources for scientific computing, HPC, and AI. It abstracts away platform complexity, enabling container scalability and allowing researchers to focus on the science. This work presents STFC Cloud's six-month deployment of Azimuth, highlighting backend automation and integrations powering a seamless, flexible service. We cover key concepts (IaaS, Kubernetes, OpenStack) to explain Azimuth's role within the STFC Cloud, detail our early-production deployment, its applications, and the tools for managing tenancies, quotas, and monitoring. Finally, we will discuss the utility Azimuth will have, using the Square Kilometre Array (SKA) as a case study, providing scientists with a turnkey solution for deploying software across sites without infrastructure expertise—accelerating research as adoption grows. **Bio:** James Cornall is a Scientific Computing Graduate within STFC's Scientific Computing Department. He recently concluded his third rotation as part of the SCD Graduate Scheme with the STFC Cloud team. His current rotation is with PSDI, with a focus on automating data annotation for use in AI.
Blair McCormick *UKRI – STFC* 	**A Modular Partitioned Solver for Conjugate Heat Transfer in Liquid Hydrogen Tanks using MUI** **Abstract:** The HASTA project is an EU-funded consortium aiming to develop digital twins of liquid hydrogen (LH_2) tanks for civil aircraft. Within the consortium, a team at the STFC is working towards a partitioned solver for Thermal-Fluid-Structure Interaction. In this approach, several solvers are used to simulate different regions (e.g. the LH_2 and tank walls) and are coupled together in a distributed memory fashion. This is achieved using the MUI library for data transfer at shared interfaces. As a part of this project, I have developed a coupling framework to simulate Conjugate Heat Transfer. The framework couples OpenFOAM for fluid dynamics with FEniCSx for solid heat conduction. Crucially, this allows OpenFOAM to be coupled for the first time without modifying its source code. Validation with an existing solver demonstrates excellent accuracy with significantly improved flexibility and extensibility marking a key milestone for developing a fully coupled simulation of the LH_2 tank. **Bio:** I am a 22-year old Scientific Computing Graduate at the STFC, currently working in the Coupled Problems in Fluids Engineering group. Prior to this role, I studied Computer Science at Durham University with a focus on Algorithmic Design and HPC systems. I am now applying these skills to the HASTA project where we are developing digital twins of

	Liquid Hydrogen tanks for use in Civil Aviation. In my spare time I love spending time outdoors and rock climbing.
Dr Dimitrios Bello *The Rosalind Frankin Institute* 	**Bridging the Gap: Aligning Project Administration with Access to Compute and Data Services** **Abstract:** High performance computing (HPC) and data services underpin a growing share of UK research, yet researchers often face delays and duplicated effort when gaining access. Administrative systems collect essential project information, but this metadata rarely flows through to the HPC and infrastructure teams who need it. The result is manual bottlenecks, inconsistent practices, and governance risks. These challenges are especially acute in multi-institutional collaborations where speed, scale, and security are most important. The Bridging the Gap project, funded through the NFCS NetworkPlus programme, is exploring how federated identity frameworks can connect research administration systems with HPC and digital infrastructure provisioning. By mapping project metadata directly into access workflows, we aim to reduce duplication of effort, accelerate time to compute, and improve compliance. This poster presents initial feedback from stakeholder interviews, the current overview for the framework and the ongoing roadmap for the project. **Bio:** Dimitrios is a researcher and software developer, in the Rosalind Franklin Institute. He is a member of the Advanced Research Computing (ARC) platform within Frankin, 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 is a member of the Digital Twin Cell project within Franklin's Artificial Intelligence and Informatics for Predictive Biology challenge. As part of this project, he is researching DL/ML approaches for the processing or enhancing of Electron Tomography and/or 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.
Willow Sparks *UKRI – STFC*	**GPU Acceleration of Monte-Carlo Electron Scattering Simulations** **Abstract:** ALC_SUTOR is a Python package for exploring scattering problems seen in electron beamlines such as Diamond-ePSCI and Diamond-HeXI. A central component of this is the Monte-Carlo integration of electron showers via spectraOPT, a CPU-based software implemented using C++ and MPI. Due to the high degree of parallelism inherent to the underlying model, it has become desirable to take advantage of GPU acceleration. We present a new Python-based implementation of the integrator, targeting CUDA GPU hardware via the use of Numba-CUDA and CuPy. The integrator maintains full feature parity with the original implementation, while supporting both single-GPU and multi-GPU execution for significant improvements to simulation

runtimes. We demonstrate this through a test study on CsPbBr3, exploring the reflection and transmission of primary and secondary electrons over a range of beamline energies. This outlined approach will assist in future theoretical investigations of electron-material interaction, while also showcasing the efficacy of using JIT-compilation for high performance computing.

