



Science and
Technology
Facilities Council

Conference Proceedings
STFC-CONF-2025-002

CoSeC 2025 Mid-Year Forum Report

Sheffield University, UK
4th June 2025

MT Hoang

July 2025



©2025 UK Research and Innovation



This work is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

Enquiries concerning this report should be addressed to:

RAL Library
STFC Rutherford Appleton Laboratory
Harwell Oxford
Didcot
OX11 0QX

Tel: +44(0)1235 445577
email: library@stfc.ac.uk

Science and Technology Facilities Council reports are available online at:
<https://epubs.stfc.ac.uk>

Accessibility: a Microsoft Word version of this document (for use with assistive technology) may be available on request.

DOI: [10.5286/stfcconf.2025002](https://doi.org/10.5286/stfcconf.2025002)

ISSN 2753-5800

Neither the Council nor the Laboratory accept any responsibility for loss or damage arising from the use of information contained in any of their reports or in any communication about their tests or investigations.



**Computational Science Centre
for Research Communities**

Community Forum

Wednesday 4 June 2025

Sheffield University



Foreword	3
Introduction	4
Agenda.....	5
Summary	6
Forum Discussion	7
Making a Successful Research Community	9
Research Software Excellence	9
Community Impact	9
Conclusions and Next Steps.....	10
APPENDIX	12
Collaborative Computational Projects (CCP) Reports	12
Materials Science	12
Biological Science	20
Computational Engineering.....	29
Tomographic Imaging.....	39
Quantum Computing.....	46
New Community Reports	49
High End Computing Consortium (HEC) Reports	57



Foreword

The UK High Performance Computing community is widely admired internationally for its high degree of coordination and resulting impact. Collaborative Computational Projects have existed for more than 40 years, and such successful open collaboration remains essential to maintain, further strengthen and build on UK leadership at the cutting edge of current HPC research - particularly now when developments are moving at pace in AI and Quantum Computing as well as established digital simulation. For many years the CCPs have themselves been coordinated by a Steering Panel (SP) of HPC community leaders and research council representatives, and this has now expanded into a wider Forum embracing all UK existing and new HPC communities to ensure most effective sharing of information, skills and insights amongst these and right across UKRI.

Following the SP model, the Forum will meet every six months with recorded Minutes updated and distributed to participants in between. This first Forum Report attempts to capture key emerging points in a succinct fashion to inform a wider audience so that the collective impact of contributions from the gathered experts and input from CoSeC can be maximised.

The Sheffield Forum was particularly effective because it not only allowed community chairs to receive latest information direct from all the council representatives and directly feedback but also provided the opportunity to share experience and informed inputs for collaborative best practice, skilling and research software developments.

The Forum was held just ahead of the very welcome confirmation of the large new national investment in both, significantly upscaled, capability supercomputer and AI compute resources. Together these will allow the HPC communities to realise many of the opportunities they have already identified for integrating AI with digital simulation, but such advances will of course also require secured continuity of community recurrent funding.



Prof. Mark Savill, Chair of CoSeC Community Forum

Introduction

The CoSeC Community Forum is a bi-annual meeting involving the communities directly involved with the CoSeC programme, including the [Collaborative Computational Projects \(CCPs\)](#), the [High-End Computing Consortia \(HECs\)](#) plus representation from all of the [UKRI research councils](#). Invitations are also extended to CCPs and HECs that currently sit outside of the CoSeC programme but maintain links for collaborative purposes.

The forum is an opportunity for the [CoSeC communities](#) to come together to report on recent activities and future opportunities as well as to discuss national strategy related to computational research. The meeting provides networking possibilities and discussion activities that help to share information and best practice between the communities as well as identify potential for future cross-community projects.



This meeting was hosted by the Chair of [CCPBioSim](#), Prof. Sarah Harris, at the University of Sheffield where the state-of-the-art facilities in its Diamond Building were put to good use.

This meeting looked to build on previous Forums by emphasising the importance of strategic discussion in the context of the national Digital Research Infrastructure, as well as providing a platform for the six new communities recently [funded through CoSeC's call](#).

Agenda

Welcome	Prof. Mark Savill (Cranfield University)	09:45
Minutes (November 2024) and matters arising	Prof. Mark Savill (Cranfield University)	09:50
<ol style="list-style-type: none"> 1. CoSeC Update (30 minutes) 2. Research Object Cataloguing (15 minutes) 3. Grand Challenges (15 minutes) 	<ol style="list-style-type: none"> 1. Dr Stephen Longshaw (UKRI STFC) 2. Dr Gemma Poulter (UKRI STFC) 3. Alison Oliver (UKRI STFC) 	10:00
Break		11:00
UKRI DRI Update	Andrea Sharpe (UKRI DRI Council)	11:30
UKRI Council DRI Session	UKRI Council Representatives	11:45
Isambard-3 & Isambard-AI	Dr Richard Gilham (BriCS)	12:30
Lunch		13:00
Community Posters and Networking	ALL	14:00
New Community Introductions	New Community Chairs	14:30
Break		15:00
Collaborative computational research: context and potential	Prof. Giovanni Ciccotti	15:15
Introduction to the Forum Discussion	Damian Jones (UKRI STFC)	15:20
<p>Forum Discussion: <i>Broken into tables each with at least 1 new community Chair. Three discussion topics (15 minutes each) with 5 minutes wrap-up from each table.</i></p> <ol style="list-style-type: none"> 1. Making a successful research community 2. Research Software excellence 3. Community impact 	<p>ALL</p> <p>(see attached information for guidance)</p>	15:30
Any other business, wrap up and meeting close	Prof. Mark Savill (Cranfield University)	16:45

Summary

The CoSeC Communities Forum convened on 4 June 2025 at The University of Sheffield, with 44 attendees joining both in person and online. The meeting brought together representatives from Collaborative Computational Projects (CCPs), High-End Consortiums (HECs), and funding bodies to discuss key developments, challenges, and strategic priorities for advancing computational research in the UK.

Key updates:

CoSeC continues to drive impactful initiatives, including two cross-cutting projects: Research Object Cataloguing, aimed at consolidating research outputs, and Energy Efficient Computing, which addresses sustainability in computational practices. A workshop on energy efficient computing is scheduled for September 2025, and the newly launched Community Collaboration Visit Fund will support knowledge exchange at National Labs. The CoSeC Impact Award and Fellowship Programme were also announced, recognising contributions and fostering emerging talent across the community.

The **Grand Challenges** Initiative is gaining momentum with submissions from CCPs under review to identify shared themes and barriers. This initiative aims to develop a unified strategy aligned with national priorities supported by dedicated web pages to showcase community goals and advocate for sustained funding. The initiative represents a critical opportunity to position CCPs as essential contributors to the UK's strategic research agenda.

The **Digital Infrastructure Programme (DRI)** has made significant strides with investments in federated data services, software development and professional support over the past 18 months. A one-year spending review plan has been submitted, and we await further clarity on future funding allocations.

Discussions also addressed ongoing **resource challenges**, including the extension of ARCHER2 and reduction of Tier 2 investment to Isambard3. Members were encouraged to explore international opportunities such as EuroHPC, which has seen limited uptake from UK researchers. The continued importance of the community access model was also emphasized as a critical approach to supporting national computational research needs.

The forum welcomed six **new CoSeC communities**, each of which presented their visions and contributions through flash talks. These communities, focused on diverse areas, such as granular materials simulation, arts and humanities, particle physics, and numerical relativity are expected to tackle critical challenges within their respective fields and to broaden the impact of computational research across disciplines. They are funded directly through a [CoSeC funding opportunity](#) from UKRI's [Digital Research Infrastructure \(DRI\) programme](#) to develop and proliferate the Collaborative Computational Project (CCP) model across research communities within UK Research and Innovation (UKRI), creating a strong and stable landscape of communities to support the concept of research computing software as an infrastructure.

Forum Discussion

The discussion sessions were highly productive, focusing on three central themes critical to the activities of CoSeC and its communities: **building a successful research community**, **research software excellence**, and **community impact**. Attendees divided into smaller groups to share insights and propose actionable strategies, with each group presenting key aspects of their discussions. The groups were provided with the following guidance ahead of each discussion:

1. Making a successful research community

This discussion should consider – but not be limited to – factors that make CoSeC communities successful such as governance structures, EDI, communication strategies, engagement with community members and so on.

Existing and established communities:

- What works well for your community?
- Are there aspects of building your community that you would recommend to a new community?
- How do you tie long-term planning for your community software with its underlying research drivers?

New communities:

- What are the plans that you are considering incorporating into your model?
- What would help you achieve success in the long term?

2. Research software excellence

At a recent internal meeting of CoSeC project leads we discussed the use of STFC infrastructure within communities to enable excellence within their community research software efforts. This looked at offerings such as the STFC Cloud, Data Analysis as a Service (DAaaS) virtual environment for training and the “Anvil” multi-platform software testing service.

Through this discussion we recognised the different approaches to DevOps (Development Operations) taken within the communities and fundamentally recognise that no single solution will meet all of the needs of all communities. Software development can be deeply personal and choice of enabling technologies used is part of that.

The purpose of CoSeC within its communities is not to dictate how research software is created and maintained but instead to understand best practice, gather knowledge and to provide training where useful. As CoSeC is a direct part of STFC Scientific Computing, it can enable the use of services and platforms run by the department (some of these are listed above), these could be considered as direct replacements for commercial solutions (e.g. use of STFC Cloud in place of AWS, Microsoft or Google commercial Cloud) or as an additional capability (e.g. the Anvil Jenkins service to enable additional Continuous Integration capability beyond GitHub Actions).

Please discuss the infrastructure options that you currently use within your community and why, these may be available to you through your own institutions (or be generally free for use), or something you enable specifically through your community funding. Do they meet your requirements? What is missing?

3. Community impact

Historically CoSeC collected four metrics on an annual basis that were published as part of its annual report and, importantly, were specifically related to CoSeC activities. These were: publications related to key research software, presentations, training days and citations.

With the changes to the CoSeC model that were introduced in April 2024, it now has a different reporting remit, while reporting on the impact of the programme is still of utmost importance, it is now equally important that the research impact from across the communities are also captured and reported.

Recently, CoSeC has joined forces with the *Hidden REF* project (Research England; <https://hidden-ref.org>) to help it uncover the Non-Traditional Outputs (NTOs) and Hidden Roles that exist in computational research. This is key to ensuring the development of the Research Technical Professional (RTP) career path that CoSeC both champions and relies upon. The collective voice of the CCPs, HECs and other similar communities is a powerful part of the solution to ensuring research impact generated by computational research is better reflected in the REF, ultimately leading to a better reflection in the national research story.

Please discuss the way that your community considers research metrics, considering questions such as:

- How do you report these locally to your REF teams?
- What do you feel is valuable to report if you could send in whatever you chose
- Are publications sufficient to represent different research output types (e.g. software, datasets, training) or do these warrant their own dedicated recognition?
- Are existing methods of disseminating research outputs sufficient (e.g. arXiv, Zenodo etc.) or is more capability needed and if so, what?

Making a Successful Research Community

The forum emphasised that thriving research communities are built on shared goals, strong leadership and inclusive governance. Regular activities such as town halls, flash presentations and training sessions were highlighted as essential for fostering trust, collaboration and constructive feedback.

Equally important is creating a welcoming and inclusive culture where all members feel valued and respected, which is crucial for sustaining long-term engagement.

Participants shared practical suggestions, including enhancing communication through tailored tools like mailing lists, actively onboarding new members through training programmes and hosting low-barrier events to encourage participation and expand the community. Efforts like these ensure that communities remain vibrant, inclusive and responsive to the needs of their members

Research Software Excellence

Research Software Engineers (RSEs) play a vital role in computational research, and their contributions should be recognized through clear career paths and leadership opportunities.

Discussions also underscored the importance of quality training, curated datasets and accessible tools to support software excellence.

Participants also explored the need to address hardware challenges, such as limited GPU access and transitioning legacy codes to modern platforms to meet current demands.

Collaborative feedback was emphasized as critical in improving software usability and user engagement, ensuring that research software continues to evolve in line with the community's needs.

Community Impact

Participants explored how to better measure and communicate the value of CCPs. There was broad agreement on the need to move beyond traditional, quantity-based metrics such as download counts and paper citations toward more meaningful measures, including case studies and success stories that highlight real world impact.

Challenges in linking research outputs such as publications and datasets to specific software tools were noted, as were difficulties in standardizing metrics across disciplines. Recognising non-traditional contributions through initiatives like Hidden REF was seen as an important step forward.

Communicating success to stakeholders, including funders and policy makers, remains critical for increasing visibility and demonstrating the value of computational research to national priorities.

Conclusions and Next Steps

This event saw a change in the format of the meeting, moving more towards the idea of the forum being an opportunity for discussion and collaborative knowledge exchange rather than a way to provide detailed overview of the technical work being done within each community. In part this is because that is now the job of this new report that will be published following each forum, but also it reflects the fact that there are now 25 communities involved as well as representatives from across the UKRI Digital Research Infrastructure (DRI) landscape. Between them the breadth of domain-specific research covered is simply more than can be reasonably covered in a day.

The forum has evolved from the CCP Steering Panel meeting into something more wide-reaching that considers not only the amazing research happening within the CCPs and HECs but also the strategic aspects of what research communities can do to help enable a stable and effective DRI. This meeting saw effective dissemination of the upcoming DRI landscape from the perspective of UKRI into the communities and a strong set of discussions about the best way to strategically position the communities to maximise their value for UK research. A highlight was during the round-table discussion session where a mixture of new and established communities Chairs as well as UKRI representatives were able to discuss some fundamental ideas about what a community is and what good looks like in terms of research software development. Outcomes from these discussions can be found earlier in this report and offer a useful insight into current thinking and where things can go next for improvement.

During my opening presentation about CoSeC, the idea of how collaborative research communities are an irrefutable necessity in the context of AI and its applications was explored. The example of how the Protein Data Bank (PDB) ultimately led to the creation of the AlphaFold system was given with the fact highlighted that, without communities like CCP4 or CCP-EM, something like PDB just wouldn't be possible. At a time where AI is seen as a major driver for the UK's economy and forms a significant part of its industrial strategy¹, it is easy to forget that the modelling, simulation and data curation that are fundamental drivers for communities like the CCPs and HECs aren't just a nice thing to have, they are fundamental. It is easy also to forget that, like the AI tools and methods built upon these foundations, the work is not done, the techniques, methods and software needed to create the outputs that AI relies upon are an ever-changing part of research and just as important to support (and fund) as anything else. This was discussed and unanimously agreed upon during the meeting and will form a major part of CoSeC's messaging moving forwards.

The new national supercomputer, Isambard-AI, was ably presented by Dr Richard Gilham from the Bristol Centre for Supercomputing (BriCS) alongside the existing Tier-2 Isambard-3 service. At the time of Richard's presentation, it was uncertain what else would complement these brilliant services in the national computing landscape, and this formed a major part of the discussion in the room given the apparent "cliff-edge" that UK computing infrastructure faced. However, as I write this, it has now been confirmed by UK government as part of its Spending Review that there would be a £750m investment² into the Edinburgh Parallel Computing Centre (EPCC) to deliver the next major resource to sit alongside Isambard-AI. This welcome news puts

¹ <https://www.gov.uk/government/publications/industrial-strategy>

² <https://www.gov.uk/government/news/scotland-to-host-uks-national-supercomputer-as-chancellor-confirms-750-million-investment>

a very different spin on the prospects of computational research in the UK and provides a renewed hope that the communities can continue to be the engine driving the creation of the software and data needed to make the rest of the investments being made in AI practically useful.

