



EPSRC Service Level Agreement with STFC for Computational Science Support

FY 2020/21 Annual Report
(Covering the period 1 April 2020 – 31 March 2021)

May 2021

Table of Contents

Background	3
CoSeC Project Office	5
CCP5 – Computer Simulation of Condensed Phases	7
CCP9 – Electronic Structure of Solids	8
CCPNC – NMR Crystallography	10
CCPQ – Quantum Dynamics in Atomic Molecular and Optical Physics.....	11
CCP-Plasma – HEC-Plasma Physics	14
CCPi – Tomographic Imaging	15
CCP SynerBI - Synergistic Biomedical Imaging.....	17
CCPBioSim - Biomolecular Simulation at the Life Sciences Interface	19
CCP WSI+ - Wave Structure Interaction.....	21
CCP Turbulence - Turbulence	23
CCP NTH - Nuclear Thermal Hydraulics	25
CCP QC – Quantum Computing.....	27
MCC – Materials Chemistry Consortium.....	28
UKCP – UK Car-Parrinello Consortium	30
UKCOMES - UK Consortium on Mesoscale Engineering Sciences	32
HEC Plasma Physics	33
HECBioSim - High-End Computing for Biomolecular Simulation	33
UK-AMOR - Atomic, Molecular and Optical Physics R-matrix Codes	36
UKTC - The UK Turbulence Consortium	38
UKCTRF - UK Consortium on Turbulent Reacting Flows.....	39
Software Outlook.....	40
Resource Planning and Yearly Costing	42
Metrics.....	42

Background

The Scientific Computing Department (SCD) of the Science and Technology Facilities Council (STFC) provides computational science support for a number of scientific communities funded by EPSRC, and organised in Collaborative Computational Projects (CCPs) and High End Computing (HEC) Consortia. This programme of work is carried out by staff at the Daresbury and Rutherford Appleton Laboratories under a Service Level Agreement (**SLA**) with EPSRC, and its main objectives are:

- Develop and maintain codes and methods to keep them internationally relevant for current and evolving scientific challenges and hardware capabilities
- Widen participation in the exploitation of methods and codes through training and scientific collaboration
- Support collaboration and coordination of the various communities to broaden and strengthen the UK-based research activities aligned with EPSRC's goals
- Provide career paths and professional development opportunities for computational scientists and engineers focused on method and software development
- Widen engagement with the broader UK and international communities engaged in developing methods and software for computational science and engineering

The nature of the support provided is tailored to the needs of the communities and can include:

- **Development of theory, algorithms, and software:** This is a key element of support for many current projects, resulting in long-term, continued expansion and updating of the software programs. It may include the consolidation of existing codes into a more sustainable community software package
- **User support and training:** This includes individual support and training as well as help to organise and conduct events such as workshops, summer schools and study weekends. Support for Centres of Doctoral Training is also offered
- **Outreach and promotion of computational science and engineering activities:** facilitate the exchange of expertise, and tools, reaching out to new communities, including experimentalists and industry, nationally and world-wide
- **Collaboration on scientific projects:** working together with scientists in the communities to advance scientific research and help nurture projects and develop new opportunities
- **Porting, optimisation, and benchmarking** on local and national computing platforms including High Performance Computing (HPC) and new architectures: evaluation of new hardware and software technologies
- **Maintenance, distribution,** license management, dissemination and demonstration of software
- **Management of scientific data:** This includes activities such as, for example, the development of visualisation and workflow management tools, database and archiving, and verification and validation activities
- **Co-ordinate and nurture existing and new communities,** from practical tasks such as organising community meetings, to representing the communities in strategic activities in the UK Research Councils and abroad.

CoSeC, the Computational Science Centre for Research Communities, brings together these activities with those in support of CCP4 (partly funded by a BBSRC grant), CCP-EM (funded by an MRC grant), and CCP-WSI (funded by an EPSRC grant): <https://www.scd.stfc.ac.uk/Pages/CoSeC.aspx>

The communities currently supported are summarised in the table below:

Project	Title	Project Chair	CoSeC Project manager
CCP9	Computational Electronic Structure of Condensed Matter	Prof Stewart Clark	Dr Leon Petit
CCP-NC	NMR Crystallography	Prof Paul Hodgkinson	Dr Simone Sturniolo
CCPi	Tomographic Imaging	Prof Philip Withers	Dr Edoardo Pasca