Personal Bio: Hi there! My name is Willow Sparks and I'm a graduate computational scientist at STFC. I joined last year after completing an MPhys in Theoretical Physics at Lancaster University, and since then I've been exploring lots of different areas of HPC. I've worked on various exciting projects over the past year, ranging from condensed matter physics to GPU accleration to big data compression. I'm always happy to chat about anything related to scientific computing and HPC, so feel free to reach out!

Charlie MacMahon-Gellér

DDN / DiRAC

Fast Workflow Rooflines for AI

Abstract: AI workflows place diverse and often competing demands on modern HPC systems, stressing compute, memory, and filesystem I/O in ways that vary significantly between applications. Accurately identifying system-level bottlenecks is essential for efficient hardware selection, workload placement, and performance optimisation, yet traditional deep-dive profiling remains time-consuming, intrusive, and difficult to scale across architectures and workflows. We present a portable workflow-roofline modelling framework designed for fast, high-level bottleneck diagnosis of AI workflows. The framework combines a curated, lightweight toolset with automated installation and wrapper scripts to enable monitoring and benchmarking on a wide range of systems. By calibrating system ceilings via microbenchmarks and mapping workflow behaviour onto both classic and workflow roofline models, our approach rapidly exposes whether compute, memory bandwidth, or filesystem I/O bandwidth forms the dominant performance constraint. We test the framework using the Dafab Field Delineation workflow, a CNN-based pipeline operating on Copernicus Sentinel-2 satellite imagery. Evaluation of our framework on an NVIDIA H100 node reveals clear filesystem I/O bottlenecks; reducing workflow write amplification improved end-to-end throughput by 21.5% and we observe a 7.5x cost efficiency advantage over a CPU-only node, based on node rental costs. Our results highlight the value of workflow-level rooflines and our curated framework as a practical, system-agnostic approach to guiding AI workflow optimisation and informing HPC system design. Further work will extend our framework to include automated benchmarking, add Slurm-ready wrappers, and broaden evaluation to transformer-based workloads and additional hardware platforms.

Bio: I am a 3rd year PhD student in Astrophysics currently undertaking a DiRAC innovation placement in partnership with DDN. My placement research has been focused on studying the dependence of various AI workflows on different system hardware components, to identify bottlenecks and aid system design and workflow optimisation. In my PhD, I look at using AI to accelerate astrophysics research, GPU based optimisations of galaxy simulations, and the effect of tidal forces on galaxy shapes.

Tom Kirkham

UKRI – STFC

DAFNI DINI

Abstract: This poster summarises work done on the Data for Infrastructure for National Infrastructure (DINI) project within the STFC DAFNI team. The work is motivated by infrastructure research and its high impact on the social, economic and environmental well-being of the UK. DINI focus on the needs of the computational modelling and machine learning communities and the challenge they face in the availability of quality data. Data is an essential prerequisite for good analysis and good decision making, but there are many barriers to the effective use of data, which are prevalent in the infrastructure systems engineering domain, including: Discoverability, Licensing, Sensitivity, Metadata, Interoperability, Reliability and Ethics.

Bio: Dr Tom Kirkham is the Science Lead for the Data Analysis for National Infrastructure (DAFNI) programme at STFC. DAFNI is a free to use platform providing a space for researchers to share data and models within collaborative workflows.

Zeyuan Miao

University of Manchester

Real-Time Digital Twin of Fusion Material Experiments Using AI and HPC

Abstract: Understanding tritium transport in fusion materials is critical for safe and efficient fusion energy system design. Gas-Driven Permeation (GDP) experiments provide essential data but are complex and require precise control. We present a digital twin framework that integrates artificial intelligence and high-performance computing (HPC) to accelerate and optimize GDP experiments. A large language model (LLM) combined with a structured knowledge graph enables rapid retrieval of historical experimental data and decision support for planning new trials. Simultaneously, real-time HPC simulations infer the true sample temperature from indirect thermocouple measurements, allowing dynamic adjustment of experimental conditions. Results are visualized in a 3D virtual environment via NVIDIA Omniverse, supporting interactive monitoring and control. This approach enhances experimental efficiency, reduces redundancy, and provides deeper scientific insight into tritium permeation behaviour.