The communities present at this forum fell into 1 of 3 categories. They were either a new community with funding until December 2026, an established CCP with funding until October 2026 or a High-end Computing Consortia with funding until December 2026. In effect, all collaborative community funding across the UKRI landscape currently has a deadline of Q4 2026. This was clearly a keen topic of discussion as one of the cornerstones of the success of the CCP and HEC models is stability in funding. All present from UKRI took away the clear message that by the next forum meeting there needs to be some clarity on what comes next and the timelines under which it will be made available.

The coming 6 months will be an exciting time for CoSeC, we have a number of opportunities such as the next round of the CoSeC Fellowship programme, an open call for nominations for our Impact award, an upcoming call for this year's CoSeC Conference as well as an ongoing call for community visit funding. Alongside this, our technical work continues across the 25 communities directly within the programme where there are a number of exciting upcoming events.

In September there will be a workshop held at the Rutherford Appleton Laboratory and organised by the STFC Scientific Computing Computational Mathematics theme around developing sustainable green research software. Work also continues on the development of a "CoSeC Catalogue", a prototype design of this was presented at this forum and by the next forum the intention is for there to be significant progress with a view to a working version hosted on our website and ready to begin population by the communities. We also continue work on the CoSeC Challenges piece, following a presentation at this forum by our Research Community Manager, Alison Oliver, this will now move to the next stage with a view to presenting a complete prototype by the time of the next forum. I strongly believe this will provide an invaluable resource to allow a snapshot of what computational research challenges look like from a cross-UKRI view, access to this kind of knowledge will be invaluable when shaping future strategic priorities for the UK's DRI.



Dr Stephen Longshaw, Director of CoSeC

APPENDIX

These reports are reproduced as provided by community Chairs in a template as a response to a request for information from the CoSeC Programme Office.

Collaborative Computational Projects (CCP) Reports

Materials Science

Report from CCP-NC for the Period 01/10/2024 to 31/03/2025

Dr. J. Kane Shenton (UKRI-STFC, CoSeC Project Lead)

Prof. Paul Hodgkinson (Durham University, CCP-NC Chair)

1. Background

Nuclear Magnetic Resonance (NMR) is a useful technique to determine chemical structure, especially in compounds of which it is hard to produce single crystals big enough for diffraction techniques, as commonly found in organic molecules. NMR Crystallography is the technique of using quantum-mechanical simulations to predict NMR spectra to a high degree of precision and combining this with experiment to open new ways of exploring structure in not yet understood crystals, such as new pharmaceuticals. CCP-NC has the objective of disseminating and promoting this approach throughout the experimental NMR community in the UK and worldwide.

2. Highlights for the current reporting period

In January 2025 we organised an Advanced Materials Search Workshop at the Cambridge Crystallographic Data Centre, bringing together domain experts from CCP-NC, CCDC, the Physical Sciences Data Infrastructure (PSDI) and OPTIMADE to discuss the best way to enable features like substructure searching in the upcoming version of the CCP-NC database. The use cases were refined and a set of action points drawn up to enable these use cases to be met.

Advancing the FAIR agenda, we continue to develop both our legacy database (completing requested improvements to the user-interface and the accessibility of the service) and our upcoming version which is based on the fully FAIR-compliant Nomad Oasis software stack. We have started implementing the key components, including the NMR schema and file parsers, collaborating closely with the Nomad developers, the PSDI and the CASTEP and Quantum Espresso code developers.

Version 0.9.2 of Soprano – our core Python library – was released in October 2024. This release came with a major new way to use Soprano: a robust command-line interface (CLI). This CLI has several tools that the NMR community has requested, including plotting 2D NMR spectra and extracting and summarising NMR parameters magres files without having to write any code. We have begun adding more tools to this, to enable additional steps to standard NMR workflows, including interfacing with spin simulation tools such as Simpson.

3. Overview of work relating to accelerated computing

CASTEP's NMR functionality is the key software for our user community, since it allows experimental NMR parameters for a given crystal structure to be calculated from first principles using Density Functional Theory (DFT). Although some other codes, most notably Quantum Espresso, now also have this functionality, CASTEP remains the most well-developed code and dominates use in the UK. Hence it is critical for the NMR crystallography community that CASTEP, and CASTEP-NMR in particular, can work effectively with the latest hardware, especially GPU acceleration. Using specialist RSE support (starting immediately after this period) we plan to port CASTEP's NMR functionality to run on GPUs.

GPU acceleration for CASTEP-NMR will be developed in the wider context of support for heterogeneous computing: adapting codes and runtime systems to exploit diverse computing environments; and training users who will not be experts in heterogeneous computing. Without support, there is the risk that users will run demanding calculations very inefficiently. The GPU port will be extensively tested, along with appropriate documentation on how to get the most out of the available hardware.

4. Current plans, developments, or specific applications of AI

The CCP-NC has two main plans around becoming AI-ready:

1. Machine-learned interatomic potentials (MLIPs) offer an efficient way to bridge the gap between DFT's feasible length- and timescales and real-world problems. Although pre-trained MLIPs for materials have been demonstrated to be effective, our community will need training and guidance to use them effectively. We have begun carrying out a series of three carefully documented case-studies that are directly relevant to our community, showcasing the potential for these methods.
2. ML is also important for calculating NMR observables from molecular dynamics snapshots, a task generally impractical with DFT. [ShiftML's](#) successful prediction of chemical shifts has already gained traction among non-computational researchers, and rapid advances are expected. However, constructing suitable training databases for ML—structurally diverse but representative of target systems—requires significant expertise. Furthermore, learning full NMR tensors, rather than simply the trace of them will be essential to obtain time-averaged NMR properties. We will approach this challenge by 1) using the existing ShiftML dataset to explore the potential to learn full NMR tensors and 2) developing an on-the-fly learning approach directly within CASTEP (building on existing on-the-fly learning of forces and energies). These projects are planned to take place in the coming financial year.

To further these goals and engage with the community, we are holding a workshop in Manchester (13-14th May 2025) to bring together NMR crystallography's problems and machine-learning's solutions.

5. Up to 5 publications that create an impact story for your community from the reporting period.

It is difficult to choose 5 papers out of the large number of high-quality papers that cited CCP-NC support from this period. We can, however, highlight the collection of 22 papers which formed the basis of the [Faraday Discussion in NMR Crystallography](#). This discussion meeting was developed over a couple of years by CCP-NC and involved a wide range of international

participants presenting work at the forefront of NMR crystallography. We expect that the papers in this volume to have a significant impact in the field. In particular, we note the number of papers in this volume and more widely that involve applications to new battery and energy storage materials; experimental NMR combined with computation is one of the most effective tools for understanding these complex materials.

6. Workshops and new opportunities

Date	Title	Description
13-14 May 2025	NMR-ML Discussion meeting	Discussion workshop including panel sessions for members of NMR and ML communities. (~50 participants)
TBC	CCP-NC Townhall	Disseminate activities and vision of CCP-NC and gather community feedback.
TBC	CCP-NC Online: running NMR calculations efficiently	Part of our lunchtime discussion seminar series. This will focus on how to run CASTEP GIPAW calculations efficiently on HPC systems.

7. Issues and problems

There has been some impact of the continued restrictions on core staff support, but we have responded creatively to the evolving circumstances and have continued to make steady progress on all our key objectives.

Report from CCP5 for the Period 01/10/2024 to 31/03/2025

Prof. Paola Carbone (University of Manchester, CCP5 Chair)

1. Background

CCP5 is the Collaborative Computational Project for computer simulation of condensed phase materials at length scales spanning from atomistic to mesoscopic levels. Founded more than 40 years ago, CCP5 has promoted the involvement of UK scientists in collaborative research achieved via software and methodology development, training, networking and outreach. It provides support for all UK scientists engaged in developing, applying and exploiting computer simulation methods for condensed matter systems. CCP5 has over 1500 UK and international members, which comprise research active academic faculty staff in 35 different UK universities and at least 18 other UK industrial, charitable or government organisations. A distinctive feature of CCP5 is its successful strategy of developing and disseminating new codes and methods for all kinds of materials problems. These include solid-state materials, polymers, colloidal solutions, liquids and mixtures, liquid crystals, surfaces and interfaces, homogeneous and heterogeneous catalysts, mineral, bio-mineral, organic and bio-molecular systems.

2. Highlights for the current reporting period

CCP5 Summer school (<https://summer2024.ccp5.ac.uk/>) took place between 14-25 July 2024 at University of Newcastle. 75 students from UK and overseas took part these years with a roughly 50/50 split, with the first two days being dedicated to computer programming classes, Fortran and Python, which students choosing one of them. Basic lectures in theory of modelling

were given, with basic molecular dynamics, monte carlo and free energy methods being covered in the mornings with practicals in the afternoon. Last three days of the school were dedicated to advanced lectures students being able to choose between First Principles Simulations, Biomolecular Simulations, Mesoscale and Machine Learnt Interatomic Potentials. Despite being introduced only few years ago Machine Learnt interatomic potentials is by far the most popular with 29 students taking it.

Annual general meeting happened this year at University of Sheffield between 4 – 6 September with more than 40 participants. It showcased a series of talks ranging from all established academics to early career researchers.

CCP5 Prize was awarded to Dr Marcus Campbell-Banerman University of Aberdeen.

The CCP5 Summer Bursaries as usually is very popular and oversubscribed with 5 bursaries being awarded this summer.

CCP5-CECAM sandpit had its second awardees this year to Dr Xioacheng Shang, University of Birmingham and Prof Gabriel Stoltz École Nationale Des Ponts et Chaussées, Paris France.

CaMML School is a machine learning for materials training course run by the Physical Sciences Data Infrastructure (PSDI) initiative in collaboration with Alchemy, with support from STFC-SCD, PSDS, CCP5 and CCP9 as a follow up to the very popular 2023 Machine learning for Atomistic Modelling Autumn School. This training is targeted towards PhD students, in particular those in the Materials and Molecular Simulations field, who have experience of coding but are not highly experienced with machine learning. The aim of this training is to introduce attendees to the latest methods of machine learning for the atomistic simulation of materials. It runs between 31/03-04/04 2025 at Daresbury Laboratory and had 43 students from all over the world.

Undergraduate bursary programme was run with five bursaries being awarded.

3. Overview of work relating to accelerated computing

There is no GPU development at the moment. Only incidental work can be reported here with respect to GPU work which came out in context of MLIPs. Benchmarking of MACE on multiple GPUs and variety of systems was undertaken. Results show MACE is competitive for system sizes of interest to majority of material scientists: tens of thousands of atoms with speeds of close to 1 ns per day.

4. Current plans, developments, or specific applications of AI

Up to a quarter of a page providing specifics around plans or direct developments and applications of AI techniques and workflows.

Machine Learning Interatomic potentials have shown great promise in achieving DFT accuracy and beyond and accessing timescales well beyond the ground truth method. In addition to benchmarking activities, reported above. The work on this work package concentrated on fine tuning strategies for metal organic frameworks and on suitability of foundation models to be used out of the box for zeolites. Results indicate that with relatively small effort, hundreds of new structures, carefully selected one can improve the dramatically accuracy of foundation models. Two papers are published or submitted.

Workflows: DL_FIELD workflow capabilities have been broadened to include translation of force field (FF) data to produce read-to-run input and data files for LAMMPS simulation package. In addition, model building features and model customisation such as solvation, molecular duplications and introduction of bond constrains are also made available for LAMMPS files. A new force field scheme has also been implemented, namely the TraPPE united atom model. This also involves program restructuring, to ensure the FF scheme is made available for all translations that are currently implemented in DL_FIELD: DL_POLY, GROMACS and LAMMPS, seamlessly integrated as one of the FF component schemes within DL_FIELD framework.

Benchmark simulations between DL_POLY, GROMACS and LAMMPS have been performed for systems of argon, LiCl, pentane, rigid and flexible water to determine input parameters for reproducible single point energies between software. This is in preparation for seamless tutorial material over different MD packages. Comprehensive literature review on packages available to build polymer simulation workflows, including input formats (e.g. SMILES, bigSMILES etc.) and efficient box building, in preparation for their implementation with DL_field.

5. Up to 5 publications that create an impact story for your community from the reporting period

The publications chosen should tell a story about the work within the community, they may highlight key topics or outputs (e.g. software or highly impactful research), or together they may form a single cohesive message about a key challenge being tackled.

Elena, A.M., Kamath, P.D., Jaffrelot Inizan, T. et al. Machine learned potential for high-throughput phonon calculations of metal—organic frameworks. *npj Comput Mater* 11, 125 (2025). <https://doi.org/10.1038/s41524-025-01611-8>

Modelling Silica using MACE-MP-0 Machine Learnt Interatomic Potentials

Jamal Abdul Nasir, Jingcheng Guan, Woongkyu Jee, Scott M. Woodley, Alexey A. Sokol, C. Richard A. Catlow, Alin Marin Elena, <https://arxiv.org/abs/2411.00436>

6. Workshops and new opportunities

Community Events sponsored by CCP5:

- 33rd European Symposium on Applied Thermodynamics (ESAT 2024) Poster session, Edinburgh, 9-12 June,
- Theory and Simulation of Soft Matter: from Virial Coefficients to Liquid Crystals, 28 November 2024, University of Manchester
- Recent Appointees in Physical Chemistry meeting, University of Edinburgh, 13-15 January 2025.
- CCP5 Annual General Meeting, 4-6 September 2024, University of Sheffield.
- Computational Molecular Science 2024, Warwick, 11-13 September 2024
- Molecular Modelling in Clay Science: from constructing clay models to performing molecular dynamics simulations. Edinburgh, 16th -17th of May 2024

- Synergies Between Mathematics, Data Science, and Molecular Simulations in Materials Science, Birmingham, 3-5 July 2024
- Adsorption Summer School (2024 edition), University of Strathclyde, Glasgow, 11-14 June 2024

7. Issues and problems

Nothing to report

Report from CCP9 for the Period 01/10/2024 to 31/03/2025

Prof. Stewart Clark (Durham University, CCP9 Chair)

1. Background

CCP9 “Computational Electronic Structure of Condensed Matter” was established in 1980 in order to advance method and code development in the UK first principles electronic structure community. CCP9 performs a key role in developing new theoretical and computational methods, in maintaining and developing software for the materials science community, and in training successive generations of researchers and academics.

CCP9 has pioneered the development of a range of materials simulation codes that underpin research in a wide variety of thematic areas and are crucial for the work of many academic and industrial research teams across the UK and internationally. CCP9 supports the large code projects CASTEP, ONETEP, CRYSTAL, Wannier90, CASINO, and QUESTAAL, each of which has hundreds or thousands of users and each addressing different needs in the materials modelling space, spanning, for example, the range from extreme accuracy for small systems to the ability to treat extended systems at lower accuracy. These codes provide efficient, precise and usable implementations of a range of theoretical methods suitable for different scientific problems; these codes are carefully designed and optimised to leverage HPC to maximise our modelling capability.

CCP9 also crucially provides a network which connects UK research groups in electronic structure, coordinating both research and software engineering activities and organising training in theory and application of the codes. Participation in the network continues to grow as electronic structure methods play an ever-growing role across materials science, computational chemistry including catalysis and electro-chemistry, and simulation of biological systems. The current network involves 70 UK academic groups.

2. Highlights for the current reporting period

The CCP9/CECAM Electronic Structure Graduate School, hosted at Daresbury Laboratory 24-28 February 2025, was a big success with over 30 participants from the UK and Europe coming together for this advanced school in electronic structure theory and application. Lectures and hands-on practical sessions from the developers of the CASTEP, CRYSTAL and QUESTAAL codes covered a range of methods suitable for different topics and ensured that the next generation of electronic structure PhDs and early postdocs are conversant in today’s state-of-the-art methods.