CCP PET-MR	Computational Collaborative Project in Synergistic PET-MR Reconstruction	Prof Kris Thielemans	Dr Evgueni Ovtchinnikov
CCP-BioSim	Biomolecular simulation at the life sciences interface	Dr Sarah Harris	Dr Tom Keal
CCP-WSI+	Wave structure interaction +	Prof Deboah Greaves	Dr Stephen Longshaw
CCP Turbulence	Turbulence	Prof Sylvain Laizet	Dr David Emerson
CCP-NTH	Nuclear thermal Hydraulics	Prof Shuisheng He	Pror David Emerson
CCP-QC	Quantum computing	Dr Viv Kendon	Dr Alin Elena
MCC	UK Materials Chemistry Consortium	Prof Scott Woodley	Dr Tom Keal
HEC-BioSim	High-End Computing Consortium in biomolecular simulation	Dr Syma Khalid	Dr James Gebbie
UKCP	United Kingdom Car-Parrinello Consortium	Prof Matt Probert	Dr Dominik Jochym
HEC-Plasma	Plasma High-end Computing Consortium	Prof Tony Arber	Dr Joseph Parker
UKCOMES	UK Consortium on Mesoscale Engineering Sciences	Prof Kai Luo	Dr Michael Seaton
UK-AMOR	UK Atomic, Molecular and Optical physics R-matrix Consortium	Prof Jonathan Tennyson	Dr Martin Plummer
UKTC	UK Turbulence Consortium	Dr Sylvain Laizet	Prof David Emerson
UKCTRF	UK Consortium on Turbulent Reacting Flows	Prof Nilanjan Chakraborty	Prof David Emerson

A brief description of each community can be found in the community-specific sections below. More information is available at <http://www.ccp.ac.uk/> for the CCPs and at <https://www.epsrc.ac.uk/research/facilities/hpc/access/highendcomputingconsortia/> for the HEC consortia. The current level of support awarded to the communities is as follows. Please note that this table shows the original level of support awarded following the EPSRC call in 2015 and the current level of support (in brackets) following the most recent EPSRC CCP call in 2019.

Community	Core support per project (FTEs per annum)
CCP9	2.6 (3.0 from 1 April 2020)
CCP-NC	1.4
CCPi	1.3 (1.5 from 1 April 2020)
CCP-PET/MR	1.25 (1.8 from 1 April 2020)
CCP-BioSim	1.25 (1.95 from 1 April 2020)
CCP-WSI+	2.0 (from 1 April 2020)
CCP Turbulence	1.75 (from 1 April 2020)
CCP-NTH	1.0 (from 1 April 2020)
CCP-QC	0.6 (from 1 April 2020)
Materials Chemistry	2.5 (2.0 from 31 Oct 2018)
HEC-BioSim	1.0 (0.8 from 31 Oct 2018)

UKCP	1.0
HEC Plasma	0.2
UK COMES	1.0 (0.6 from 31 May 2018)
UK-AMOR	0.2 (from 1 April 2018)
UKTC	0.3 (from 30 June 2018)
UKCTRF	0.13 (from 07 Jan 2019)

CoSeC Project Office

The objectives of the Project Office are to provide coordination through the CCP Steering Panel and programme oversight through periodic reports, and input to reviews, to EPSRC and the SLA Steering Committee, support for the organisation of networking and training events, and manage the progress of the programme staff, L&D and liaise with the communities supported, in particular to help increase the impact of the programme.

Summary Report (1 April 2020 – 31 March 2021)

During this reporting period we have made some changes to the structure of the project office, dividing the work into four work packages that, although separate, also include a large degree of overlap and interaction. The four work packages are “**Project Management**” lead by Damian Jones, “**Impact**” lead by Dawn Geatches, “**Technical**” lead by Stephen Longshaw and “**Strategy/People/Relations**” lead by Barbara Montanari.

Project Management: This work package exists to provide management of the CoSeC programme, including financial monitoring and forecasting, reporting and wide-ranging activities such as the coordination of the CoSeC Steering Committee and CCP Steering Panel meetings, support for CoSeC related activities such as conferences, seminars and training courses and interaction with STFC’s Scientific Computing Project Management Office for sharing of best practice and moving towards a standard, project management approach within the organisation.

Project office work has continued as normal despite the covid lockdown and associated issues. The 2019/20 annual report was submitted in June and the 2020/21 interim report was submitted in November. Internal, quarterly meetings with STFC based CoSeC project leads have taken place via zoom. Monthly monitoring of staff effort bookings and spend has continued as planned, and financial forecasts for financial year 2021/22 have been created.

Impact: The aim of this work package is to build awareness of CoSeC’s valuable role in the scientific landscape of computational research in the UK, increasing its visibility as the hub of a community network for current and future members and funders.

The impact work continued to highlight and promote the presence and outputs from the CCP and HEC communities supported by CoSeC, which over the last year comprised: 20 news articles, many of which are composites of news from several CCPs and/or HECs; approximately 40 tweets; 5 case studies. The news articles included the announcement of the inaugural winner of the CoSeC Impact Award (CIA), and the contribution of several CCPs and HECs to the research into the SARs-Cov2 virus. @CoSeC_community now has 216 followers on Twitter, a steady increase from 159 in May 2020, and the most popular tweets (i.e. those that are ‘liked’) contain photos of people. Three of the five case studies featured the work of the CIA winner (Dr. Natalie Tatum) and the two runners-up, and an additional REF-style case study was written for Natalie’s personal use. The 2021 CIA applications opened and closed with 12 applications, a modest increase from last year’s nine, although the panel of judges remarked on this year’s high standard of applications.