Bio: Zeyuan Miao is a research associate focused on developing surrogate modelling and digital twin techniques to improve traditional simulations. During his Master's and PhD at the University of Manchester, he concentrated on enhancing the efficiency and accuracy of predictions related to the structural integrity of weldments, leveraging methods such as automated data generation and machine learning models including Artificial Neural Networks (ANNs) and Physics-Informed Neural Networks (PINNs). More recently, Zeyuan has applied his expertise to real-time digital twin frameworks for fusion material experiments, integrating AI, high-performance computing (HPC), and immersive visualization to accelerate Gas-Driven Permeation (GDP) experiments. He is also interested in using surrogate models for fusion power plant design and aims to advance simulation techniques across a range of scientific and engineering fields.

Luke Webb *Queen Mary University of London* 	**Using digital twins to analyse liquid-solid impact on soft polymers** **Abstract:** Wind turbine blades degrade rapidly from leading edge erosion, increasing the levelized cost of energy. Inspired by polymeric impact-resistant coatings on blades, this study aims to develop a validated digital twin of a droplet impacting a soft polymer. This is intended to provide understanding of the liquid-solid impact damage mechanisms and a bridge from singular impact to long-term damage. Validation data is captured from a single jet impact test, using localised spectrum analysis to capture the surface displacement, strain and acceleration fields of the substrate. The acceleration field informs the material model, capturing the effect of ultra-high strain rates on the polymer. The numerical model closely replicates the experimental results. Model development centres on domain coupling and shock to quantify droplet damage potential. Results show high uncertainty in the predicted pressures, requiring further testing for perspective at the damage-scale. **Bio:** A 2nd year materials science PhD student in Dr Wei Tan's Mechanics of Composite Materials group. Motivated by the development of sustainable technology, my research focusses on the issue of leading edge erosion in wind turbine blades, with a background in numerical modelling and data science. I specialise in developing impact-resistant materials through the use of digital twins to better understand fluid-structure damage mechanisms.
James Legg *University College London* 	**A digital home for your research group with MyAccessID** **Abstract:** The National Federated Compute Services NetworkPlus initiative (https://nfcs-networkplus.ac.uk/) is “engaging compute service users engaging compute service users and providers who are focused on the three pillars of governance, technology, and community, to develop a roadmap for a national federation of compute services to support the UK’s digital research infrastructure.” The Universities of Bristol and Cambridge, and The Hartree Centre and UCL currently have a joint project (https://nfcs-networkplus.ac.uk/projects/exploring-the-governance-requirements-for-enabling-uk-dris-to-adopt-myaccessid/) funded by that Network on the adoption of MyAccessID (https://wiki.geant.org/spaces/MyAccessID/pages/290717698/MyAccessID+Home) as a federated identity verification system. I am the Co-lead of a subproject at UCL’s Centre for Advanced Research Computing, to build a test tenancy in our private cloud utilising MyAccessID to verify logins. This technology provides a simple way for users from other universities to have accounts for and log in to compute facilities belonging to a particular university. Our model for the compute facility is a “tenancy”, a set of computer resources and storage, allocated to a research group in our private cloud. The reality is that such groups have external team members who, unless they can log in, are excluded from day-to-day access to the group’s data and compute resources. In some cases, it may be permissible for external users to upload and download data, but this is tedious to operate and can lead to delays and version control problems. The subproject will provide a case study in how to set up such a facility which will document not only the technical arrangements but the regulatory and organisational issues that system operators will need to take into account. A particular novelty is that