Advances in modelling defects and interfaces workshop. The workshop, held at the Institute of Physics, London, was organised to honour the vast catalogue of achievements of Professor Alexander Shluger to modelling defects and interfaces in solids and nanosystems; his many contributions to computational materials science

(<https://thomasyoungcentre.org/event/advances-in-modelling-defects-in-solids-workshop/>)

Developments in CASINO by John Trail include a new implementation of forces which is now part of the main code base and is being tested for more complex molecules in comparison with alternative high-accurate methods (CCSD); John Trail has also developed and a new many-body pseudopotential technique that promises to improve calculation accuracy considerably.

XML I/O functionality has been added to ONETEP linear scaling DFT code by Manuel dos

Santos Dias. The main objective is to develop an improved restarting of aborted or stopped calculations, to increase the user friendliness of complex workflows for materials properties calculations. This is well advanced with major code developments already merged into the master repository.

GPU porting of ONETEP functionality has begun: the main objective is to speed up the band structure unfolding with spin-orbit coupling by moving its computation to GPUs. The initial scoping stage has been completed, and the implementation stage will start from April 2025 onwards.

Developments in CRYSTAL by Barry G Searle include energy saving optimisations undertaken as part of the new DRI work programme.

Calculation exceeding 1M basis functions using CRYSTAL on archer2. The calculation, performed by Barry G Searle, was of a mesoporous silicate comprising 70,000 atoms and ran for ~6h on 4096 nodes employing the EPLA massively parallel eigensolver.

Compton scattering calculations using QSGW with the QUESTAAL code have been compared with DFT and GW in elk in a collaboration between Jerome Jackson and Dugdale group in Bristol.

CASTEP 25.1.1 has been released for academic users; this is an important code in the CCP9 ecosystem as demonstrated by academic group licences for CASTEP passing the milestone of 2000 active user groups worldwide at the end of 2024.

Ongoing Collaborations in Materials Science continue for the CCP9 support team including various projects in magnetism by Manuel dos Santos Dias whose collaborations with Juelich and Diamond include a joint experimental-theoretical project using the I21 beamline at Diamond to study altermagnetic magnons in CrSb using RIXS. The measurements took place in November, and the first results of this collaboration are available as a preprint (<https://doi.org/10.48550/arXiv.2503.02533>) and a paper on magnetic anisotropy energy of oxidised adatoms: *Phys. Rev. B* **110**, 224409 (2024). Jerome Jackson's collaborations with Georgetown, Case Western and NREL in the USA resulting in a paper on the Electronic structure and exchange interactions in altermagnetic MnGeP2 (<https://doi.org/10.48550/arXiv.2505.02934>).

3. Overview of work relating to accelerated computing

GPU optimisation of CASTEP is being conducted by two DRI-funded postdocs at the Durham University and the University of York, working to accelerate the response function code that underlies many spectroscopy calculations. GPU developments are ongoing in ONETEP, led by Manuel dos Santos Dias at STFC.

The CCP9 post-doc Patrick Williams (Durham) is in collaboration with CCP-NC (Ben Durham, Durham) and UKCP-HEC (Matt Smith, York), working in a team for GPUing Castep perturbation methods. Transitioning from openACC to OMP.

4. Current plans, developments, or specific applications of AI

Neural Network DFT (NN-DFT) is a combination of neural networks and DFT that extends the variational principle of physics with the loss minimisation rule of deep learning. NN-DFT is a Physics-Informed Neural Network (PINN), as the NN learns the physics in the DFT calculations in an unsupervised way, accelerating the energy minimisation. On the other hand, Reinforcement Learning is a deep learning approach used in particular in the optimal control problems for logistic companies: the NNs learn the optimal and fastest way in a variational approach using a reward and punishment approach. Paolo Trevisanutto, part of the CCP9 team at STFC and together with colleagues from STFC's Machine Learning for Science group is implementing a reinforcement learning algorithm in NN-DFT to speed up the variational energy minimisation.

5. Workshops and new opportunities

CECAM/ALC/CCP9 Spectroscopy Masterclass, 1-5 September 2025

This event is a 5-day in-person school in theoretical spectroscopy at Rutherford Appleton Laboratory, 1-5 September 2025. The school will teach the calculation of different spectroscopies including core (such as XAS, XMCD), valence band (ARPES) and magnetic excitations (as probed by INS) using first-principles electronic structure codes: SPR-KKR, MsSpec and Questaal. Lectures will provide an introduction to the theory as well as a high-level overview of the current state of the art in spectroscopy modelling. Afternoons will be dedicated to practical, hands-on teaching sessions designed for newcomers to electronic structure theory. The aim of the practical sessions is to enable experimentalists to carry out calculations independently. The final afternoon is dedicated to open discussion on how best to provide theory support at the large-scale facilities. The event is funded as a CECAM Flagship School, with sponsorship from the Ada Lovelace Centre and the UK's CCP9 network.

For more information, see <https://www.cecamlc.org/workshop-details/cecamlc-ccp9-spectroscopy-masterclass-1439> or email scd_events@stfc.ac.uk Please register here: <https://cvent.me/Ov28O3> (deadline 25 July 2025).

Lennard-Jones Centre-CECAM Meeting 2025 in Cambridge

The Lennard-Jones Centre-CECAM Meeting 2025 will place from 1st–5th September 2025 at the Ray Dolby Centre, Cavendish Laboratory, University of Cambridge. This event will celebrate advances in electronic structure and Density Functional Theory, aligning with the 70th anniversary of the Theory of Condensed Matter Group in Cambridge and the 150th anniversary of the Cavendish Laboratory. The meeting will serve as a tribute to the significant contributions

of Cambridge scientists such as Professors Mike Payne, Volker Heine, and Richard Needs, whose work has shaped generations of researchers both within the UK and internationally.

The meeting will feature invited session leads presenting key topics in the field, accompanied by early career researchers and contributed talks, ensuring a dynamic and forward-looking discussion. A historical session will concentrate on the legacy of UK researchers, while an industry half-day will focus on commercialisation and industry collaboration.

For more details and registration, please visit: <https://ljc.group.cam.ac.uk/dft-2025>

7. Issues and problems

Nothing to report

Biological Science

Report from CCPBioSim for the Period 01/10/2024 to 31/03/2025

Dr. James Gebbie-Rayet (UKRI-STFC, CoSeC Project Lead)

Prof. Sarah Harris (University of Sheffield, CCPBioSim Chair)

1. Background

CCPBioSim is the Collaborative Computational Project in biomolecular simulation at the life sciences interface, bringing together chemists, physicists and chemical engineers as well as researchers from all branches of "molecule-oriented" biochemistry and biology. Simulations help to analyse how enzymes catalyse biochemical reactions, and how proteins adopt their functional structures e.g. within cell membranes. They contribute to the design of drugs and catalysts, and in understanding the molecular basis of disease. Our aim is to involve experimentalists and computational specialists in this work, sharing the belief that the best science can be done when theory and experiment are closely integrated. CCPBioSim engages with early career researchers and non-experts through the provision of tutorials and workshops enabling them to become proficient and productive users of biomolecular simulation techniques. We are also actively engaged in developing new advanced methods, which in future will be used by our community to deliver new and exciting science.

2. Highlights for the current reporting period

A complete overhaul of the cloud-based CCPBioSim training platform has been delivered. The CCPBioSim training catalogue contains 22 courses across range of topics from beginner through to advanced material. Each of which has its own software environment built as a docker container. This meant that as the portfolio grew the workload to maintain and keep them cutting edge became significant, another major issue is that these containers only worked with our cloud system and not on users own computer. A new platform has been delivered that has introduced modern gitops lead CI principles to establish an automated multi-architecture build system with automated test pipelines that can deliver constantly automatically updating software environments in a way that the deployed version is always bleeding edge. The

automated testing means that when updates to software from around the world breaks core teaching components, we can respond to it immediately and on an ad-hoc basis rather than having to do this for all training courses every 6 months. This means training delegates are always exposed to the bleeding edge in methods and software, and we can start to grow our training portfolio significantly further without the extra maintenance workload. A dashboard with links to all of the repositories is here <https://github.com/jimboid/biosim-k8s>.

A current focus on software development for novel methods is in the development of CodeEntropy, a novel code for the calculation of structural and solution-based entropy from molecular dynamics trajectories. Work this year has focussed on bringing the initial fusion of two separate codes (for structure and solution entropies) into a single united and cleanly written codebase. A significant refactor of the code has been complete with the adoption of a single purpose single output class-based approach for organising the code. Progress has been made on the development and addition of new scientific theory for general purpose orientational entropy calculation. Significant improvements have been made to the code readability, code documentation and also to the parameterisation and output of the code. The improvements to input/output are with clean readability and interoperability and the later adoption into FAIR databases such as the BioSimDB a programme in the UKRI funded PSDI. Significant improvements to the software engineering practices in the development of the code have been introduced, by the addition of CI pipelines for automated testing, documentation building and code quality correction. We intend to launch this code at our CECAM flagship workshop on entropy in Vienna in August 2025. Code can be found here <https://github.com/CCPBioSim/CodeEntropy>

An MDAnalysis workshop was held at UCL on 10 May 2024 (jointly sponsored by the Thomas Young Centre) with 30 participants.

The CCPBioSim Annual Conference was held in Newcastle 1 – 3 July 2024. There were 140 participants.

CCP5 Summer School held in Newcastle 14-25th July 2024. CCPBioSim contributed training materials and lectures for the advanced biomolecular simulation part of the summer school. There were 70 participants.

CCPBioSim co-organised a mesoscale modelling conference in Trento with CECAM 26th-29th August 2024:

<https://www.cecaml.org/workshop-details/biomolecular-simulations-at-the-mesoscale-1330>

The Training Week was held in Sheffield 2 – 6th September 2024. There were 42 participants in person and online. The week had 6 workshops: Software Skills (including managing environments, code quality and testing); Reproducible simulations with Aiida; Introduction to ChemShell for QM/MM simulations; Introduction to MDANSE; Intermediate MDAnalysis; and Introduction to OpenMM. Material from the software workshop can be found here: https://ccpbiosim.github.io/software_workshop/.

Multiscale conference held in Manchester 31 March until 2nd April. A conference jointly organised with CCP5 covering simulations of a multiscale nature at the materials-biomolecular simulation interface. There were 100 participants.

GW2 Genomics to whole cell workshop held in Sheffield 23 - 25th October.

In this reporting period we ran 11 Industry seminars held online around once a month and are very popular with around 80-100 people attending each talk. This year we have had talks from AstraZenica, Sygnature, GSK, Quantum Bio Inc, Acellera, Cresset and Schrodinger as well as a number of special seminars from experimental facilities and software organisations. These talks covered a wide array of topics spanning drug discovery, molecular simulation and HPC, AI and Quantum computing.

Our online cloud training gateway has available 22 courses, and collectively within the reporting period have delivered 5,127 training sessions across the world with around 2,500 of these in the UK.

CCPBioSim along with CCPN, CCP4 and CCPEM have jointly been successful with a bid into the UKRI DRI funded community extension to CoSeC to form a collaborative project for Integrative Molecular Biology (DRIIMB). CCPBioSim have appointed two staff, Hima-Bindu Kolli to work jointly between CCPBioSim and CCPN in November 2024 and Robert Welch to work jointly between CCPBioSim, CCP4 and CCPEM in April 2025. Work in this period has largely been at the CCPBioSim-CCPN interface due to the second appointment not happening until April. Bindu has been focussing her efforts around training and familiarity with the NMR software tools and methods, data standards and formats from CCPN and those from CCPBioSim as well as training around the FAIR data tooling developed for biomolecular simulation in the PSDI. Initial work has identified exemplar data sets and workflows for simulations conducted on NMR obtained structures, these will be used to inform the development of future tooling for the automation of biomolecular simulations for the purpose of gaining a dynamics-based understanding of NMR structure data.

3. Overview of work relating to accelerated computing

CCPBioSim as a consortium does not conduct directly activities that purely focus upon heterogeneous computing. However, the community has strong direct connections to the HECBioSim consortium which focuses on access to current UK HPC resources and optimising and benchmarking codes for such platforms, and the Exabiosim initiative which is focused on next generation simulation capabilities under the exascale HPC regime. CCPBioSim are obviously an important influence in the activities of both of these initiatives since its members are the underlying driver as consumers of existing and future UK HPC. Our community as a whole have been engaging in the porting to and further optimisation of our codes for around a decade already so have significant experience in this area. Our community plan to engage with the next generation platforms such as ISAMBARD-3 and DAWN for applications around AI driven molecular simulations. We will continue to develop our new methodologies with these HPC capabilities from the concept stage. As part of the UKRI-DRI funded DRIIMB we plan to share our expertise in HPC across the communities (CCPBioSim, CCP4, CCPEM and CCPN) that have come together to form the combined community.

4. Workshops and new opportunities

A list of upcoming workshops and meetings (up to 12 months in advance) for addition to the publicly visible CoSeC event planner.

Where applicable, highlight events of potential interest for other communities and opportunities for cross community working.

- CCPBioSim Annual Conference – Southampton 14-16th July 2025
- CCPBioSim will contribute to the CCP5 Sumer School July 20-31st 2025
- CECAM Flagship Workshop “Entropy of Soft Matter” – Vienna 27-29th August 2025
- CCPBioSim Training Week – Sheffield TBA October 2025

5. Issues and problems

Lack of clarity on whether there will be future funding calls makes it difficult to focus on long term planning.

Report from CCP4 for the Period 01/10/2024 to 31/03/2025

Prof. Ivo Tews (University of Southampton, CCP4 Chair)

1. Background

CCP4 provides and supports an integrated suite of programs for determination of macromolecular structures by X-ray crystallography. CCP4

- develops cutting edge approaches to experimental determination and analysis of macromolecular structures as a community-based resource;
- supports the development and integration of novel software into one suite;
- serves the widest possible research community, embracing academic (not for profit) and industrial (for-profit) research;
- offers education and training to scientists in experimental structural biology and encourages the wide dissemination of new ideas, techniques, and practice.

2. Highlights for the current reporting period

The CCP4 £2m BBSRC started 1.8.2024 and includes four work packages:

- **WP1** A statistical framework for analysing structural change and its representation in data (time and/or state); Hough, Evans, Winter, Diamond Light Source; PDRA Rachel Tang has started 30/09/2024.
- **WP2** Joint refinement of related structures; Murshudov, LMB Cambridge; Agirre, York; Tews, Southampton; Orville, Southampton, Diamond Light Source; PDRA Martin Malý has started 22/11/2024.

- **WP3** Advance computational methods for data processing and structure solutions to exploit advantages in electron diffraction; Krissinel, Waterman; Research Complex at Harwell; PDRA tba.
- **WP4** Exploiting Deep Learning-based structural bioinformatics for crystallography; Rigden; Liverpool; PDRA Aderik Verspooels has started 16/09/2024.

The collaboration contract between the seven participating institutions was signed 14/02/2025. The grant team consisting of all hired PDRAs and all grant PIs met at a first kick-off meeting 27/11/2024. To foster collaboration, a slack channel was created in January 2025. The grant team meet twice in person at the CCP4 study weekend (January) and at the Developers meeting (July), where they also report to the CCP4 executive.

Following the CoSeC communities forum 23rd May 2024 and the CoSeC townhall in Manchester 1st/2nd July 2024, CCP4 responded to the 2024 CoSeC CCP Bridging Call together with CCPBioSim, CCPEM, and CCPN. The proposal on “Digital Research Infrastructure for Integrative Molecular Biology (DRIIMB)” was awarded in October 2024; the contract is now ready for signature. Several posts have been associated with this proposal. See also sections 3 and 4.