Work on re-designing the CoSeC webpages continues within the strictures of UKRI's organisational policies, which means CoSeC will not have its own website, although it has sufficient freedom to build a cohesive and easily navigable collection of webpages.

Technical: The goal of this work package is to better understand the technical and scientific landscape of CoSeC and create a framework to support its development in a way that is able to evolve. This knowledge will be used to identify overlap between projects, or research. In particular the aim is to facilitate and foster cross-project collaboration where this might not otherwise happen.

In this time period this work package has started a number of new initiatives around the goal of knowledge and skills sharing. A new series of thematic working groups has been created, initially around the themes of website technology and scientific code coupling, each working group has a designated facilitator and will run for as long as is practicably useful. Outputs include a new set of website platform technologies being trialled and a CoSeC-wide review of current website approaches, the coupling group has created an internal report and planned a new publicly available seminar series to start in May 2021, initially with 3 talks to run at two-month intervals. New working groups around other topics will be created as needed.

A number of other items have been considered within this work package, specifically around the goal of better connecting the scientific landscape of the CoSeC communities, as well as exploring a number of ideas to elevate the external status of CoSeC. In particular, the idea of delivering an annual scientific review across the communities is under consideration, this would be in addition to the existing annual report and would focus on detailed scientific output that represents the landscape of communities that CoSeC supports. An annual dissemination event is also being considered, initially as a significant part of an existing appropriate event like STFC's annual Computing Insight UK (CIUK) conference, but with the aim of eventually creating a CoSeC specific event along the lines of the UK Fluids conference which was created on the back of the EPSRC funded UK Fluids Network initiative. Finally, once the pandemic situation allows, a number of activities designed to foster internal collaboration on both scientific and technical topics across CoSeC are being considered and will be delivered as part of upcoming whole-team meetings.

Strategy/People/Relations: This work package will address the wider ranging aspect of CoSeC activities including its leadership, interaction with CoSeC communities and funding agencies, and CoSeC's relationship with other national and international computing initiatives for research. The work package will also manage staff development, performance and recruitments.

Strategy: In collaboration with STFC and Southampton University colleagues, Barbara Montanari prepared and submitted a Statement of Need for a Physical Sciences Data infrastructure (PSDI) in response to the EPSRC call for Large Scale Infrastructure. Most chairs of CCPs and HECs, as well as the Chair of the CCP Steering Panel, (alongside leaders of experimental institutions and other initiatives such as, for instance, the UK Catalysis Hub) supported this statement explicitly and provided examples to illustrate the need and the impact that this infrastructure will have. Barbara also provided support for the CCP5++ networking grant proposal (Alin Elena is co-I, Barbara is a CCP5++ Steering Group member), which was successful.

People: During the reporting year, Barbara was appointed member of CECAM Council, of the ExCALIBUR hardware and enabling software Steering Committee (being already a member of the ExCALIBUR Steering Committee), and was part of the EPSRC RSE Fellowships panel. Two recruitment campaigns took place (for CCPi and CCP-NC) and one member of CoSeC staff was supported through promotion. The reorganisation of the CoSeC Project Office into four distinct work packages and distribution of their leadership across a team has contributed to the professional growth of Damian Jones, Dawn Geatches and Stephen Longshaw through empowerment.

Relations: Barbara gave an invited talk at the Royal Society Yusuf Hamied Workshop for India and the UK about the UK software ecosystem and challenges of exascale computing. This should provide opportunities to create new collaborations with India. Barbara and Stephen Longshaw presented a CoSeC update at the six-monthly HEC Chairs Meeting. CoSeC organised and presented at the six-monthly CCP SP Meeting and Barbara attended the CCP9 community meeting, the CECAM Council meeting, and the Psi-k annual community meeting. Regular meetings to discuss CoSeC funding took place between EPSRC and STFC.

Significant Impact (1 April 2020 – 31 March 2021)

Development of CoSeC news articles and case studies has continued to help improve the CoSeC brand and communicate the depth and quality of CoSeC led work to a wider audience. We also had twelve applicants for this year's CoSeC Impact Award, an increase from nine in the first year.

Outline Plans 2021/22 (1 April 2021 – 31 March 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

Moving forward into 2021/22 the project office will continue to explore new ways of raising awareness of the CoSeC brand and communicating the impact of the work that CoSeC funded CCPs and HECs are involved in. The current five year CoSeC SLA cycle will come to an end on 31 March 2021 so, as we move into 2021/22, we will work with EPSRC to agree the funding model and level going forward. The 2022 CoSeC Impact Award will be launched towards the end of 2021.