	we shall be working with MyAccessID's newly appointed agent in the UK as they develop their system to handle and approve registration of services using MyAccessID into the MyAccessID system. The poster comprises a central illustration of the model facility together with text listing likely organisational and regulatory issues, links to useful resources and an exhortation to get involved in the community our project is building to support adoption of this technology. **Bio:** James Legg has been working at UCL's Centre for Advanced Research Computing for 3 years now on their Project Contender, evaluating novel architecture machines, and was recently awarded a PhD from UCL's CDT for Data Intensive Science. For many years, he worked in patents for electronics, semiconductors, computer systems, communications and sensors and as partner in charge of IT systems at his law firm, but has also worked manufacturing superconducting microcircuits and as a concrete block dispatcher. Earlier studies were physics, semiconductors and intellectual property at Cambridge, Imperial and Queen Mary. He learnt Linux and Python on a Raspberry Pi and still has his original BBC Micro.
Larisa Dorman-Gajic *UK Atomic Energy Authority* 	**Fusion of Talent: Celebrating the Many Roles of Women in Computing** **Abstract:** In November 2025, the UK Atomic Energy Authority hosted a one-day Women in Computing event uniting women, underrepresented genders, and allies from academia, research, and industry. The importance of events like these is underscored by the Lovelace Report 2025, which found that each year 40,000-60,000 women leave UK tech roles. This highlights systemic issues that persist in computing, starting with stereotyping from a young age causing fewer women to join computing in the first place, followed by the exodus of women during their mid-career. The day featured networking, talks, a poster session, and a panel on the barriers faced, fostering open discussion and connection. The poster reflects on lessons learnt from the event, exploring ways to help women and other minorities in computing feel more welcome. **Bio:** Larisa is a Data Engineer in the Computing Division of the UK Atomic Energy Authority. She has worked on various projects in the data management of facilities and research projects within the organisation since she started in 2023. She studied Theoretical Physics (MPhys) at the University of Edinburgh where she learnt how to programme. After graduating she worked as a Data Analyst in industry before entering back into the research space.
Arjun Shergill *University of Birmingham*	**GPU-Accelerated Simulation of Butterfly Wing-Wake Interactions** **Abstract:** The flight of insects has attracted significant research attention due to their remarkable manoeuvrability at low Reynolds numbers, enabled by employing unsteady vortex dynamics and specialised wing morphology. Of these, butterflies are particularly interesting: They are able to change the aspect ratio and wing sweep angle of their

flexible, partially overlapped wings during flight, and different species use distinct wing shapes that produce unique aerodynamic behaviours. This work models three butterfly species with rigid wings – the Monarch, Swallowtail and Indian Leaf – using a GPU-accelerated Immersed-Boundary Method combined with an Overset Mesh. This approach allows one to capture high-fidelity vortex structures efficiently, avoiding the high cost of traditional mesh methods such as Adaptive Mesh Refinement (AMR) and overcoming the limitations of CPU-based solvers. Dual parallelism with OpenMP and OpenACC further improves the performance. Simulations reveal a consistent formation of Leading-Edge, Wing-Tip, and Trailing-Edge vortices. The swallowtail maintains an attached leading-edge vortex (LEV), while the Monarch and Indian Leaf show inboard LEV detachment due to their forward-swept forewings. This aligns with previous research. Overall, the present study demonstrates that GPU-accelerated IBM-Overset methods can accurately resolve complex flapping-wing problems. Future work will incorporate wing flexibility, body motion, and large-scale multi-body simulations.

Bio: I am a first-year PhD student in Aerospace Engineering at the University of Birmingham, working under the supervision of Dr Chandan Bose on the development of a GPU-accelerated flow solver for simulating bio-inspired flight. I previously completed an MEng in Aerospace Engineering at the University of Birmingham, during which I completed an undergraduate research project on the wake dynamics of tandem pitching airfoils, also supervised by Dr Chandan Bose. The present study focuses on the unsteady aerodynamics of butterfly flight, a problem that demands large computational resources due to the complexity of vortex structures. Current CPU-based HPCs cannot resolve these unsteady vortices. To address this limitation, my work focuses on the use of GPU-acceleration with dual parallelization of OpenAcc and OpenMP. This research aims to further develop our biological understanding of the butterfly family, and provide insight into bio-inspired Micro Air Vehicles (MAVs).