3. Overview of work relating to accelerated computing

The CoSeC post (Daniel Celis Garza) associated with WP4 in the DRIIMB CoSeC bridging grant is a pilot in heterogeneous computing. The post explores how Molecular Replacement (MR), the dominant method of Macromolecular Crystallography structure determination within CCP4, can use the large database of AI generated structural templates (>600 million). The availability of such large template databases requires faster computational search tools and GPU acceleration. As well as effectively managing resource, acceleration can improve on the success rate by more accurate model positioning and faster exploration of search domain conformations and MR scenarios. The post explores how to reformulate MR algorithms for massive parallelisation on GPU devices, targeting a solution where translational and conformational searches are performed simultaneously across all available GPU threads. We recognise the CCP-EM community face equivalent challenges when fitting atomistic information into cryo-EM reconstructions. Given the larger maps from multi-component complexes potential benefits of GPU-enabled speed-ups are even greater, which will be explored in collaboration between the two CCPs.

4. Current plans, developments, or specific applications of AI

Two of the work packages of the BBSRC funded CCP4 grant explore AI:

- **WP3** Methods for electron diffraction data for the refinement of macromolecular structures, using machine learning approaches to filter out dynamical scattering, and procedures for taking crystal defects and inelastic scattering terms into account.
- **WP4** Exploiting Deep Learning-based structural bioinformatics: use of covariance-based distance and contact predictions to validate protein and nucleic acid structures; development of rational editing of RNA models for molecular replacement.

Additionally, in the DRIIMB CoSeC bridging grant we will explore the potential of AI/ML-based assessments of candidate Molecular Replacement solutions to outperform the conventional scoring methods based on root mean squared deviations (RMSD) template and target datasets.

5. Up to 5 publications that create an impact story for your community from the reporting period.

In 2025, the primary publication for CCP4 (Acta Crystallogr D Struct Biol. 2023, 79:449-461. doi: 10.1107/S2059798323003595) has been cited 172 times. There are many high impact publications on scientific application of the software, including in Nature-, Science- and Cell-group journals. Further citations accumulate also in the field of AI use, in particular with respect of the use of AlphaFold in structural Biology (FEBS Open Bio 2025, 15:202-222, <https://doi.org/10.1002/2211-5463.13902>; Acta Crystallogr D Struct Biol. 2025, 81:4-21, <https://doi.org/10.1107/S2059798324011999>; Bioinformatics 2025, 41: btaf115, DOI 10.1093/bioinformatics/btaf115).

6. Workshops and new opportunities

See here: <https://www.ccp4.ac.uk/workshops/>

The CCP4 study weekend held in January each year is a highlight event and of interest for other communities; attendance gives opportunity for cross community (net-)working, but online participation is also possible. The meeting gives a good overview on current trends.

Meetings coming up are:

2025

- APS 23-30 June
- Developers meeting – July
- SWSBC Sussex 21-22 July
- Thailand July
- ECM Poznan 25 August
- ECM computing school 22 – 24 August
- Northern Meeting September
- Chile September
- Diamond Light Source 24-28 Nov
- AsCA, Taipei, 1 Dec

2026

- Study Weekend January
- Developers meeting – July
- SWSBC Portsmouth July

- BCA Summer School, York 31 July – 8 Aug
- Northern Meeting September

7. Issues and problems

The CCP-QCr proposal for a new CCP submitted to CoSeC in 2024 was not successful. However, CCP4 needs to integrate new methodology in quantum crystallography. The aim is to reach convergence of crystallographic techniques with quantum mechanics and directly use these methods in structure refinement of macromolecular structures. While the necessary hardware to achieve sufficient quality crystal structures or repeatedly run quantum chemical calculations was generally out of reach, this is changing through technological advances. While resource hungry when addressing larger problems, methodology needs to be advanced so that quantum mechanical calculations are applied to all fields of crystallographic research, across fields like material sciences, chemistry, and biology. Funding is urgently sought.

CCP4 has started now published a web page on ED&I: <https://www.ccp4.ac.uk/ccp4-equity-diversity-inclusion/>. We offer a document on best practice for hiring staff. The short run-times of some contracts (e.g. CoSeC funding) are not aligning with career planning for younger researchers. The CCP4 Equity, diversity & inclusion team are now working on the ED&I roadmap.

Contract issues like multi-party agreements take time and do not allow us to hire immediately into research grants that have been awarded. This is true for both, the BBSRC and CoSeC contracts.

Report from CCPEM for the Period 01/10/2024 to 31/03/2025

Prof. Martyn Winn (UKRI-STFC, CCPEM Principle Investigator)

Prof. Maya Topf (Leibniz Institute of Virology, CCPEM Chair)

Dr. Tom Burnley (UKRI-STFC, CoSeC Project Lead)

1. Background

CCP-EM brings together a diverse and rapidly growing community of researchers involved in cryogenic electron microscopy. Software development is focussed on a suite of software that implements structure determination by single particle analysis and subtomogram averaging, including the interpretation of maps in terms of atomic models. The suite includes contributions from many collaborating groups in the UK and worldwide. Recently, we have implemented the ccpem-pipeliner framework which supports metadata management in workflows, archiving and deposition in the public repositories Empiar/EMDB/PDB. The Doppio user interface sits on top of the ccpem-pipeliner framework, with the Javascript framework allowing web-based operation and remote submission of compute jobs. CCP-EM organise and/or support a wide variety of training activities, with the annual Spring Symposium attracting over 300 in-person delegates. Machine learning techniques are ubiquitous in the current suite and in our future plans, for example implementing operations such as object recognition, classification and denoising, to name a few.

2. Highlights for the current reporting period

Management: Following an election on the 1st of May, we are pleased to announce that the new Chair of CCP-EM is Maya Topf (Centre for Structural Systems Biology, Hamburg) effective immediately. The new Deputy Chair is Giulia Zanetti (Birkbeck/Crick).

New website: We moved to a WordPress platform for www.ccpem.ac.uk, which went live in March 2025.

Software: Doppio v1.2 was released 29/01/25, and v1.3 on 14/05/25. We continue to add new functionality, such as CryoDRGN jobs for heterogeneous reconstruction, CryoVAE and CryoDANN jobs for particle denoising and filtering, ProSMART job for generating atomic distance restraints for protein and nucleic acid chains, and MetalCoord job for generating restraints for coordinated metal ions. In particular, v1.3 included the complete workflow for sub-tomogram averaging based on Relion v5 and associated tools.

General updates to Doppio include an interactive image viewer to allow browsing of images and volumes in job results. We continue to improve the feedback to users on structure quality and have added a table of percentiles for model-to-map fit scores in the Model Validation job.

We now use Conda Constructor to make Doppio installer packages, and additional software packages can be installed with one click from the Doppio GUI. Many model building jobs make use of CCP4 programs, and these have been updated to be compatible with CCP4 9.0. The underlying code base has been updated to use Python 3.9, Node 20 and Electron 35.

Full release notes here: https://www.ccpem.ac.uk/docs/doppio/release_notes.html

DRI-IMB: CCP-EM is a member of the “Digital Research Infrastructure for Integrative Molecular Biology” (DRIIMB) collaboration, alongside CCP4, CCPN and CCPBioSim, funded by the CoSeC bridging call to look at methods for integration across structural biology.

As part of Workpackage 2 “Linking MD simulations to cryo-EM”, Joel Greer is working on tools for assessing heterogeneous reconstruction in cryoEM which can describe discrete or continuous variation in molecular structures. Together with collaborators in the US, a community challenge will be launched in summer 2025. Joel has been working with colleagues from CCPBioSim to generate ground truth data from MD simulations of selected biomolecules. We will also contribute to WP4.2 “Scoping common data formats”, together with CCP4, which will look at joint refinement of atomic models against cryoEM and crystallographic data, using the interactive Moorhen graphics platform.

The 11th **Spring Symposium** was held at the East Midlands Conference Centre, 23-25 April 2025. The Diamond Biological Cryo-Imaging User Meeting took place on the first day. The CCP-EM Scientific Organisers were Laura Spagnolo and Rene Frank.

There were 376 in-person delegates (20 countries, 4 continents) and 613 virtual delegates (unique Zoom views, 29 countries, 6 continents). We received £21k sponsorship which was used to support 32 student bursaries. Raw recordings of talks were available online immediately, followed soon after by polished videos on the CCP-EM YouTube channel <https://www.youtube.com/@ccpem7128/playlists>.

Other workshops: (1) Icknield model building workshop, RAL, Nov 2024; (2) EMBO Integrative structural biology, EMBL Hamburg, Nov 2024; (3) CERN Inverted School of Computing, March 2025; (4) Czech Society for Structural Biology, March 2025. In addition, we have held several Doppio Road shows around the UK.

Training courses are usually run using the STFC Data Analysis as a Service (DAaaS) training system, running on STFC Cloud VMs (with or without GPU) and with Doppio and example data pre-installed. User access managed through DAaaS training website and proxy. The Doppio app runs in the user's local browser giving a faster response than a remote desktop. We find this a very good way of running training courses.

FAIR data: We are members of the Chan Zuckerberg Institute for Advanced Biological Imaging working group on tomography metadata. We have developed the TomoBabel library (with CZII, EMBL-EBI and CNB-CSIC Madrid, <https://github.com/TomoBabel>) and are now working on metadata converters for specific software packages.

Green: Tom Burnley is part of a working group with others in the cryoEM field (from Leeds, Cambridge, Thermo Fisher) to quantify the carbon usage per experiment, and make recommendations to help users to be environmentally sustainable.

3. Overview of work relating to accelerated computing

Different components of the CCP-EM software suite can make use of GPUs, OpenMP and MPI, although in general the software doesn't scale to more than a few nodes. To-date, concerns have centred on memory and storage requirements, but work is planned to look at ways to speed up key steps in the cryoEM pipeline. Work in the Excalibur project "ExaBioSim" has included initial benchmarking of the Relion software for reconstruction. A 6-month pilot project with Oxford RSE Group and DLS/STRUBI (Feb-Aug 25) has carried out further benchmarking and GPU profiling of Relion (benchmarks have been made public). Future work will look to optimise use of GPUs, and to provide the user with recommended program parameters based on available hardware and a database of previous runs.

4. Current plans, developments, or specific applications of AI

AI/ML techniques are in widespread use for interpreting large cryoEM datasets. There are several on-going AI projects including: (1) affinity-VAE for automatic clustering and classification of objects in cryoET tomograms; (2) cryoDANN (Domain Adversarial training of Neural Network) for particle selection as a refinement of the 2D classification step for particle set cleaning; (3) a noise2noise implementation for denoising tomograms; (4) benchmarking of AI tools from a US collaborator (BBSRC funding).

5. Up to 5 publications that create an impact story for your community from the reporting period

"PERC: a suite of software tools for the curation of cryoEM data with application to simulation, modelling and machine learning" B Costa-Gomes, J Greer et al., arXiv preprint, March 2025
<https://arxiv.org/abs/2503.13329>

“Affinity-VAE: incorporating prior knowledge in representation learning from scientific images” M Famili, J Mirecka et al., arXiv preprint, March 2025 (also ECCV2024 proceedings).

<https://doi.org/10.48550/arXiv.2209.04517>

“Roodmus: a toolkit for benchmarking heterogeneous electron cryo-microscopy reconstructions” M Joosten, J Greer, J Parkhurst, T Burnley and A J Jakobi *IUCr* November 2024

<https://doi.org/10.1107/S2052252524009321>

6. Workshops and new opportunities

- EM Facilities Meeting, Warwick, 2-3 June 2025
- CCP-EM SPA Image to Model Data Processing, KMUTT Bangkok, 25-27 June 2025
- Integrative structural biology, Padua, Italy, 30 June – 5 July 2025
<https://www.unipd.it/en/integrated-structural-biology>
- India model building workshop, Bangalore, India, 28 July – 2 August 2025
- "Computing in Crystallography Forum" satellite of ECM35, 22-24 August 2025
<https://ecanews.org/sig-09/computing-in-crystallography-forum/>
- EMBO cryo-EM image processing, Birkbeck, London, 9 - 16 September 2025
<https://meetings.embo.org/event/25-cryo-em>
- CCP-EM Icknield model building workshop, RAL, Oct 2025 (TBC)
- 12th Spring Symposium, EMCC Nottingham, 22-24 April 2026
<https://www.ccpem.ac.uk/symposium/>

7. Issues and problems

Nothing specific.

Computational Engineering

Report from CCP-WSI for the Period 01/10/2024 to 31/03/2025

Prof. Deborah Greaves (University of Plymouth, CCP-WSI Chair)

1. Background

The CCP in Wave Structure Interaction (WSI) began in October 2015 and was extended in October 2020 through the successful CCP-WSI+ proposal. The aim is to bring together and expand the national community of academics and industrial partners, working collaboratively

on activities with a shared objective of building a Numerical Wave Tank (NWT) facility for the simulation of free-surface flow phenomena and fluid-wave interactions with complex structures. A range of networking activities, including focus groups, workshops, and road-mapping exercises, will provide a framework for innovation and the development of strategic software. The goal is to develop a robust and efficient computational WSI modeling tool capable of reliably quantifying wave-induced loads and the corresponding structural mechanics to supplement laboratory and field measurements early in the design process of associated structures. Throughout the project, with the assistance of CoSeC support (2.0 FTE), numerical modelers and experimentalists will combine their expertise to develop an open-source framework using state-of-the-art code coupling and parallelization practices. This will enable the inclusion of highly nonlinear and multi-physics effects in simulations while reducing the overall computational effort required. The software, held in a central code repository, will be professionally engineered, maintainable, and future-proof. It will also be tested and validated against measurement data from new fundamental benchmarking experiments. The project will provide advanced training in computational science and software development, as well as deliver outreach activities for schools and the general public.

2. Highlights for the current reporting period

During this reporting period, a series of virtual meetings of the CCP-WSI Working Group were held to discuss the CCP-WSI and HEC-WSI progress and activities.

Further meetings were held, between the CCP Chair and the CoSeC team associated with the CCP-WSI+ funding, focusing on the three key work packages: code coupling, parallel optimization, and sustaining outputs.

A two-year CCP-WSI Bridging Fund project co-led by the University of Plymouth and CoSeC is underway, aiming to enhance predictive capabilities in wave-structure interaction through AI-driven modelling. With CoSeC Core support, the work focuses on (i) developing AI-based reduced order models for FOWT simulations, and (ii) applying dimensionality reduction techniques to high-fidelity CFD datasets. The project includes over 35,000 node hours on ARCHER2 allocated via HEC-WSI. Key deliverables include training datasets, comparative studies on decoupled/coupled AI-ML model architectures, and journal publications on reduced-order modelling and dimensionality reduction methods.

The CCP-WSI Blind Test Series 5 has attracted 12 participants from 15 universities across 6 countries, submitting a total of 17 solutions. The work has granted 2 special issues on The Thirty-fifth (2025) International Ocean and Polar Engineering Conference (ISOPE 2025): “CCP-WSI BLIND TEST: Sloshing” and “CCP-WSI: Intelligent CFD”. So far, 9 conference papers have been accepted and will be presented on ISOPE conference 1st 7th June.

The CCP-WSI focus group workshop 5 has been held in Newcastle, UK on 2nd April 2025. A total of 46 participants attended the event, including 26 in person and 20 online. This workshop focused on the 2 key themes: “Challenges and Opportunities in Using Machine Learning for Applications in Offshore Renewable Energy” and “Identifying the Grand Challenges in Numerical Modelling for WSI”. This event featured 7 invited speakers, 2 groups for the Panel Discussion.

The “Machine Learning for WSI” training event (4 – 5 Nov 2024) was jointly organised by both CoSeC team and the University of Plymouth team, attracting 14 in-person participants and a peak of 29 on-line attendees.