CCP5 – Computer Simulation of Condensed Phases

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 35 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 450 UK members and over 550 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.

Our core software support covers numerical energy minimisation, classical molecular dynamics and Monte Carlo simulation, ranging from atomistic to multi-scale molecular systems. An increasing effort is exerted to tackle major challenges in cutting edge parallel simulations, linking atomistic and higher level models with first principles (quantum), spanning longer time- and length-scales by means of coarse-graining and mesoscale modelling so as to provide reliable multi-scale simulation protocols.

CCP5 major software and methodology support includes five active projects which together account for over 4,000 active licence holders worldwide in 80 countries. DL_POLY is a general purpose, classical, particle dynamics program. DL_MESO is a general purpose Dissipative Particle Dynamics program. DL_MONTE is a general purpose particle Monte Carlo program. ChemShell is an advanced command line environment with tools and methods for modelling materials systems simultaneously in classical and quantum terms. DL_FIELD is a cheminformatics program for conversion of materials structures from XYZ/PDB description to structure and force-field model files suitable for input into DL_POLY, DL_MESO and DL_MONTE.

CCP5 also provides funding for undergraduate student bursaries, workshop and conference funding and international visitor tours in the UK as well as an extensive range of training events including the annual CCP5 summer school.

Summary Report (1 April 2020 – 31 March 2021)

Please note: CCP5 was not renewed as part of the 2019 EPSRC CCP call. However, there were a number of activities that were able to continue unfunded. They are reported here for completion.

The majority of planned events were postponed due to COVID. The CCP5 summer school was oversubscribed like usual. Some students were offered places this year for next years school with the hope that this will be in person. The CCP5 40th Annual General Meeting was a virtual event attended by over 300 participants from the UK, US, Japan, India and Europe. In total, attendance showed nearly a four-fold increase on previous in-person AGMs. Talks covered computational research on a wide range of topics, from the performance of tyres, through iron deficiency anaemia, to quantum computing for scientific calculations. A highlight of the day was the announcement of the winner of the inaugural biennial CCP5 Prize and Lecture Award: Prof Kostya Trachenko. STFC Executive committee elections were held with two new members being appointed, J Skelton, University of Manchester and A Logsdail University of Cardiff.

The CCP5++ grant was funded this will allow networking activities to go for another three years starting with September 2021. PI is Prof Paola Carbone from University of Manchester. There were elections for a new member of the executive committee with a record number of four candidates standing in the election. Dr David Wilkins from Queen's University of Belfast was elected. For the first time we used a totally anonymised system of voting, helios.

The CCP5 AGM this year will be morphed in an online event with the MMMhub and CCP9. Provisional dates are in the week of 13 September. Prof Paola Carbone and Dr Alin Elena are part of the organising committee and executive committee member Prof Chris Lorenz in the Scientific committee from CCP5.

A one day event to launch the new grant for the network is planned for second half of September at Cosener house with a restricted number of participants due to covid. The CCP5 Summer School will go online this summer July 11-22, with a record number of applicants this year, XXX. A new advanced course was added: Machine Learning for Interatomic Potentials, which is lead by Prof Gabor Csanyi from University of Cambridge. The CCP5 Bursaries programme continues this year with the call to close on 1st of May 2021.

CCP9 – Electronic Structure of Solids

The Collaborative Computational Project on Computational Electronic Structure of Condensed Matter (CCP9) brings together UK researchers in the field of ab-initio electronic structure calculations of condensed matter systems. The field includes the study of metals, semiconductors, magnets, and superconductors from atomistic quantum mechanical calculations. The activities of CCP9 encompass such areas as magneto-electronics, photonics, nanotechnology, high-temperature superconductors, novel wide band gap semiconductors, and magneto-calorics. CCP9 provides a network that connects UK research groups in electronic structure, facilitates UK participation in the larger European Ψ k Network, and supports a number of cutting edge computational codes.

Summary Report (1 April 2020 – 31 March 2021)

Wannier90: One of the major tasks within the new CCP9 which started on April 1st is the Wannier90 library project, with the CoSeC team, consisting of Jerome Jackson, Barry Searle, Martin Plummer, and Leon Petit, working in close collaboration with the code developers, Arash Mostofi (Imperial College

London) and Jonathan Yates (University of Oxford). The goal of this project is to create a parallelized and thread safe library version of the Wannier90, a code that is already interfaced to many community codes and is used to derive advanced materials properties. The parallelization of the library will result in a speed-up and facilitate high throughput calculations. Interfacing to Wannier90 in library mode will make it more accessible to a larger part of the community, and scientists will eventually be able to access the library through the CECAM electronic structure library (ESL).