Hibah Saddal

University of Birmingham

Modelling the Gust Response of Bio-Inspired Morphing Wings Using a Partitioned Strong Coupling Approach

Abstract: This study presents an open-source framework for the coupled modelling of the fluid–structure interaction in bio-inspired wings subjected to gust disturbances at a low Reynolds number. Three wing configurations are investigated: a conventional NACA0012 airfoil, along with bio-inspired peregrine falcon and barn owl wing sections.

The incompressible fluid flow, governed by the Navier-Stokes equations, is solved using a finite volume approach. In contrast, the hyperelastic structure, governed by the Neo-Hookean model, is solved using finite element analysis. Two-way coupling between the fluid and the structure is facilitated through a partitioned strong coupling approach, using radial basis function mapping and the Interface Quasi-Newton Inverse Least Squares (IQN-ILS) acceleration scheme.

Simulations examine the influence of wing morphology, the dimensionless bending rigidity and gust transition speed on the aerodynamic response and wake dynamics.

Simulations were executed on the ARCHER2 high-performance computing system using scotch domain decomposition.

	Bio: I am a second-year PhD student in Aerospace Engineering at the University of Birmingham, working with Dr Chandan Bose on modelling and simulation of fluid-structure interaction systems with variable flexibility. Prior to studying engineering, I completed a BSc in Mathematics from Coventry University, followed by an MSc in Applied Mathematics from the University of Birmingham. My PhD project focuses on modelling bio-inspired fluid-structure interaction problems, using a computational fluid dynamics approach with open-source codes. My research requires a coupled numerical framework for studying two-way fluid-structure interactions. Given the computational intensity of these simulations, High Performance Computing (HPC) resources are crucial. Some research problems include the implementation of covert-inspired flexible flaps attached to the upper surface of airfoils to enhance lift and delay stall at high angles of attack, as well as the investigation of flexible morphing wings as an alternative to conventional rigid wings for improved aerodynamic performance. The findings of these studies have applications in the design of unmanned aerial vehicles.

HAMMERSPACE

NVIDIA Inception Program

Your Compute Is Fast. **Your Storage Should Be Too.**

Stop letting storage bottlenecks slow down your HPC workloads.

Hammerspace Delivers:

- A Parallel Global File System for **extreme performance**
- Hybrid cloud data orchestration **across HPC sites and clouds**
- A Unified global namespace to **break down storage silos**

Learn more at **HAMMERSPACE.com**

Sam Ezeh

Durham University

ALARM: Automated Latency-Aware Rank-Reordering in MPI

Abstract: Rank-reordering optimises communication-heavy HPC workloads by assigning MPI ranks to physical cores such that most communication happens within a compute node. In the past, this technique has been used to achieve a 78% reduction in MPI_Allgather latency and a 30% reduction in execution time in an example N-body application. However, unlocking these gains typically involves complex manual interventions, such as defining virtual topologies, rewriting the code-base or creating an MPI machinefile. We present a generic, implementation-agnostic tool that fully automates the profiling and rank-reordering procedure. Our tool first runs the application in a profiling mode to infer the application's communication pattern. It then uses this profile to automatically reorder ranks on subsequent executions. This tool does not require any source code modifications or binary recompilation and supports existing applications written in Fortran and C/C++. In the spirit of the power of community, our automated tool empowers a broader range of scientists to achieve significant performance gains, democratising an optimisation technique that was previously accessible only to HPC experts.

Bio: Sam Ezeh is a Computer Science (MEng.) student at the University of Durham. He specialises in High-Performance Computing and low-level systems and competed in the IndySCC 2024 Student Cluster Competition on the winning UKSCC team.

Sam returns to the CIUK community this year after competing in last year's student cluster competition with the ClusDur Durham HPC society team. Drawing on performance engineering experience from internships at Qube Research & Technologies and Google, his poster presentation, ALARM (Automated Latency-Aware Rank-Reordering in MPI), introduces a tool designed to fully automate MPI rank-reordering for communication-heavy workloads without source code modification.