The WSI community, in collaboration with OpenCFD Ltd, developed modifications enabling OpenFOAM to interface with the Multiscale Universal Interface (MUI) library. This allows coupling with external solvers without modifying OpenFOAM’s source code. The contribution, now integrated upstream, is expected to appear in OpenFOAM v2506, greatly enhancing modular multi-physics and multi-scale simulation capabilities.

The CCP-WSI+ interacts with the eCSE-funded gpuFoam project lead by Prof Gavin Tabor (CCP-WSI co-I). gpuFoam is a research software engineering project to port the popular CFD code OpenFOAM to GPU platforms, of considerable interest given the use of OpenFOAM in WSI. The gpuFoam project is advancing well with core aspects of the code (main matrix solver routines and some boundary conditions) having been ported already; the aim is to complete and release at least two OpenFOAM solver codes, icoFoam and simpleFoam, by the end of the project (summer 2026).

The CCP-WSI YouTube Channel now has 1,723 views and 46 subscribers. The CCP-WSI Organisation on GitHub has 9 repositories, 112 members and 12 followers. The CCP-WSI Organisation on GitHub current contains 16 benchmarking test cases.

The combined CCP-WSI, HEC-WSI and SIG-WSI mailing list has 210 members. The HEC-WSI currently supports 23 sub-groups: 8 Institutions and 15 projects, with 60 active users from 21 organisations. In this reporting period, the CCP-WSI website has received 16358 page views from 2506 users over 3419 sessions. The WSI Catalogue has received 3837 page views from 634 users over 1564 sessions.

The CCP-WSI Catalogue: CAT-WSI is being shared and similar catalogue approaches are being adopted by other CCPs. Poulter, G. (2024) 'Research Object Cataloguing', presentation at the CoSeC Communities Forum, 7 November 2024.

3. Overview of work relating to accelerated computing

Over 440,560 CPUs have been used through 35,641 jobs since January 2023. 123,683 CUs have been used through 50,234 jobs for the period Nov 2024 to April 2025.

4. Current plans, developments, or specific applications of AI

The CCP-WSI team has planned the upcoming CCP-WSI Hackathon, with the schedule currently under discussion.

The CCP-WSI Focus Group Workshop 5 facilitated a broad discussion on grand challenges in the field. Building on these outcomes, the CCP-WSI team plans to convene a meeting to discuss and finalise the roadmap for integrating AI into Wave-Structure Interaction (WSI) research.

Building on CoSeC’s evaluation of AI/ML approaches, the Research Technical Professional (RTP), with support from CoSeC, plans to integrate AI models to replace traditional solvers

within a coupled CFD–CSM computational model of floating offshore wind turbines (FOWTs). The coupling will be achieved via the ParaSIF framework and embedded within a FOWT surrogate model. Both implicit AI-based surrogates and non-intrusive reduced-order modelling techniques will be explored to enhance computational efficiency.

UoP is supporting 2 summer internship students to create immersive visualisation of the coupled CFD simulation of a floating offshore wind turbine.

5. Up to 5 publications that create an impact story for your community from the reporting period.

Yu, S., Ransley, E., Qian, L., Zhou, Y., Brown, S., Greaves, D., Hann, M., Holcombe, A., Edwards, E., Tosdevin, T., Jagdale, S., Li, Q., Zhang, Y., Zhang, N., Yan, S., Ma, Q., Tagliaferro, B., Capasso, S., Martínez-Estevez, I., Göteman, M., Bernhoff, H., Karimirad, M., Domínguez, J.M., Altomare, C., Viccione, G., Crespo, A.J.C., Gómez-Gesteira, M., Eskilsson, C., Verao, G., Andersen, J.F., Palm, J., Niosi, F., Dell’Edera, O., Sirigu, M., Ghigo, A., Bracco, G., Cui, F., Chen, S., Wang, W., Zhuo, Y., Li, Y., Peyrard, C., Benguigui, W., Barcet, M., Robaux, F., Benoit, M., Teles, M., Ntouras, D., Manolas, D., Papadakis, G., Riziotis, V., Zheng, Z., Lei, W., Wang, R., Chen, J., Shao, Y., Visbech, J., Bingham, H.B., Engsig-Karup, A.P., Zhou, Y., Cai, Y., Zhao, H., Shi, W., Li, X., Zeng, X., Xue, Y., Zhuang, T., Wan, D., Engel, G., Tierno, M., Ducrozet, G., Bouscasse, B., Leroy, V., Ferrant, P., Barajas, G. and Lara, J.L., Modelling the Hydrodynamic Response of a Floating Offshore Wind Turbine – a Comparative Study, Feb 2025, In: Applied Ocean Research. 155, 104441.

Poulter, G. (2024) 'CAT-WSI: A Catalogue of Research Objects for Wave Structure Interaction', presentation at the CoSeC Annual Conference, 4 December 2024, Manchester Central Convention Complex.

Liu N, Scarlett G, Davidson J, Windt C, Forehand D, Tabor G & Jia L (2025) 'High-Fidelity Modelling of a Hinged-Raft WEC: a CFD Approach', Proceedings of the 44th International Conference on Ocean, Offshore and Arctic Engineering (OMAE2025), Vancouver, Canada, 22–27 June 2025.

Tan R, Mahfoze OA & Liu W (2025) 'Hydrodynamic Study of Sloshing in a Horizontally and Vertically Excited Cylinder', Proceedings of the 35th International Ocean and Polar Engineering Conference (ISOPE2025), Oslo, Norway, June 2025.

6. Workshops and new opportunities

CCP-WSI Hackathon (Date TBC): A community-focused, hands-on coding event currently under planning. The event is expected to take place within the next 12 months, with details to be confirmed in upcoming CCP-WSI coordination meetings.

Enhancing CFD Visualisation for Immersive Visualisation Suite (IVS) (Summer 2025):

A two-month internship project aimed to improve the real-time visualisation of time-dependent CFD simulations within the Babbage IVS. The project explored Unity-based workflows to address latency issues in ParaView during animated sequences. A functional VR-ready

demonstration was developed, alongside a reproducible workflow to support future outreach and stakeholder engagement.

International Workshop (2026 Q3): One international workshop will be hosted to build engagement with PRACE/Euro HPC Exascale and the international standards (IEA, IEC, ITTC, etc) community.

CCP Conference: CCP-WSI will collaborate with other CCP communities to deliver a joint CCP Conference, providing a platform for sharing best practices, fostering cross-community collaboration, and strengthening collective impact.

CCP-WSI Blind Test Workshop (2025 Q3): A dedicated workshop will be held in Q3 2025 to disseminate results and insights from Blind Test Series 5.

7. Issues and problems

Nothing to report

Report from CCP Turbulence for the Period 01/10/2024 to 31/03/2025

Prof. Sylvain Laizet (Imperial College London, CCP Turbulence Chair)

1. Background

Our daily life is surrounded - and even is sustained - by the flow of fluids. Blood moves through the vessels in our bodies, and air flows into our lungs. Fluid flows disperse particulate air pollution in the turbulent urban as well as indoor environments. Fluid flows play a crucial role for our transportation and our industries. Our vehicles move through air and water powered by other fluids that mix in the combustion chambers of engines. Many of the environmental and energy-related issues we face today cannot possibly be tackled without a better understanding of the dynamics of fluids. From a practical point of view, fluid flows relevant to scientists and engineers are turbulent ones; turbulence is the rule, not the exception.

To date, a complete theory of fluid flow phenomena is still missing because of the complexity of the full equations describing the motion of a fluid. Their understanding and control are however crucial to improve technologies especially with minimal ecological impact as well as to anticipate events, in many areas ranging from engineering applications (e.g., industrial process, propulsion and power generation, car and aircraft design) to environmental sciences and technologies (e.g., air quality, weather forecasting, climate predictions, flood disasters monitoring). Significant progress has been made recently using high performance computing, and computational fluid dynamics is now a critical complement to experiments and theories.

The CCP Turbulence is aiming to (i) considerably enhance the UK capabilities to simulate complex turbulence problems that were until very recently beyond imagination, (ii) offer user support, training and networking activities and (iii) enable capability computing on emerging hardware platforms. The software developments and collaborative activities will give UK researchers a unique opportunity to be the first to explore new physics and to answer basic questions regarding the physics and modelling of turbulent flows found across a range of engineering, physiological and geophysical applications.

To perform numerical experiments that push the boundaries of turbulence research, purposely developed, highly scalable simulation codes need to be used. As different types of flows are governed by different equations and boundary conditions, a number of algorithmically different flow solvers are needed to ensure that computational resources are optimally used. Four flagship flow solvers exist in the UKTC: XCOMPACT3D, OPENSBLI, NEKTAR++ and code_saturne. These four codes are well-established, widely used, open source, highly scalable, with the ability to perform turbulence-resolving simulations of a wide range of turbulent flows. They are routinely used for simulations with thousands of nodes for production runs. Each flagship solver has been developed with the requirement to produce high-quality results with as much accuracy as possible and as few computational resources as possible. The variety of flow solvers within the UK Turbulence Community is crucial for robustness and reliability, to support innovation, and to maximise impact by using the most efficient and most accurate algorithm for a given flow configuration.

Moreover, CCP Turbulence is also supporting who enabling libraries OPS and 2DECOMP&FFT that allows researcher to build their own fluid solver running on an efficient fashion on thousands of cores and GPUs.

2. Highlights for the current reporting period

The CCP Turbulence is playing a pivotal role by directly supporting the development, software integration, and maintenance of XCOMPACT3D and OPENSBLI, along with their core libraries 2DECOMP&FFT and OPS. Due to funding limitation, it is challenging for the CCP Turbulence to support more flow solvers. However, the CCP Turbulence is now supporting SENGA+ (via the bridging call), mainly because the UKTC has recently absorbed the UK Consortium on Turbulent Reacting Flows (which did not secure funding during the last HEC call).

Highlights for the current period:

1-Maintenance and Support: This task is dealing with software sustainability. With a constantly evolving HPC landscape, it is crucial to make sure that the features already available in our flow solvers are optimised, reliable and robust enough to be compatible with the latest versions of libraries and compilers. We have developed a proper CI infrastructure for this.

2.Code Optimisation: Continuous improvement of flagship flow solvers and libraries (XCOMPACT3D, OPENSBLI, 2DECOMP&FFT and OPS) to achieve exascale readiness, with a specific focus on GPU porting and performance enhancement.

3.Expanded Computational Capabilities: Integration of SENGAs as a core supported code, including advanced particle modelling functionalities within OPS. An MHD and a Lagrangian Particle Tracking module has been added to XCOMPACT3D.

3. Overview of work relating to accelerated computing

The focus of the CCP Turbulence is to prepare the UK turbulence community to the exascale era. All the activities of the CCP Turbulence described in the section below are targeting heterogeneous computing platforms for its finite-difference flow solvers.

4. Current plans, developments, or specific applications of AI

No activities related to AI as we are missing CoSeC support (0.5 FTE).

5. Up to 5 publications that create an impact story for your community from the reporting period

Jian Fang, Sylvain Laizet, Alex Skillen, A high-order finite-difference solver for direct numerical simulations of magnetohydrodynamic turbulence, Computer Physics Communications, Volume 307, 2025, 109400, ISSN 0010-4655, <https://doi.org/10.1016/j.cpc.2024.109400>.

6. Workshops and new opportunities

- Delivery of several invited talks at prestigious conference and for prestigious series of seminars
- Organisation of few hackathons about our flow solvers and supported libraries.
- Release of an updated version of the 2DECOMP&FFT library & Xcompact3d & OpenSBLI & SENGAs
- 2025 UK Turbulence Consortium annual meeting held in March 2025 in Cambridge with over 115 attendees.

7. Issues and problems

- Limited options to benchmark our software (at scale) on GPU nodes in the UK.
- End of the Tier2 systems in the UK and end of ARCHER2 (Nov 2025 or Nov 2026), with no new supercomputers planned in the foreseeable future in the UK.
- Lack of clarity about the future of the CoSeC support for CCP projects (after 2026).

Report from CCP-NTH for the Period 01/10/2024 to 31/03/2025

Prof. Shuisheng He (University of Sheffield, CCP-NTH Chair)

Dr. Wei Wang (UKRI-STFC, CoSeC Project Lead)

1. Background

Nuclear energy plays a vital role in the transition to a low-carbon future. As outlined in the UK Government's 2024 Civil Nuclear Roadmap, nuclear is expected to provide 25% of the country's electricity by 2050. However, challenges such as cost-effectiveness and safety remain barriers to large-scale deployment. Advanced reactor designs, including Advanced Modular Reactors (AMRs), aim to address these issues, but it requires cutting-edge tools to analyse complex physical processes. Nuclear thermal hydraulics (NTH), which focuses on the study of heat and fluid flow in reactor systems, is fundamental to the safe and efficient operation of reactors. These systems involve multi-scale, multi-physics behaviours, where small-scale effects like turbulence and boiling can have large-scale impacts on performance. Traditional NTH analysis relies heavily on empirical methods and simplified models, which are increasingly being supplemented by advanced simulation techniques such as computational fluid dynamics (CFD).

CCP-NTH was established to drive innovation in this space. It supports the development of robust, validated modelling tools and fosters collaboration across academia, industry, and national laboratories. CCP-NTH is structured into two core work packages: WP1 focuses on community building, including organising workshops, training, and benchmarking activities and WP2 is methodology and code development and maintenance, which includes high fidelity modelling and simulation and robust (reliable, affordable and user-friendly) CFD. Under WP2, the CCP has developed CHAPSim, a UK-based high-fidelity CFD code designed for direct numerical simulation (DNS) of thermal hydraulic flows, now available as open source. The project also supports the development of efficient engineering tools based on RANS approaches, such as coarse-grid CFD for complex reactor components. The long-term vision of CCP-NTH is to build a sustainable suite of simulation tools, incorporating next-generation computing technologies like AI, machine learning, and accelerator-based platforms.

2. Highlights for the current reporting period

Ongoing development and maintenance of the community DNS solver CHAPSim2 have continued to expand its high-fidelity simulation capabilities, which are essential for advancing research in nuclear thermal hydraulics and fusion energy applications. One key enhancement is **the addition of high-order cylindrical coordinate** which allows CHAPSim2 to accurately model pipe and annular flows without relying on complex boundary reconstruction methods like immersed boundaries. Another enhancement is the **implementation of a smooth, accuracy-preserving turbulence inlet and convective outlet**, enabling realistic simulations of thermally developing flows. In collaboration with STFC and the University of Sheffield, **multiphase flow modelling** has been introduced to support the development of next-generation reactor safety analysis tools. To better capture surface roughness effects in supercritical fluids, **an immersed boundary method** has also been implemented in collaboration with STFC and University of Sheffield. Furthermore, **a magnetohydrodynamics module in Cartesian coordinates** was developed and validated in partnership with University of Sheffield, STFC and UKAEA, enabling the simulation of liquid metal flows in fusion reactor environments. Throughout development, **high-performance computing efficiency** has remained a priority, supporting the goal of sustainable, energy-efficient green computing.

Training sessions and workshops have also been held to introduce CHAPSim2 functionalities and discuss ongoing improvements and applications.

A benchmark study has been conducted to evaluate the performance of RANS models for simulating Reactor Vessel Auxiliary Cooling Systems (RVACS). This collaborative effort brought together participants from both industries, such as EDF Nuclear Services, EDF R&D, and Frazer-Nash Consultancy, and academia, including STFC, the University of Sheffield, and the University of Manchester. The study focused on comparing various solvers, turbulence models, near-wall treatments, and simulation strategies. By systematically evaluating these approaches, the benchmark has provided valuable insights into best practices and helped build confidence in selecting appropriate numerical models and mesh resolutions for simulating complex passive cooling systems. **Two technical workshops** were held as part of the initiative to facilitate discussion and share experience among participants. All benchmarking data has been collected and will be made **publicly available** via STFC following the publication of the associated research paper.