The first phase consisted of a thorough investigation of the scaling and performance of the existing code, including testing of compilers and performance on Archer. Benchmarking calculations have been run on a number of testcases that were devised for parallel profiling. The second phase, restructuring the library calls, has progressed well, with the code having improved workflow and readability. A whole range of variables no longer needs to be duplicated and have been moved from the module being passed directly now through the argument list. This major milestone is in the process of being committed to the public repository. The type classification, with variables being collected in new datatypes is ongoing. The CoSeC team has weekly discussion sessions, and progress/future tasks are discussed during monthly zoom discussions with the Wannier90 developers.

Materials Science: In collaboration with University of Leeds and King's College, Jerome Jackson has been studying magnetism in Yttrium iron garnets. Their work serves as proof of concept that the QSGW methodology, as implemented in the CCP9 flagship code QUESTAAL, can be run highly efficiently for large systems > 80 atoms. Yttrium iron garnets play a key role in advances such as spin-caloritronics and superconducting magnetic qubits, and the description of correlations in these complex materials at the level presented here has not have been achieved in earlier studies. This milestone has major implications for the scope of QUESTAAL. The study has recently been published as invited paper in connection with a special issue of the IOP journal Electronic Structure. Kun Cao has been involved in a combined theory/experimental study of magnetic excitations in transition metals Fe and Ni via resonant inelastic X-ray spectroscopy (RIXS). Using time dependent density functional perturbation theory he computed the corresponding magnon spectrum in support of the experimental study. The work has been published in PRB. As part of the on-going collaboration with Ames laboratory and Warwick University, Leon Petit has been studying the electronic structure and magnetic phase diagrams of Eu₂In. The material is characterized by giant magnetocaloric effect under a magnetic phase transition that shows however no thermal hysteresis. This kind of materials are very valuable for technological applications such as magnetic cooling. The study finds the origin of the effect to be electronic in nature, and has resulted in a publication in PRB. A new collaboration of the DL group (Z. Szotek and L. Petit) with colleagues from India looked into the electronic structure of YbCdSn. Using a combined theoretical and experimental approach, this work lays the foundations to a fundamental understanding of the material properties, eventually leading to fit-for-purpose material design in fields such as quantum computing and low-consumption spintronic devices.

CRYSTAL: Barry Searle has finalized the OpenMP parallelisation of Electron Nuclear forces. This feature is part of the CRYSTAL21 beta pre-release code. The OpenMP parallelization of CRYSTAL presents a major upgrade of the code. An important aspect of this work has been the code modernisation necessary to support OpenMP, which also improves the code maintainability for the developers. The code which has been supported by CoSeC and used by the CCP9 community for many years is now ready for the next big release planned for late 2021. Continuing work beyond the CRYSTAL21 release, Barry has furthermore implemented the OpenMP extension to the spin-orbit part of CRYSTAL and has been benchmarking the code on Archer 2. Barry has been teaching at the CRYSTAL hands-on course, September 2020, with 50 people being registered for this on-line edition of the course.

The CoSeC team had plans for 4 additional hands-on courses all of which however unfortunately had to be postponed due to Covid19. An online Community meeting was held in November 2020.

Significant Impact (1 April 2020 – 31 March 2021)

Yttrium Iron garnet calculation with QSGW as implemented in CCP9 flagship code QUESTAAL. This calculation was attempted in 2018 and was found to be nearly infeasible on the available hardware due to excessive memory, IO and computational requirements. Algorithmic improvements, together with a more flexible memory management allowed efficient parallelization across multiple levels of processes and threads, utilising a number of multi-GPU accelerated nodes. The effort culminated with the ability to obtain the single cycle self-energy in under 10 minutes on an Nvidia DGX-1 node with 8_V100 cards. Typically convergence of the QSGW potential is achieved in < 15 iterations. Due to queuing system restrictions the presented results were performed on 16 Marconi100 nodes totaling 64 Nvidia V100 GPUs, reaching peak performance of 425 TFLOPs. As a result Loads of time, 46 million CPU hours from PRACE for continuation of project on rare-earth garnets (JJ co-PI). <https://arxiv.org/abs/2009.14601>

Outline Plans 2021/22 (1 April 2021 – 31 March 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

With respect to the Wannier90 project the aim is to have delivered a fully functional and tested library to the community. With respect to QUESTAAL, the BSE spin susceptibilities will be incorporated and work will start on the DMFT verification. CRYSTAL parallelization will be extended to the Coupled-Perturbed Kohn-Sham integral functions for phonon intensities and dielectric constants. Various beyond DFT materials studies will be published.

CCPNC – NMR Crystallography

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.

Summary Report (1 April 2020 – 31 March 2021)

The core results delivered during this time period has been the completion and opening to the public of the CCP-NC magres file database, which will provide a central hub for computational NMR data storage and access from now on. The database has been completed and made public at <https://www.ccpnc.ac.uk/database/> in September 2020, after a major refactoring of the code base in order to increase maintainability and satisfy more of the user demands. While maintenance is ongoing, the core software infrastructure is now in place and running. This database can be a great asset to the UK's NMR computational community, reducing the need for duplication of efforts and giving everyone an easy way to make their data publicly available, as is best practice for the sake of reproducibility and sustainability.