Edoardo Altamura

UKRI – STFC – Hartree Centre

Qiskit Machine Learning: an open-source library for quantum machine learning tasks at scale on quantum hardware and classical simulators

Abstract: We present Qiskit Machine Learning (ML), a high-level Python library that combines elements of quantum computing with traditional machine learning. The API abstracts Qiskit's primitives to facilitate interactions with classical simulators and quantum hardware. Qiskit ML started as a proof-of-concept code in 2019 and has since been developed to be a modular, intuitive tool for non-specialist users while allowing extensibility and fine-tuning controls for quantum computational scientists and developers. The library is available as a public, open-source tool and is distributed under the Apache version 2.0 license. According to the results of the Unitary Foundation's Annual Quantum Open-Source Software Survey, Qiskit ML is the most used open-source quantum application library in 2025, selected by over 27% of respondents. It also topped this ranking among Qiskit Community packages. Moreover, Qiskit ML recently received

	its 1000th contribution, has been forked over 400 times, and has surpassed 1 million downloads. **Bio:** Dr. Edoardo Altamura is a research scientist and quantum software engineer specialised in quantum algorithms and their applications to the natural sciences. He is a Principal Investigator in several award-winning STFC–HNCDI quantum computing projects, and a core maintainer of Qiskit Machine Learning, one of the world's leading open-source libraries for quantum machine learning, with over one million downloads and thousands of active users. With over six years' experience in research, Edoardo has led successful cross-sector projects with partners from industry (AstraZeneca, IBM, Algorithmiq) and academia (University of Cambridge, Manchester, UCL, Dresden). At the STFC Hartree Centre, he has driven the development of high-impact quantum software, hybrid algorithms, and open-source tools that have contributed to over £1M in collaborative R&D value and received recognition from IEEE and UKRI. He also serves as a Visiting Research Fellow at the University of Cambridge (Yusuf Hamied Department of Chemistry, UK), where he co-advises PhD students and leads collaborative work on quantum machine learning and quantum chemistry methods at the utility scale. Edoardo earned a PhD in computational astrophysics at the University of Manchester (Jodrell Bank Centre for Astrophysics), where he developed large-scale hydrodynamic simulations of galaxy clusters using high-performance computing. In this field, he is most cited for his publications on the entropy-core problem and the rotational kinetic Sunyaev–Zel'dovich effect, which advanced understanding of gas dynamics in massive galaxy clusters.
George Roe *UKRI – STFC* 	**GUILT** **Abstract:** For those who use High-Performance Computing systems, the environmental cost of their work is often overlooked. With no built-in tools for measuring or displaying carbon emissions on most systems, users are left in the dark about their energy consumption and impact on the climate. Here I introduce GUILT, a command line interface that provides users with data analytics about their carbon emissions. Designed to work on any SLURM-based cluster, GUILT can analyse current and past jobs to identify their environmental cost and suggest alternative times to run jobs to minimise emissions. This project forms part of the STFC Scientific Computing Department Energy Efficiency Team strategy to transition the department towards net zero. **Bio:** I am George Roe, a degree apprentice software engineer at the Scientific Computing Department at STFC. I am currently studying towards a Degree from Manchester Metropolitan University in Digital and Technology Solutions. I am conscious of how scientific research can impact the climate and thus have been working with the SCD Energy Efficiency Team to produce free tooling to aid users in reducing their carbon emissions.