3. Overview of work relating to accelerated computing

During the reporting period, we have initiated GPU porting of the community DNS code CHAPSim2. Initial discussions and strategic planning have focused on how to use GPU acceleration to enhance simulation performance while maintaining code flexibility. The approach will primarily use pragma-based offloading (OpenMP and/or OpenACC) within Fortran, ensuring portability and ease of integration with the existing codebase.

Planned activities include:

- (1) Identifying key computational components for acceleration, such as differential operators and FFT routines;
- (2) Integrating GPU-accelerated libraries, such as cuFFT and the GPU-enabled version of 2decomp&FFT;
- (3) Refactoring computational loops to enable parallel execution by restructuring operations to exploit iteration independence and improve memory access efficiency;
- (4) Optimizing memory access by minimizing global memory use and maximizing shared memory and register utilization to enhance overall performance.

These efforts aim to build a scalable, high-performance GPU-enabled version of CHAPSim2 to support demanding simulations in DNS.

4. Current plans, developments, or specific applications of AI

During the reporting period, no direct AI/ML activities have been undertaken. However, we have developed a roadmap to explore the application of AI and machine learning to improve the prediction of thermal-hydraulic phenomena in reactor systems, particularly in strongly heated and mixed convection flows relevant to passive cooling and decay heat removal.

As a first step, we will conduct a detailed literature review and feasibility assessment to identify suitable AI/ML techniques for our domain. Several areas of potential impact have been identified:

- (iii) Data generation and management: A major barrier to AI/ML application in this field is the lack of high-quality datasets. To address this, we plan to use CHAPSim to generate high-fidelity DNS data for representative configurations, such as forced and mixed convection in pipes under atmospheric and supercritical conditions. These datasets will be curated within a structured data management system to support future model development and training.
- (iii) (ii) Reduced-order and surrogate modelling: We will explore the feasibility of building surrogate or reduced-order models to accelerate simulations of complex convection systems. This includes cases involving conjugate heat transfer, where the differing time scales between fluid turbulence and solid thermal response pose significant modelling challenges.
- (iii) (iii) Improving flow physics models: Finally, we aim to apply techniques such as clustering and artificial neural networks to improve key components of RANS modelling, including buoyancy effects, turbulent heat flux models (e.g., SGDh, GGDh), and the Boussinesq approximation—areas known to contribute to modelling uncertainty.

5. Up to 5 publications that create an impact story for your community from the reporting period

The community solver CHAPSim2 is a high-fidelity Direct Numerical Simulation tool that has emerged as a powerful resource for understanding complex thermal-hydraulic phenomena in next-generation nuclear systems. Developed to simulate turbulent flows in reactor-relevant geometries with heat transfer, CHAPSim2 addresses a critical need for predictive capabilities in regimes where conventional models fall short, particularly under supercritical pressure conditions and conjugate heat transfer environments. The impact of CHAPSim2 is exemplified by its use in the recent study by He et al. (2024) [1], published in the International Journal of Heat and Mass Transfer, which investigated the effect of conjugate heat transfer on turbulence and thermal transport in an upward heated pipe flow at supercritical pressure. This work demonstrated that solid wall conduction significantly alters local heat fluxes and turbulence characteristics, findings that challenge assumptions commonly used in reactor design and simulation. Further extending its impact, CHAPSim2 was featured in the 11th International Symposium on Supercritical Water-Cooled Reactors [2], where it supported detailed analysis of mixed convection flows. These results contribute essential insights for the design of advanced nuclear reactors, especially where passive safety systems depend on complex heat transfer behaviour. Through such applications, CHAPSim2 not only enables deeper scientific understanding but also helps build benchmark datasets critical for validating lower-fidelity tools and guiding the development of robust, physics-informed models. Its ongoing development, including GPU acceleration and AI-ready data generation, ensures that CHAPSim2 will remain central to the UK and international efforts in clean nuclear energy innovation.

[1] He, Jundi, Wei Wang, Bing Xu, and Shuisheng He. "Impact of conjugate heat transfer on the turbulence and heat transfer in an upward heated pipe flow at supercritical pressure." *International Journal of Heat and Mass Transfer* 233 (2024): 126004.

[2] Chinembiri, Kenneth, Shuisheng He, and Wei Wang. "A study of the effect of pyramid roughness on turbulent heat transfer in supercritical water." 11th International symposium on supercritical water-cooled reactors: Pisa, Italy, February 3-5, 2025. Pisa University Press, 2025.

6. Workshops and new opportunities

Upcoming Meetings and Events

- Organising CCP-NTH Annual Technical Meeting, 19–20 June 2025, Derby
- Organising CCP-NTH Advisory Committee Meeting, 20 June 2025, Derby
- A group active participation in 21st International Topical Meeting on Nuclear Reactor Thermal Hydraulics (NURETH-21), 1–5 September 2025, Busan, Korea

Workshops and Training Activities

- CHAPSim Users' Workshop, 1 day, July 2025 (TBC)
- CHAPSim Developers' Workshop, 1 day, October 2025 (TBC)
- CCP-NTH Research Day (mini-workshop), monthly from September 2025 (TBC)

7. Issues and problems

The additional 1 FTE supported by the bridging funding was not in place until March 2025, which has delayed the planned work. Furthermore, due to the current recruitment freeze at STFC, the assigned personnel may require some initial training to undertake the relevant tasks, which could further slowdown progress.

Tomographic Imaging

Report from CCPi for the Period 01/10/2024 to 31/03/2025

Dr. Martin Turner (University of Manchester, CCPi Co-Investigator)

Dr. Edoardo Pasca (UKRI-STFC, CoSeC Project Lead)

1. Background

The collaborative computational project in tomographic imaging (CCPi) was initiated in 2012 to support the UK non-clinical computed tomography (CT) community with the aim of: ‘Developing software and the necessary training to increase the quality and level of information that could be gleaned from X-ray projection data.’

Specifically, three suites of open-source software tools have been developed by CCPi in response:

1. pre-processing software for image calibration and noise reduction;
2. reconstruction software able to reconstruct a 3D/4D volume images from ‘sub-optimal’ datasets; and
3. segmentation / quantification techniques to extract relevant quantities from 3D data.

This extensive python ecosystem supports CT users across a wide range of imaging devices including lab X-ray, national mid-scale research facilities and synchrotron X-ray and neutron CT. Our CoSeC support enables us to develop, maintain, and promote this CCPi Core Imaging Library (CIL) as well as a range of other related software; notably interactive visualisation and the [iDVC](#) (Digital Volume Correlation) code initiated as a result of CCPi funded international visits.

2. Highlights for the current reporting period

We report several community engagement events in the reporting period, spanning multiple code bases and reaching 150 users. This high number of participants to the training sessions is a testament of the high need of training in CT by the community, as outlined in the [2018 EPSRC X-Ray CT roadmap](#), and of the high standard of the training that CoSeC provides.

The larger event has been the [CIL User Meeting 2024](#) at RAL, with 59 participants from 20+ national and international institutions. The event brought together researchers and scientists to present their use of CIL and it has been preceded by a training day (for the first time with introductory and advanced parallel sessions) and followed by a hackathon to help participants to get started or use advanced methods in their research.

CIL was also presented with a short introductory training at the [ESRF User Meeting 2025](#) with 24 participants at the European Synchrotron Research Facility in Grenoble. Another smaller session was carried out at the University of Lyon and INSA-Lyon, with 10 participants. Contacts were made with the developers of OpenRTK, another software package that could be used in CIL as another engine for the physics model. Finally, we held a very popular [online training session for CIL](#) split in 3 different tracks (introductory, iterative and advanced) which was attended cumulatively by 92 users.

In the reporting period CIL has seen 2 [stable releases](#): 24.2.0 and 24.3.0, with general bugfixes, the addition of the PD3O algorithm, which is a more flexible variant of the “classic” PDHG, a refactoring of the base class for the data structure aimed at easier maintenance of the code base. The [iDVC](#) software for digital volume correlation (DVC) was presented at the international conference of the Digital Image Correlation society [iDICs](#) in France.

A fortnightly [CCPi Show & Tell](#) online meeting is organised to strengthen the community bonds and foster exchange of ideas and information.

3. Overview of work relating to accelerated computing

[Stochastic optimisation](#) algorithms show faster convergence than the standard deterministic counterparts, due to their lower sensitivity to noise and shallow optimisation landscape. In the past few months, we developed in CIL a stochastic optimisation framework that enables very flexible configurations of stochastic optimisation algorithms, such as stochastic gradient descent, SAG/SAGA, SVRG/LSVRG, with custom preconditioning, step size and momentum.

A publication on the use of another stochastic optimisation algorithm in cardio-respiratory motion compensated MRI is under review and demonstrated a nearly 5-fold speed up of convergence.

Work is in progress for the development of an efficient FDK algorithm that minimises the idle time between CPU and GPU data processing, increasing the efficiency of the code, enabling reconstruction of large datasets (3k3 and more) with standard workstations.

Planned activities on use of accelerated computing relate to enabling the use of full GPU pipeline for smaller problems, by making the CIL data container agnostic on the backend (currently bound to CPU by NumPy).

4. Current plans, developments, or specific applications of AI

Planned activity is to enable the efficient use of AI libraries such as PyTorch within the CIL pipeline. This will require work on the underlying structure of CIL data containers. A [workshop](#) and following hackathon was planned with CCP-SyneRBI and took place in early April. The workshop was attended by CCPi and CCP-SyneRBI developers and by developers of other scientific software such as ASTRA.

5. Up to 5 publications that create an impact story for your community from the reporting period

Under review: “Efficient motion-corrected image reconstruction for 3D cardiac MRI through stochastic optimisation”, E.Pasca, M. Duff et al with PMB.

6. Workshops and new opportunities

A list of workshop and conferences with a participation of CoSeC or CCPi:

- [Workshop and hackathon](#) on Efficient integration of SIRF/STIR/CIL with Pytorch, 7-9 April 2025 UCL, London
- [CIL Training](#) at [ToScA NA](#), 12/05/2025, AMNH New York USA
- [ToScA NA](#), 13-14 May 2025, AMNH New York USA
- [NoCTURN](#), 15-16 May 2025, AMNH New York USA

- [iDVC training](#) at the University of Bath, 22/05/2025 Bath
- 8th Workshop on [Advances on X-Ray Imaging](#), Diamond Light Source, 24/06/2025 RAL
- [Dimensional X-Ray CT \(dXCT\)](#), 25-26/06/2025 WMG, University of Warwick, Coventry
- [ToScA UK & Europe](#), 10-12 September, ESRF Grenoble France
- [IBSim-4i](#), 20-24 October 2025, IOP London
- CIL User Meeting 26, 10-14 February 2026 (TBC) Rutherford Appleton Laboratory

7. Issues and problems

Nothing to report

Report from CCP SyneRBI for the Period 01/10/2024 to 31/03/2025

Dr. Evgueni Ovtchinnikov (UKRI-STFC, CoSeC Project Lead)

Prof. Kris Thielemans (University College London, CCPSyneRBI Chair)

1. Background

For medical imaging, the UK is a globally leading country. The Collaborative Computational Project in Synergistic Biomedical Imaging (CCP SyneRBI), established in 2015 as CCP in Positron Emission Tomography and Magnetic Resonance imaging (CCP PETMR) and extended in 2020 under the new name, aims at bringing together the best of the UK's imaging expertise to capitalise on the investment in this area. Research has shown that the use of MRI intermediate results can improve PET image quality, and vice versa (in some cases), with some scanners capable of acquiring MR and PET data simultaneously. These techniques are now under investigation for PET and Single Photon Emission Computed Tomography (SPECT) in the context of Molecular RadioTherapy (MRT). The main deliverable of the project is an open-source reconstruction software framework we named SIRF (Synergistic Image Reconstruction Framework). SIRF is simple enough in use for educational and research purposes, thus reducing the “barrier for entry” for new contributors to imaging research and development, and at the same time powerful enough to process real scanner data. SyneRBI also contributes to other open-source projects, including STIR (Software for Tomographic Image Reconstruction) for PET and SPECT.

STFC CoSeC support for this CCP currently focusses on developing the SIRF and associated code base that provides an easy-to-use Python environment built around existing Open-Source reconstruction software. This includes maintaining network, website, running workshops and training courses, on top of the software engineering effort that contributes to SIRF development, testing, deployment and documentation.

2. Highlights for the current reporting period

2.1 PET Rapid Image Reconstruction Challenge and associated software, data and workshop.

Our major activity in the current reporting period was organising the PET Rapid Image Reconstruction Challenge (PETRIC) which ran over the summer and autumn 2024 (1 June to 30 September). Its primary aim was to stimulate research into the development of fast PET image reconstruction algorithms applicable to real world data. Motivated by the success of the clinical translation of regularised image reconstruction in PET and other modalities, PETRIC concentrated on reconstruction in a Maximum A Posteriori setting. The participants had access to a sizeable set of phantom data acquired on a range of clinical scanners. The main task for the participants of the challenge was to reach a solution which would be close to the converged target image (in terms of standard clinical measures) as quickly as possible (as measured in terms of computation time). This task therefore required a balance between algorithm design and implementation optimisation. An example solution which reaches the target image quality but takes a long time was provided at the beginning of the challenge.

2.1.1 PETRIC data.

The PET raw data was pre-processed to enable researchers to take part even if they have little experience in handling real world data. The pre-processed data is available at <https://petric.tomography.stfc.ac.uk/data/> (hosted by STFC). We are currently uploading raw and pre-processed data onto Zenodo to ensure future availability.

2.1.2 PETRIC software.

Our Open-Source software SIRF, STIR and CCPi Open-Source software CIL were provided to develop and test the algorithms. The participants were instructed to use STIR (via SIRF) projectors in their implementations, together with provided multiplicative and additive projection data, so that the reconstructed image quality and timing performance only depended on the reconstruction algorithm. Participants had access to the source code, a docker image as well as the possibility to run reconstructions on the STFC Cloud.

The PETRIC software includes instructions and pre-processing scripts for PET raw data preparation. These scripts can be used for preparing data from supported PET scanners to obtain files in a standard format, suitable for reconstruction in standardised pipelines or for experimentation.

In addition, the PETRIC software contains code for computing metrics on the reconstructed images, as well as scheduled running submissions, including computation of the metrics and timing, with results made accessible via a TensorBoard server.

All software is available at <https://github.com/SyneRBI/PETRIC/> and <https://github.com/SyneRBI/PETRIC-backend>.

2.1.3 PETRIC Workshop.

The PETRIC challenge culminated in a workshop held in November at the IEEE Medical Imaging Conference (MIC), Tampa, FL, USA where invited international experts presented recent

advances in image reconstruction, covering developments in CT, MRI and PET. This was followed by an overview of the PET Rapid Image Reconstruction Challenge, presentations of the 3 top-ranked teams and an award ceremony for the winning teams.

The workshop was attended by ~100 researchers from across the world. All presentations and recordings can be found on the PETRIC workshop page on the SynerBI website <https://www.ccpsynerbi.ac.uk/petric/>.

2.2 Domain-adaptation of image reconstruction algorithms

SynerBI has developed, adapted and tested optimisation algorithms to the image reconstruction problem, in collaboration with CCPi. These are now implemented in the Core Imaging Library (CIL) and maintained by CoSeC. As illustrated by some of the submissions to PETRIC, some of those algorithms are not robust to changes in the data and/or acquisition models and more work is needed to make them ready to be used by a wider non-expert community. In this reporting period, we have investigated the performance of various implementations of the PDHG algorithm (and its stochastic version) which has led to suggestions for “preconditioning” strategies making the algorithm more robust. We anticipate that this work will lead to recommendations with sample implementations available in CIL.