In addition to the database, a lot of work was done in collaboration with Paul Hodgkinson on the development of tools and scripts for the treatment of the problem of the averaging of dipolar couplings in solids in presence of molecular motions. These have resulted in a number of additions to the Soprano library as well as standalone scripts. The main goal is to help managing simple cases of dynamical disorder in NMR when studying in particular organic molecular crystals. This is particularly relevant for the NMR groups in Durham and Warwick as well as other partners involved e.g. in pharmaceutical research. Some scripts are now collected in the GitHub repository

https://github.com/stur86/dipolar_averages; in particular, vanVleckCalculator.py is a script that satisfies the original need for dipolar coupling averaging in disordered or dynamical systems and is now being checked against experiment.

Development of the 3D crystal rendering app crystvis-js is nearly complete. The app is now in a usable state and will soon be made available on the NPM repository. It will constitute the core of the MagresView 2.0 web application.

The UI of MagresView 2.0 is being developed and tested, with various frameworks being explored as possibilities. A non-functional template will be prepared and assessed before moving on to developing the final software.

Significant Impact (1 April 2020 – 31 March 2021)

The release of the CCP-NC database is a key step for the computational NMR community, as it provides a central reference point for those who produce NMR data with simulations to store it and make it public, and for those who need it to look for precedent results and avoid duplication of efforts. It is our intention to make it a hub that serves the purpose to consolidate the computational NMR community as well as encourage the use of the Magres file format.

Outline Plans 2021/22 (1 April 2021 – 31 March 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

While plans for 2021/22 are not finalised (the objectives are set at 6-monthly Steering Group meetings), the proposed work for 2021/22 is aligned with the work packages of the new funding period for CCP-NC. Maintaining and updating the core software supporting the community - mainly MagresView, Soprano and the new CCP-NC database - remains a key mission of the project. In particular, the development and deployment of MagresView 2.0 remains a key goal to pursue. Moving beyond that, the CCP-NC aims to expand by hiring a new member in the upcoming year, which would increase its ability to work on multiple different software and research projects. The focus will be on more software, scripts and methodologies to treat NMR of dynamical and disordered systems with computation.

CCPQ – Quantum Dynamics in Atomic Molecular and Optical Physics

The remit of CCPQ is to develop code for quantum dynamics simulations, solving the time-dependent Schrödinger equation for both heavy (nuclei) and light (electrons) particles. It has 3 main community codes covering different areas; 1) TNT: A package coding Tensor Network Theory for coherent many-body nuclear dynamics, 2) R-Matrix Suite: A set of programs for electron (positron) -atom and -molecule scattering, (ultrafast) laser pulse interactions and related problems. The package is also currently being adapted for ultracold molecule-molecule (atom-atom) interactions, bound states and resonances (RMAT_REACT), mainly through HEC UK-AMOR but with CCPQ CoSeC support and 3) Quantics: A package based on the MCTDH algorithm for molecular quantum dynamics.

The collaboration has also had related code development in the field of coherent dynamics in the CCE (Cluster Correlation Expansion) code and in the field of anti-matter scattering related to the R-Matrix work. It also supports codes and research in low-energy antimatter interactions (positron and antihydrogen collisions with atoms and molecules), with input in this area from Aberdeen and Belfast, the latter employing many-body-theory diagrammatic techniques.

Summary Report (1 April 2020 – 31 March 2021)

Please note: CCPQ was not renewed as part of the 2019 EPSRC CCP call. However, there were a number of activities that were able to continue using underspent effort from previous years.

In 2020-2021 CCPQ has residual CoSeC funding of 0.12FTE after it was not renewed in EPSRC's 2019 CCP Call. For practical purposes, for work besides basic administration, this was combined with the support for HEC UK-AMOR and the highlights reported in the UK-AMOR report are joint for UK-AMOR and CCPQ. For completeness these are repeated here:

The main coding highlight for this reporting period was the completion of Andrew Sunderland's memory optimization work for the UK-AMOR code UKRMol+ serial and OpenMP-threaded GBTOlib library as requested by the developers. This highly optimized code is now applicable to a wide/full range of electron-molecule applications and is part of the UKRMol+ GIT repository (final debugging and testing of the optimization extended to the new MPI parallel version of the library was delayed due to COVID-related disruption). Andrew Sunderland has also obtained 3 person months work from PRACE to work on PFARM (part of the Prace 6iP Applications benchmark suite: <https://repository.prace-ri.eu/git/UEABS/ueabs>) on several international Tier 0 platforms between now and October 2021.