Ashraf Hussain *UKRI - STFC*	**From Metrics to Mitigation: Visualizing and Reducing the Carbon Footprint of STFC's Ada Cloud Computing Service** **Abstract:** This project presents a suite of interactive visual tools designed to monitor, forecast, and reduce carbon intensity and energy usage. The Carbon Intensity Forecast visualizes predicted carbon intensity over a 48-hour period, enabling users to schedule activities during low-impact periods, with highlighted windows indicating the greenest operating times. Complementing this, Usage Graphs display electricity consumption and associated carbon emissions, broken down by project, machine, experiment, or user—revealing inefficiencies such as excessive idle usage in ISIS workspaces. A GitHub-style heatmap offers an intuitive, color-coded overview of daily carbon footprint trends, helping users quickly assess whether their operations are becoming greener over time. Additional features include workspace tracking, allowing individuals to monitor energy and carbon metrics in real time, and machine size usage tracking, providing average emission data to encourage efficient resource choices. Together, these tools aim to raise awareness, promote sustainability, and support data-driven decision-making within the STFC community. **Bio:** I'm Ashraf Hussain. I'm a third year Software Engineering Apprentice at STFC.
Sampreet Kalita *University of Strathclyde Glasgow* 	**A family of Frameworks for Efficient Multi-mode Bosonic Computing and Applications** **Abstract:** The Lossy Computational Models project entails the development of natural models of computing with photonic systems where their losses are utilized as features. As such, together with concrete mathematical formalisms, the validation of such computational models also requires efficient numerical simulations with several bosonic modes. This has led us to develop an ecosystem of numerical tools. We call this family of Python frameworks PyBos! We have recently analyzed single-photon interferometric effects with lossy photonic devices containing hundreds of spatial modes. A notable output of our investigation is two distinct loss models of beam-splitters, which lead to two unique photon-count statistics. Our work also throws light on how photon losses and their (in)distinguishability influence the success of optimization tasks for boson sampling and machine learning. Such unconventional approaches to computing can serve as blueprints for energy-efficient, nature-inspired alternatives to mainstream computing. The poster introduces the current members of the PyBos family - 'base', 'sampler' and 'parser' - and the ingredients in their speedup, highlights the domain and focus of our implementation and demonstrates the different scenarios we have studied using PyBos. **Bio:** Sampreet Kalita is a Research Associate in the Department of Physics at the University of Strathclyde Glasgow, currently working on the Lossy Computational Models project. He holds a Doctorate from the Indian Institute of Technology Guwahati in the physics of macroscopic phenomena arising out of light-matter interactions. Before joining the University of Strathclyde, he was a guest researcher at the Max Planck Institute for the Science of Light, where he investigated co-operative phenomena in quantum systems and applications of reinforcement learning for ground state cooling. In the past years, he has also contributed to the development of multiple open-source software frameworks for quantum systems.

Gokmen Kilic *Research Software Engineer* 	**The DiRAC Resource Allocation Framework: Maximizing Scientific Research and Innovation Through Strategic HPC Resource Allocation** **Abstract:** The DiRAC (Distributed Research utilizing Advanced Computing) facility represents the UK's premier national High-Performance Computing (HPC) infrastructure dedicated to theoretical cosmology, nuclear physics, astrophysics, particle physics, and solar and planetary science. The DiRAC Resource Allocation Committee (RAC) provides a transparent, peer-reviewed mechanism for UK researchers to access these computational resources. This presentation outlines the RAC process, detailing how researchers can successfully navigate the application system to secure computing time and Research Software Engineering (RSE) support for their projects. **Bio:** Gokmen holds a BA in Computer Engineering and an MBA from Black Sea Technical University, and a Ph.D. in Finance from the University of Bingol. Prior to his current role, he was an IT Engineer/Lecturer at the University of Artvin and a Visiting Academic at Durham University Business School. Gokmen now works as a Research Software Engineer (RSE) in the Physics department at Durham University, supporting DiRAC projects.
Raska Soemantoro *The University of Manchester*	**Quantifying the Cost of Geometric Complexity in Monte Carlo Neutronics Simulations** **Abstract:** Monte Carlo neutronics simulations are important in fusion energy for evaluating performance and shielding requirements but are computationally expensive, with significant implications for runtime and energy consumption in HPC environments. Whilst complex geometries are widely accepted as more computationally expensive than simple ones, the relationship between complexity and computational cost has not been studied quantitatively for neutronics simulations. Quantifying complexity is challenging because there is no universally accepted metric, and complexity can be subjective. Some works have proposed complexity measurements for CAD geometry that correlate human-judged complexity with measurable quantities such as number of faces or vertices. Here we extend this concept by applying such measurements to OpenMC test cases, examining whether metrics correlating with human-judged complexity also correlate with computational cost. Using OpenMC on ARCHER2, a parametric study was done, varying surface features on CAD geometry subjected to a fusion-relevant neutron source. Results demonstrate that as geometric complexity increases, the transport phase dominates an increasingly larger fraction of total computational time. **Bio:** I am a PhD student and research assistant interested in computer-aided design and machine learning. My thesis is entitled ‘Automated de-featuring of CAD geometries for simulation of complex systems’.
Benjamin W Butler *University of Strathclyde*	**Building nature-inspired models of computation** **Abstract:** The development of unconventional and nature-inspired computing platforms is one way to address many of the challenges which face modern computing, not least the need for greener compute power. In the Lossy Computational Models (LoCoMo)