2.3 Progress on a standard for PET Raw data

SynerBI continues to contribute towards the development of a standard for PET raw data called PETSIRD, an effort coordinated by the Emission Tomography Standardization Initiative (ETSI), see <https://etsinitiative.org/>. We are working towards a draft release end 2025.

On 3-4 Nov 2024, we helped promoting and arranging the 2nd ETSI hackathon in Tampa, FL, USA just after the IEEE MIC 2024, with about 25 local and 10 remote participants. For UK participants, T&S support was available via SynerBI. Main outcomes included use-case software for displaying, handling and reconstructing PETSIRD data as well as prototype converters from vendor formats to PETSIRD. CoSeC contributions included deployment of the PETSIRD Python suite to PyPI. We continue to develop our extensions to STIR to read raw data from vendor formats (work on progress on GE RDF8 and Siemens PETLINK) as well as writing PETSIRD data.

Software is available at <https://github.com/ETSIInitiative/PETSIRD> and <https://github.com/ETSIhackers>.

2.4 Raw Data Sharing and Searching (XNAT)

For both PET and MR raw data, enabling easy data access and the searching of acquisition parameters will facilitate research. The XNAT open-source platform for research image sharing is being extended, as a part of this project, to host searchable MR and PET raw data parameters and associated data. Simon Doran (Institute of Cancer Research) who has extensive experience of XNAT has been providing advice.

3. Overview of work relating to accelerated computing

We achieved in the reporting period a substantial improvement of algebraic operations, with operations on PET/SPECT acquisition data made up to 3 times faster, and image data algebra up to 15 times faster. We plan to achieve still further performance improvements of PET/SPECT data algebra by employing the Python Array API, as well as handling of CUDA GPU data structures.

4. Current plans, developments, or specific applications of AI

The above-mentioned optimised access to data should have a major impact on the tensor algebra performance of PyTorch, which we plan to use in our work on AI applications to image reconstruction. In

addition, we have added a set of functions facilitating the integration of SIRF with PyTorch, together with examples of some networks for image reconstruction, see <https://github.com/SyneRBI/SIRF/pull/1305>. This work will form the basis for future extensions with sample implementations of the currently most promising methods incorporating advanced learning methods in reconstruction, where we will be targeting PET in the first instance with other modalities following.

5. Up to 5 publications that create an impact story for your community from the reporting period

Ferri, T.; Caracciolo, A.; Ghisio, F.; Piroddi, M.; Pandocchi, M.; Fiorini, C.; Carminati, M.; Pascali, V.; Protti, N.; Mazzucconi, D.; Grisoni, L.; Ramos, D.; Ferrara, N.; Thielemans, K.; Borghi, G. Design and Validation of a SPECT Prototype for Treatment Monitoring in BNCT and First Experimental Tomographic Results. *IEEE Transactions on Radiation and Plasma Medical Sciences* 2025, 1–1. <https://doi.org/10.1109/TRPMS.2025.3562079>.

This work discusses a monitoring system and data processing for Boron Neutron Capture Therapy, a novel method for targeted cancer treatment currently under investigation. The work includes an extension of SPECT collimator modelling in our open-source software STIR and its evaluation based on Monte Carlo and measured data. The first author was sponsored by SyneRBI for an exchange from Milan to UCL.

6. Workshops and new opportunities

After the end date of this reporting period, SyneRBI and CCPi collaborated on organising a workshop and hackathon on “Integration of libraries for physics informed deep learning in imaging”. More information will be included in the next progress report.

SyneRBI will be sponsoring the PET is Wonderful conference [PET is Wonderful 2025](#) to enable the recording and dissemination of the educational sessions. We are also planning a training day after the conference.

On July 14-18 2025 Emission Tomography Standardization Initiative (ETSI) will hold its 3rd Hackathon (virtual-only). The main scope of this hackathon event will be to continue the development of meaningful and practical use cases for the PETSIRD standard since our last ETSI hackathon in Tampa, utilizing all latest features and elements definitions added to the standard over the last 6 months. The focus in this event will be the completion of the in-silico

(i.e. simulated) demonstration of the utility of PETSIRD through an end-to-end simulation study of realistic cylindrical and brain PET phantom data. Plans for a 4th Hackathon (hybrid) are in progress.

7. Issues and problems

The main stumbling block affecting our software development remains the installation of SIRF and its pre-requisites under various Operating Systems. While we provide a Virtual Machine and docker containers, this is still a hurdle for many users. We also have not yet succeeded in the Windows installation of Gadgetron.

Quantum Computing

Report from CCP-QC for the Period 01/10/2024 to 31/03/2025

Prof. Viv Kendon (University of Strathclyde, CCP-QC Chair)

Dr. Alin Elena (UKRI-STFC, CCP-QC Project Lead)

1. Background

CCP-QC was started in 2019 with a remit to bring about close co-operation between other CCPs and the quantum computing community, rather than build a separate computational community for quantum computing. CCP-QC aims to:

1. Build an active research community encompassing CCP members interested in enhancing their simulations by adding quantum computing capability to their code, and quantum technology researchers working on applications of quantum computing to simulations;
2. Generate small projects supported by CoSeC staff time to develop methods appropriate to specific applications, leading to proof-of-concept demonstrations on early quantum hardware. This will also develop capacity in CoSeC for quantum computing;
3. Collaborate with the National Quantum Computing Centre (NQCC) and the Quantum Computing (QCi3) Hub in developing early applications of quantum computing;
4. Provide training in quantum computing for researchers who are expert in computational science but lack quantum computing knowledge;
5. Support career development of early career researchers through subsidised meetings, annual awards for best presentation, and a pairing scheme to link those working in computational science with those working in quantum computing and quantum algorithms;

6. Promote cross-CCP networking to share knowledge on early applications of quantum computing, enabling the widest possible early adoption of quantum enhanced computational science;
7. Provide information on early quantum computing applications in academic research to the wider community through their website maintained for the life of the CCP. The website will advertise contact details and opportunities to join the community through meetings, training days, and signposts for collaboration.

CCP-QC does not aim to develop or curate its own body of code, rather, any code developed will belong to the CCP community for that application. Pointers to those applications, and information about what works and why will be available on the website, to guide those developing further applications.

2. Highlights for the current reporting period

Supported by CCP bridge project funding, there are two work packages running in parallel, one around Lattice Boltzmann in the Computational Engineering Theme and another around Green's functions formalism for electronic structure within the Computational Materials and Molecular Science Theme.

Lattice Boltzmann: As part of our broader efforts to explore the role of emerging technologies in computational science, initial work has begun to assess the potential of quantum computing for accelerating the Lattice Boltzmann Method (LBM). During the reporting period, initial activity has focused on laying a strong foundation for this new research direction. This has included an in-depth literature review spanning quantum computing fundamentals, quantum algorithms, and their relevance to fluid dynamics, as well as more targeted studies on existing attempts to connect quantum computing with LBM and related numerical methods. To support this, time was spent learning the fundamentals of quantum computing and getting started with tools such as Qiskit. Informal catchups with collaborators helped share ideas and discuss current challenges and possible directions, while participation in the Workshop on Quantum Computing for Fluids, held in Edinburgh on 27 March 2025, helped build a better understanding of ongoing research and connect with others working in the field. Although this work is still at an early stage, these initial steps have helped highlight important questions and technical challenges and have prepared the ground for more focused investigations over the next year—particularly into how quantum computing might speed up complex flow simulations.

Electronic Structure: During this period, the Green's function project on Rydberg hardware was developed along two different lines of research: that concerning software, the software-hardware interface and the software itself. These areas of research are certainly interconnected. With regard to the software, Green's function can be calculated in two ways: one method is the Variational Quantum Algorithm, which calculates the time evolution of the wave function according to McLachlan's variational principle. The second option uses the subspace variational quantum eigensolver to determine the eigenvalues of the excitations and construct the imaginary part of Green's function from the Lehmann representation.

Regarding the software-hardware interface that maps the qubits into fermions, we are searching for the most suitable transformation to render the 2D Hubbard model (our basic Hamiltonian)

feasible with Rydberg atom-type hardware. Currently, two approaches appear more feasible: the Jordan–Wigner (J–W) transformation and the Verstraete–Cirac (V–C) transformation. The J–W transformation has the advantage of a reasonably implementable experimental setup, but it involves transforming the Hubbard U interaction term into non-local terms. In contrast, the V–C transformation retains the locality of the interaction term but requires a more complex experimental setup. We are currently investigating which transformation is better suited to the Green's function software implementation in Rydberg atom quantum computers.

Both work packages collaborated with the EXCALIBUR project on Quantum Enhanced and Verified Exascale Computing (QEVEC) and the Quantum Technologies project on Quantum Algorithms for Nonlinear Differential Equations (QuANDiE); both grants ended on 31 March 2025.

3. Overview of work relating to accelerated computing

See above, all work we do in CCP-QC can be subsumed under the accelerators headline.

4. Current plans, developments, or specific applications of AI

Currently not applicable to CCP-QC.

5. Up to 5 publications that create an impact story for your community from the reporting period

1. Comprehensive review article:

Quantum algorithms for scientific computing R Au-Yeung, B Camino, O Rathore and V Kendon
2024 Rep. Prog. Phys. 87 116001 <https://doi.org/10.1088/1361-6633/ad85f0>

2. Ongoing work on benchmarking and standards for quantum computing led by NPL and the NQCC has created a comprehensive review and repository of code:

A Review and Collection of Metrics and Benchmarks for Quantum Computers: definitions, methodologies and software, Lall et al, <https://arxiv.org/abs/2502.06717>

6. Issues and problems

Nothing to report.

New Community Reports

CCP-volumeEM

Dr. Martin Jones (The Francis Crick Institute, CCP-volumeEM Chair)

The scope of CCP-volumeEM is to formalise a community of researchers and software engineers working in computational aspects of volume electron microscopy (vEM), which is an umbrella term for a group of EM techniques that are used to image biological material through continuous depths of at least one micrometre, at nanometre resolution. Existing communities such as volumeem.org primarily deal with the imaging and sample preparation aspects of the field, but in surveying the community it became clear that the computational aspects of data storage, handling, and analysis have become the major bottleneck in many vEM workflows.

Our goals for the community are split into three work packages:

WP1: Community building - expanding upon related initiatives by identifying the relevant stakeholders across imaging and computational communities. We are planning two community events, the first in November 2025 for consultation on the draft roadmap, the second in Autumn 2026 to further develop the roadmap, establishing a 5-year vision. At least one community hackathon will also be held.

WP2: Research Software and Data Management - establishing a set of best-practices, resources, and exemplar projects to demonstrate scalable and flexible vEM data handling and analysis principles. This will include the integration of cutting-edge technologies and next generation file formats (NGFF) and the promotion of FAIR data principles throughout. Ensuring these outputs are accessible to non-computational end-users will be a central focus.

WP3: Training and Skills - establish a community of practice centred around key open-source analysis platforms such as Fiji and Napari, as well as the vEM specific tools that use these platforms.

We have filled the posts of two RIAs (one at the Crick, one at the Rosalind Franklin Institute, 0.2 FTE each) and a project manager (at the Crick, 0.25 FTE). For WP1, a community survey has been designed and shared, with team-members presenting a poster and/or short talks about CCP-volumeEM at the [CCP-EM Spring Symposium](#), the [Volume EM Gordon Research Conference](#) in Barcelona, as well as an accepted flash-talk at the upcoming [Microscience Microscopy Congress](#) in Manchester. The first community event has been scheduled for 3rd November at the Francis Crick Institute, with a provisional programme consisting of a set of talks followed by breakout groups to discuss the roadmap.

In WP2, a set of vEM specific software environments have been containerised and discussions have been held to deploy these as training platforms via STFC's DAaaS resource. In parallel, work on creating an exemplar workflow to demonstrate a generalised toolkit for 3D data reconstruction via stitching, registration and alignment has begun, using NGFF and modern python tools to flexibly and efficiently cover the range of requirements.

The information required to successfully deliver WP3 will be extracted from the survey responses, and the containers developed in WP2 will be used to standardise and simplify the practical requirements for delivering training.

CCC-ParaSolS

Collaborative Computational Community in Particulate Solids Simulations

Dr. Kevin Hanley (University of Edinburgh, CCC-ParaSolS Chair)


www.ccc-parasols.ed.ac.uk | ParaSolS@ed.ac.uk

CCC-ParaSolS is a community open to all UK-based scientists and engineers from academia and industry with a common interest in simulating particulate solids (granular materials) for a variety of applications. These granular materials include natural soil deposits, pharmaceutical powders, food ingredients (e.g., salt, flour, agricultural grains), and aggregates and cement used in construction. The most popular particle-scale simulation method is the Discrete Element Method (DEM), although there are many others.

Currently, developments of particulate simulation methods happen in discipline silos, inhibiting full exploitation of methodological advancements. CCC-ParaSolS aims to address this issue by creating an overarching, multi-disciplinary community which embeds diversity and inclusiveness in its composition, governance, and activities. CCC-ParaSolS will promote the use of opensource software for granular simulations, deliver bespoke training on DEM code and HPC usage, host physical network events at different UK locations, undertake high-priority code development projects based on community needs analyses, and create a five-year vision for the community. Through these activities and more, CCC-ParaSolS will establish a strong foundation for a long-term Collaborative Computational Project in particulate solids simulations.

Since January 2025, the community has been established with > 70 members, supported by a [comprehensive website](#) and [LinkedIn presence](#). Governance structures have been put in place for CCC-ParaSolS. Following an [introductory webinar](#) in March, the [first hybrid Network Event](#) took place in Edinburgh from 14–16 May with almost 40 in-person and 10 online attendees. This event included training, delivered by developers, on the use of three popular open-source DEM codes (LAMMPS, MercuryDPM and YADE), facilitated workshops, a panel session, and ample opportunities for networking. A photo from this highly successful event is below.

In the next six months, we will continue to expand the community. Progress will be made on the development of code benchmarking cases. A second Network Event is envisaged in October 2025 at which training on the use of HPC (ARCHER2) will be delivered by experts from EPCC, gaps in granular simulation capabilities will be identified and prioritised for code development in year 2, discussions will take place on code benchmarking progress and the role that AI/ML will play in particulate solids simulations, and the first version of our collective five-year vision for the community will be co-created.



In-person attendees enjoying the first CCC-ParaSols Network Event in Edinburgh

CCP-DCM

Data-centric Computational Mechanics

Prof. David Ham (Imperial College London, CCP-DCM Chair)

1. Community Scope and Ambitions

This is a [collaborative computational community](#) centred around the FEniCS and Firedrake finite element systems. [FEniCS](#) and [Firedrake](#) are world-leading software frameworks for the numerical solution of partial differential equations (PDEs). These numerical simulations are essential for advancing science and engineering across a very broad range of disciplines. FEniCS and Firedrake are used to develop solvers in the geosciences (ocean, atmosphere, cryosphere, geodynamics), nuclear fusion (plasma, tritium transport, breeding blankets), physiology and medicine (brain, heart), and many more besides. Our role is to support and develop the core simulation toolchain to enable current and new users to continue to push the limits of the possible in simulation scale.

2. Work done so far

Establishing a CCP

- Website up at <https://ccp-dcm.github.io>. An ac.uk domain is pending o Management committee has met monthly.