Another highlight is the award of an ARCHER2 eCSE project to work on UK-AMOR's RMT code for laser atom/molecule interactions (Martin Plummer will be the main worker, the PI is Dr Andrew Brown, QUB and UK-AMOR). This work, commencing in October, will improve the parallel efficiency of RMT for the large-scale jobs enabled by the introduction of arbitrarily polarisation of the laser (and future planned new functionality).

Martin (as planned co-I) supported Dr MM Law (Aberdeen) in an ambitious New Horizons proposal which would extend the RMAT_REACT framework and code-base to treat antimatter rearrangement collisions (antimatter chemistry) relevant to low-energy experiments planned by the ALPHA project and others at CERN when work recommences after the May 2021 restart. While this was not successful it provides a basis for a responsive mode application for a new 3-year project. This new application will now be underway in Q2 2021.

While the second half of the reporting year has mainly been consolidating development work relating to the eCSE project and the AquA-DIP grant, a major highlight has been the revival and reorganisation of the joint UK-AMOR/CCPQ and HEC/CCP-PLASMA workshop, 'Atomic and Molecular Data Needs for Plasma Applications', postponed from April 2020 (with Jonathan Tennyson joining the organisers). This three day online workshop, to take place on 13-15 April 2021 with ~160 registered participants, will have 23 talks, 3 informal discussion sessions, a demonstration by Quantemol Ltd and a poster session (a lightning introduction session and a more detailed as-live breakout rooms for each poster). The posters and talks will be preserved on the CCPQ website.

Finally, Dr Andrew Brown (QUB) has been awarded a Research Software Engineer fellowship to unify diverse AMO codes (concentrating on the various 'R-matrix' packages) into a multipurpose effective single resource across the range of AMO physics. His proposal was supported by UK-AMOR generally and by CoSeC support (practical commitment from MP) in particular and will be a major continuation and expansion of the CoSeC/SLA code support and curation work given over the years for CCPQ and UK-AMOR.

Significant Impact (1 April 2020 – 31 March 2021)

We repeat here the Impact statement from the UK-AMOR report, which counts as joint for UK-AMOR and CCPQ:

Impact of CoSeC support work in this period will take longer than the reporting year to appear.

The main immediate impact is the new eCSE funding for work on UK-AMOR's RMT code. [The earlier CCPQ/UK-AMOR work on adapting PFARM for RMT_REACT will have impact in the form of new collision calculations in the near future, while longer term impact is expected from the work on the bound-state code.] The earlier work on PFARM (within CoSeC and as part of PRACE) has resulted in new PRACE funding for AGS and the PFARM code being part of the Prace 6iP Applications benchmark suite.

CoSeC support enabled Martin Plummer to be involved with a successful 3-year EPSRC grant application (the 'AQuADIP' project) which commenced in May this year, with PDRAs based at QUB and UCL, linking RMT with a semi-classical approach. A result of this grant in this reporting period was the international workshop 'Quantum Battles in Attoscience, co-funded by CECAM, which switched from being an in-person event to an online event (organised by Prof Carla Faria, UCL and colleagues with some help from QUB) and was extremely successful. It had much wider international participation (>300 attendees) than the in-person event would have had, as well as a highly successful disciplined approach to discussing opposing viewpoints and approaches to aspects of attosecond science (with follow-up articles on this approach to organisation in Physics World). The workshop is also an excellent template for future workshops in a post-COVID environment.

In particular the 'Atomic and Molecular Data Needs for Plasma Applications' workshop, with ~160 participants, should have strong impact in strengthening and forming links and collaborations between the AMO and Plasma physics communities.

Andrew Brown's well-deserved RSE fellowship will involve work that follows on from general CoSeC/SLA support over past years and the present to develop and unify the AMO (R-matrix in particular) code packages for wider use. AB told MP that the fellowship was "... a recognition of the value of AMO software in general, the R-matrix codes specifically, and their potential impact over the long term. You have been such a key part in maintaining the community and the codebase. I am excited to work with you over the next five years to create something that will deliver benefit for the R-matrix community and AMO physics more generally."

Thanks to Andrew Sunderland's CCPQ support and follow-up PRACE work, the PFARM benchmark suite, available for public download from the Unified European Accelerator Benchmark Suite repository, has been restructured with the CPU and GPU versions streamlined (now both versions are included in the one source code download and configuration). Scripts to automate benchmarking of the atomic and molecular datasets have been published and the accompanying documentation has been revised and expanded. Testing of the benchmark by third-parties is currently underway, shortly leading to a new release of the UEABS.