	project, we are developing a methodology for abstracting a model of computation from the natural physics (or chemistry, or biology) of a substrate or smart material, which will allow future researchers and engineers to answer the question: what are the computational capabilities of this device? **Bio:** Benjamin W Butler is a theoretical physicist in the Computational, Nonlinear & Quantum Optics group at the University of Strathclyde. He is currently working on the development of nature-inspired computational models which incorporate photonic loss and gain as features, rather than bugs. He completed his PhD at the University of Glasgow in 2025, where his research focussed on structured light and atom-photon interactions.
Robin C. Laurence *Department of Mechanical and Aerospace Engineering, The University of Manchester* 	**The digital, made physical... made digital, again! Providing the tools for residual stress optimisation in metal additive manufacturing** **Abstract:** Metal components made though additive manufacturing are becoming ever more present in industrial high value manufacturing. Before wide spread adoption, especially to safety critical roles, the manufacturing process must be both better optimised and well understood. One such additive manufacturing process is wire arc additive manufacturing (WAAM) which is capable of producing large scale metal components (on the order of meters) through specific application of molten metal, which then solidifies building up a component layer by layer. This process results in high temperatures and high temperature gradients in these components, leading to both distortion during the build and potentially high levels of residual stress, which could lead to component failure. These stresses can be modelled using finite element methods, providing insight into internal stress distributions, but in order for the modelling to remotely accurate many details of the manufacturing process must be recorded. A work flow is presented, using an example component, showing how all the necessary information can be captured through use of a digital twin which combines data from many inputs, storing them all in a single location, and how then, combined with other software the finite element model can be prepared. The finite element job is carried out by a HPC facility resulting in overall time and energy savings. **Bio:** Dr. Robin Laurence is a research associate in mechanical and aerospace engineering at the University of Manchester. They earned a DPhil in Materials Science from the University of Oxford, focusing on residual stress in polycrystalline diamond cutting tools. Their current research explores residual stress in metal additive manufacturing and multi-pass welding, encompassing both experimental methods, such as neutron diffraction, and computational modelling predictions. To facilitate the computational work they have become involved in digital twinning of additive manufacturing cells.
Jinjiang Li *University of Manchester*	**Fluid–Structure Interaction of Bristled Insect Wings: Understanding Flow Blockage and Stiffness Effects** **Abstract:** This study investigates how bristle stiffness and spacing influence flow leakage in miniature insect wings operating at low Reynolds numbers (<40). Using OpenFPCI—an open-source Fluid–Structure Interaction solver coupling OpenFOAM-extend and

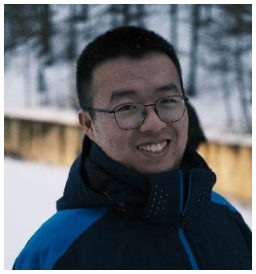

ParaFEM—we simulate seven parallel bristles attached to a common base within a wind-tunnel-like domain. The solver, packaged in a Docker container for flexible HPC deployment, enables reproducible and scalable simulations. Results show that bristle stiffness significantly affects aerodynamic performance. This work is the first to integrate bristle deformation into bristled-wing aerodynamics, advancing understanding of flow–structure interactions in tiny fliers and promoting open, collaborative HPC-based biomimetic research.

Bio: I am a PhD candidate in Mechanical Engineering at the University of Manchester. My work focuses on the application of digital twin technology in manufacturing and product design, exploring how advanced simulations and real-time data integration can enhance efficiency and innovation in these fields.

Jimin Lee

University of Liverpool

Lightweight Vision at the Edge: Efficient Ablation-Aware Training under Quantised Edge Constraints

Abstract: Deep convolutional neural networks (DCNNs) are highly performant on computer vision tasks, but their memory and compute requirements often exceed what is available on edge devices. Such deployments often prioritise a specialised subset of classes where classical fine-tuning often bears the risks of overfitting and poor generalisation. Ablation-aware techniques on the other hand have been shown on larger architectures to be effective at reducing this risk across longer training schedules however the effectiveness of such techniques post-quantisation, common for edge deployments, is much less well understood. In this work, we investigate the effects of ablation-aware techniques during fine-tuning of small architectures and data subsets as well as their impact on performance after post-training quantisation.

Bio: 2nd Year Computer Science @ University of Liverpool with research interests in efficient deep learning, model compression and Natural Language Processing.

See you in 2026…