Initial events

- Firedrake USA user & developer meeting in Waco, TX. February '25 o Firedrake seminar, Rutherford Appleton Laboratory. March '25 o Firedrake and FEniCS developer retreat. Cambridge. March '25
- Firedrake course at African Institute of Mathematical Sciences. Kigali, Rwanda April '25 o Use case hackathon bringing together Firedrake and FEniCS developers and users. Totnes. May '25

3. Plans for the next 6 months

- Scoping meeting with key stakeholders. London. June '25
- FEniCS '25 user and developer workshop. Groningen, Netherlands. June '25
- Firedrake '25 user and developer workshop. Leeds, September '25
- User community survey. Summer '25
- Firedrake and FEniCS new user tutorials, Autumn '25

CCP-TEPP

Towards a Collaborative Computational Project in Theoretical and Experimental Particle Physics

Dr. Ed Bennett (Swansea University, CCP-TEPP Chair)

1. Community scope

CCP-TEPP covers Theoretical and Experimental Particle Physics. Computationally intensive Particle Theory includes Lattice Quantum Field Theory (including Lattice QCD), a major consumer of HPC resources. Particle Experiment includes many aspects of the software pipeline supporting collider experiments such as those at the Large Hadron Collider: software for event generation, detector simulation, triggering and data acquisition, tracking and reconstruction of particle collisions, and subsequent analysis. Particle phenomenology straddles the two areas.

2. Goals and ambitions

Aim for 2025 is to collect community input and prepare a 5-year roadmap for software needs in the field; in 2026 to begin progressing aspects of this roadmap in collaboration with CoSeC. Ambition to improve cross-discipline communication and coordination, reduce the amount of duplication of effort where appropriate, and raise baseline computational skill level to improve productivity.

3. Work done so far

Workshops in Edinburgh (2025-04-29–30) focusing on software for particle theory, and in Warwick (2025-06-01–02) on software for experimental particle physics, collecting input to the roadmap.

4. Work planned over the coming 6 months.

Collate notes from the two workshops above into a draft roadmap. Circulate this draft roadmap for comments from the community and adjust responding to feedback.

Knowledge Exchange event at the Hartree Centre in September, bringing together computational researchers from across theory and experiment, to share best practices and discuss approaches to delivering on the roadmap.

CCP-UKNR

UK Numerical Relativity

Prof. Eugene Lim (King's College London, CCP-UKNR Chair)

1. What is CCP-UKNR: UK numerical relativity community?

UK Numerical Relativity (<https://www.uknumericalrelativity.org/>) is a community driven group of about 40-50 astronomers and physicists whose research involves the use of computational methods to solve Einstein's equations of general relativity of gravitation, which is a set of non-linearly coupled partial differential equations. As a major component of the international LIGO-Virgo-KAGRA gravitational wave detector collaborations and the space-based LISA interferometer experiment, UK scientists played leading roles in the rapidly growing gravitational wave science. Furthermore, the UK leads or plays key roles to several internationally recognised numerical relativity codes ExaGRyPE (Durham), BAM (Cardiff), METHOD (Southampton), MHDueT (Nottingham), Einstein Toolkit (ETK Collaboration) and GRChombo/GRTEclyn (GRTL Code Collaboration). GRChombo and BAM were benchmark codes for the DIRAC-3 acquisition and roll-out.

The present CCP-UKNR Chair is Prof Eugene Lim (King's College London), with co-Leads Dr Katy Clough (QMUL), Prof Mark Hannam (Cardiff), Dr Geraint Pratten (Birmingham), Dr Patricia Schmidt (Birmingham) with Technical RSE Dr Miren Radia (Cambridge/DIRAC). The CoSeC Leads are Prof David Emerson and Dr Tyrone Rees.

2. What are our present challenges?

- Community Building:** UK historically is a pioneer in the development of numerical relativity techniques, starting with the famous Gregynog Conference organised by Prof Bernard Schutz at Cardiff over 40 years ago, but divergence in methodologies and emphasis have led to a fragmentation of the community. Fortunately, in the last 15 years or so the community has undergone a renewal, with the development of new UK-led codes (such as GRChombo/GRTEclyn and ExaGRyPE), and the hiring of new international scientific leaders of MHDueT and BAM. A key goal of the CCP is to bring together all these different voices to form a sustainable community of friendly competitors and collaborators, and to advocate for issues relevant to us.
- Transitioning to a GPU world:** With future digital infrastructure expected to be heavily reliant on GPUs, the community is actively transitioning its core codes to support GPUs. For example, ExaGRyPE is built with Indigenous GPUs support, while GRTEclyn and MHDueT is completing their transition into using the AMReX software framework with support for GPU offload, with an expected roll-out by the end of 2025.
- Code Development/Training/Exploitation:** Codes are increasingly complex, due to both the addition of more sophisticated physical modeling, and advances in digital infrastructure technology. A key challenge is to develop a sustainable strategy for the roll out of new codes to the community (e.g. training) and the development of skills that are required to maintain and refine the codes.

- **Long term strategic planning:** The CCP aims to serve as a community forum for the discussion of long-term health and goals of UKNR, especially with regards to the everchanging digital infrastructure landscape. A key topic of discussion is code/library consolidation, and the possibility of the development of a UK-led AMR library.

3. Activities

- First Annual UKNR Community meeting (QMUL, Sept 10-11, 2024, <https://sites.google.com/view/uknr24>). In this meeting, the community agreed to apply for the CCP, which was then awarded.
- Community Townhall (KCL, Feb 26, 2025): The main stakeholders of the community meet to discuss the plans and ideas for moving forward with the CCP.
- Code Benchmarking Exercise (6 months, starting June 2025): To be led by Dr Miren Radia (Technical Lead), the community will undertake code benchmarking for the NR codes that are actively in use by the community. We aim to take a “capability snapshot” of the present codes running on both CPUs and/or GPUs, in order to understand where we stand as a community with respect to the present digital landscape, and also to inform the design and scoping of the next generation systems such as DIRAC-4.
- Collaboration with CoSeC : A meeting was held on April 9th between the CoSeC and members of the CCP to discuss collaboration opportunities. We are presently scoping out possible projects which can be jointly undertaken.
- First UKNR International Meeting at Gregynog (June 17-19, 2025, <https://sites.google.com/view/nrgregynog>): We will hold the first UKNR International meeting, bringing together stakeholders from the UKNR community across the world to discuss the science, the technology and the future of UKNR in the coming digital infrastructure landscape.
- Second Annual UKNR Community meeting/Townhall (Birmingham, Autumn 2025): In addition to discussing the science and other issues pertaining to the UKNR community, we will also present the first code benchmark results and update the community on the progress of the CCP.

CCP-AHC

A Collaborative Computational Project serving Arts, Humanities, and Culture Researchers

Dr. Eamonn Bell (Durham University, CCP-AHC Chair)

1. Community scope

We focus on DRI and software that supports computationally intensive research and innovation in the arts, humanities, and culture areas, which is typically funded by in the UK by UKRI’s Arts and Humanities Research Council (AHRC). Previous efforts to gather community requirements targeting AHRC-facing researchers provide some insights into HPC and cloud computing usage

to date, where 38% of survey respondents used or were moving to “HPC/Cloud” in 2021. We aim to increase this proportion. The culture within the community around research outputs, skills, and institutional support were viewed as barriers to adoption of large-scale compute of this kind. Addressing these needs are particularly urgent given the widespread interest in and adoption of AI methods within the field. We seek to identify projects with the potential for development through engagement with RSE and computational science resource within CCP-AHC and CoSeC during the scoping project and over the longer-term future of this project and related initiatives. Our technical scope includes any code, pipelines, or workflows from within the community that is running or has the potential to run on UK-funded HPC and advanced computing resource, including those supported by e.g. GPUs and other novel accelerators.

2. Your goals and ambitions

The goal of CCP-AHC is to support the sustainable and efficient development of software, pipelines, and workflows used by arts, humanities, and culture researchers who make use of UK-based digital research infrastructure (DRI). It will do so by disseminating and implementing the Collaborative Computational Project (CCP) model that has successfully been used by many other scientific software communities over the past several decades, opening up CoSeC to a new area of research and innovation. By introducing the CCP model to our community, which is not currently well-known, we aim to increase engagement with DRI and large-scale compute. We will work closely with data services, skills, and people development projects funded by the AHRC programme for an infrastructure for Digital innovation and curation in Arts and Humanities (iDAH), including the Digital Skills in Arts and Humanities Network ([DISKAH](#)). We also seek to engage with representatives of recent AHRC infrastructure investments, including DiSSCo, CoSTAR, and RICHeS, and comparable European and international efforts, with a view to widening engagement with DRI.

3. Work done so far

- Project website is launched at www.ccpahc.ac.uk, including a blog and resources page with c. 200 30-day active users at time of writing
- Mailing lists created ([CCP-AHC-ANNOUNCE](#), [CCP-AHC-DISCUSS](#)) with 70 and 40 subscribers, respectively at time of writing
- Open call for codes, workflows, and pipelines available for contributions now online at <https://www.ccpahc.ac.uk/activities/codes-eoi/>
- CCP-AHC Town Hall 2025 (<https://www.ccpahc.ac.uk/activities/town-hall-2025/agenda/>) held in Durham on 22 May 2025 with c. 25 attendees in-person and 13 online, followed by first Advisory Group meeting on 23 May 2025
- Past and planned dissemination of project at HPC-SIG, Durham HPC Days 2025, PASC’25, RSECon25, and CIUK 2025
- 25+ one-to-one meetings between Delivery Team members and key national stakeholders since January 2025

- Design agency ([curious12](#)) commissioned to develop a visual identity for CCP-AHC, and design directions shared within project team

4. Work planned over the coming 6 months

- Forthcoming announcement of monthly, online community calls to identify prospective CCP-AH “working group” members and/or working groups
- Report of requirements gathered during Town Hall 2025 discussions in round-table format, including draft roadmap and work plan for 2026
- Refresh of call for codes, pipelines, and workflows in light of Town Hall 2025 and Advisory Group feedback
- Four, in-person regional engagement events planned in Scotland, Northern Ireland, Wales, and England. (Save the date for an event at University of Edinburgh, afternoon of 6 November 2025!)
- Technical scoping exercises including free code review sessions facilitated using project resources from early 2026
- Town Hall 2026 planned for September 2026

High End Computing Consortium (HEC) Reports

The UK Turbulence Consortium

Prof. Sylvain Laizet (Imperial College London, UKTC Chair)

The UK Turbulence Consortium (UKTC) brings together complementary expertise and coordinates activities to look at coherent, rational, and strategic ways of understanding, predicting and controlling turbulent flows using High Performance Computing. The consortium is crucial for the UK in order to augment and unify the research efforts of its participants and to communicate its findings to a wider audience.

Firstly, funded in 1995, the UKTC has been through six highly successful iterations. It has seen significant growth since its inception, from 5 original members to more than 80 members over 30 UK institutions today. In the last 4 decades, the UKTC has (i) demonstrated its ability to convert access to national High-End Computing (HEC) resources into internationally-leading research (hundreds published papers since 1995 with thousands of non-self-citations), (ii) established its international competitiveness, (iii) helped its members to leverage and secure substantial funding from governmental bodies and industry, (iv) allowed the discovery of new fluid flow phenomena, leading to new ways of improving the beneficial effects and reducing the negative effects of turbulent flows, and (v) facilitated the design of more sophisticated turbulence models redefining industry standards.

Four flagship flow solvers exist in the UKTC: [OpenSBLI](#) for compressible flows with a finite difference approach, [Xcompact3d](#) for incompressible flows with a finite-difference approach, [Nektar++](#) for compressible/incompressible flows with spectral/hp element approaches and [Code_Saturne](#) for compressible/incompressible flows with a finite volume approach. These four codes are well-established, widely used, open source, highly scalable, with the ability to perform turbulence-resolving simulations of a wide range of turbulent flows. They are routinely used for simulations with thousands of nodes for production runs.

Work done in the past 12 months:

Maintenance and support activities for the UKTC four flagship solvers: With a constantly evolving High Performance Computing landscape, it is crucial to make sure that the features already available in the UKTC flagship solvers are optimised, reliable and robust enough to be compatible with the latest versions of libraries and compilers. This has been achieved with the support of CoSeC staff.

Software development activities towards exascale computing: (with funding via the [ExCALIBUR](#) initiative and the CCP Turbulence): The UKTC flagship solvers are being re-designed so that they can leverage at scale modern hybrid supercomputers based on CPUs and/or GPUs. A notable activity has been an initial testing of the scalability at scale of [Code_Saturne](#) and [Xcompact3d](#) on Marenostrom 5 in Spain. This has been possible thanks to EuroHPC and would not have been possible with the current UK computing infrastructures.

Organisation of the 2025 UKTC annual meeting in Cambridge (March 2025) with over 115 attendees: Keynote lectures were delivered by Oriol Lehmkuhl (Barcelona Supercomputing Centre) on “Current trends on numerical simulation for the design of new and disruptive aerodynamic configurations using supercomputers” and Richard Gilham (Bristol) on “Isambard- AI and Isambard 3- construction and user experience”. UKTC members showcased their work with 45 oral presentations and 20 posters.

Organisation of several hackathons and user meetings for the UKTC flagship solvers.

Work planned over the coming 6 months:

- Continuation of the maintenance and support, and software development activities for the UKTC flagship solvers, in particular for performance at scale on modern GPUs.
- Preparation and organisation of training, hackathon and user meetings for the UKTC flagship solvers.

NERC Atmospheric and Polar Consortium

Dr. Grenville Lister (University of Reading, Atmospheric and Polar Consortium Lead)

1. Nature of the consortium

The Atmospheric and Polar consortium constitutes those projects that run models and/or workflows on ARCHER2 and JASMIN to support work aligned with NERC strategy and vision for these disciplines. It is one of 3 consortia, the others being Mineral and Geophysics and Oceanography and Shelf Seas. The vast majority of projects are NERC funded either through research grants or through National Capability programmes. The consortium partners with the associated JASMIN atmos consortium.

2. Role of the consortium

The consortium is responsible for managing and administering ARCHER2 compute time and storage and by extension JASMIN disk and tape storage. Projects in the consortium run independently of each other and interact with the consortium only inasmuch as the consortium handles computational resources on their behalf. The consortium does not decide on the science activities of its constituent projects.

Users are requested to submit applications to the NERC HPC Steering Committee in either March or October. Following submission deadlines, the HPC Steering Committee, which includes the consortia leads assess the applications and allocate compute resource accordingly.

3. Consortium administration (2025)

ARCHER2

- Managing 77.5% of the NERC ARCHER2 allocation (~6M node-hrs)
- Servicing 52 projects in 2025-2026 with 179 users

- From 21 UK Institutions

JASMIN

- Managing 11.8PB of disk and 25PB of tape
- Servicing 180 projects with ~580 users

Example projects

- CANARI - Climate change in the Arctic-North Atlantic Region and Impacts on the UK
- TerraFIRMA - Assessing a range of climate mitigation strategies.
- HRCM - High Resolution Climate Modelling
- TBT - Assessing past, present and future changes in global mountain water resources.
- SUNSET - Simulating under ice shelf extreme topography



**Computational Science Centre
for Research Communities**

Report produced by Mai Tuyet Hoang

Science and Technology Facilities Council

Rutherford Appleton Laboratory,
Harwell Campus, Didcot, Oxfordshire OX11 0QX, UK
T: +44 (0)1235 445000 F: +44 (0)1235 445808

Sci-Tech Daresbury,
Daresbury, Warrington, Cheshire, WA4 4AD, UK T:
+44 (0) 1925 603000 F: +44 (0) 1925 603100

Contacts

Damian Jones, Programme Office Manager

Georgia Lomas, Events Manager

CoSeC@stfc.ac.uk