Outline Plans 2021/22 (1 April 2021 – 30 September 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

The remaining residual CoSeC support for CCPQ ends in March 2021 (while CCPQ goes into hibernation). Thus there are no official milestones. MP will maintain the CCPQ website as part of ongoing UK-AMOR support (UK-AMOR's website is run from UCL). As part of the limited UK-AMOR CoSeC support, and through other current funding sources, MP will concentrate on UK-AMOR projects and further funding attempts in the atomic, molecular and optical collisions physics areas, including low-energy antimatter, of CCPQ that are related to UK-AMOR. The academic members of CCPQ involved in strongly correlated systems (the TNT code and similar work) are involved in various Quantum Computing related areas and projects/networks, including CCP-NC. Andrew Brown's RSE Fellowship is a major boost for continued general curation and user-friendly expansion of the UK AMO codebase.

CCP-Plasma – HEC-Plasma Physics

CCP-Plasma includes researchers from UK universities, the Culham Centre for Fusion Energy (CCFE), the Rutherford Appleton Laboratory (RAL) and AWE. The study of plasma physics covers a huge range of scales and applications. It is core to the development of laboratory experiments such as fusion power, new light sources and the next generation of particle accelerators. On the largest scales it is fundamental to our understanding of astrophysics.

CCP-Plasma was established in 2007 with the aim of pooling the collective expertise across these disparate subjects and developing core plasma physics simulation codes, and training packages, for UK science. CCP-Plasma supports three core codes EPOCH, BOUT++ and GS2. We also have one Flagship project for the development of a radiation-hydrodynamics ALE code for laser plasmas (Odin). 1FTE of CoSeC core support is split evenly between the codes GS2 and BOUT++, which both focus on modelling plasma in magnetic confinement fusion devices. This CoSeC support is funded 20% from the HEC-Plasma Consortium until 28/5/2023.

The Plasma High-End Computing (HEC) Consortium is an EPSRC funded collaboration of UK plasma physics researchers. The Plasma HEC supports research in the simulation of plasmas, including those plasmas of relevance to magnetic confinement fusion, laser-plasma interactions, and laser fusion energy. The software development includes a commitment to optimizing key codes, with input from EPCC, RAL and Warwick Computer Science, and developing new physics packages as required to maintain the UK's role in laser-plasma physics and all approaches to fusion for energy. The consortium supports meetings and software development. It also manages a block allocation of national super-computer time to be used for plasma physics research.

Plasma Summary Report (1 April 2020 – 31 March 2021)

Please note: CCP-Plasma was not renewed as part of the 2019 EPSRC CCP call. However, there were a number of activities that were able to continue using underspent effort from previous years.

CCP/HEC-PLASMA is currently without official support as Dr Joseph Parker is now employed by Culham Centre for Fusion Energy on the ExCalibur project. Dr Martin Plummer (UK-AMOR and CCPQ) is maintaining contact with Dr Ben Dudson and Dr Colin Roach following joint organisation of a (postponed due to COVID) joint workshop on atomic, molecular and optical (AMO) physics data for plasma physics. Organisation of this workshop was revived in the 2nd half of the year, now as an online workshop (with more than 120 participants) to take place 13-15 April 2021. Contact will continue in the next reporting year regarding possible further AMO work for HEC-PLASMA, and contact will also be kept with Joseph Parker, who is continuing some work related to CCP/HEC-PLASMA CoSeC work in his new post. Thus some development and testing work has continued on ARCHER in the c01 account (highlighted below). Continued limited ARCHER c01 time was successfully negotiated with EPSRC by MP to allow CoSeC related PLASMA computations to continue (on ARCHER and ARCHER2).

Simulations of the tokamak edge require the inversion of elliptic operators to obtain the electrostatic potential from the vorticity and density. Work on operator inversion in BOUT++ in the last reporting cycle allowed CCFE to undertake realistic, fully 3D simulations of turbulence in tokamak plasmas. However, the inversion remains the performance and scaling bottleneck. This is because the 3D solver is based on iterations in two dimensions, with a direct solver used in the remaining dimension. The direct solver requires all-to-all communication, and therefore does not scale. We have developed a novel iterative solver that combines inversions using the Thomas algorithm local to a processor, with a multigrid solver to solve a reduced system that connects the processors. Our algorithm is faster than cyclic reduction (the previous best algorithm in BOUT++) and retains good scaling efficiency to twice as many cores. Improved scalability will allow realistic simulations of tokamak turbulence to utilize more cores, significantly reducing time-to-solution and allowing code users to perform more high-fidelity simulations.

Significant Impact (1 April 2020 – 31 March 2021)

Our new Laplacian inversion algorithm improved the scalability of BOUT++, allowing realistic simulations of tokamak turbulence to utilize more cores. This significantly reduces the time-to-solution and allows code users to perform more high-fidelity simulations.

The April 2021 joint meeting (with UK-AMOR/CCPQ): “Atomic and Molecular Data Needs for Plasma Applications” will have major impact to be reported on next year (it has over 120 registrations and will bring plasma and AMO scientists together to improve realistic modelling and understanding of processes in plasmas).

Outline Plans 2021/22 (1 April 2021 – 30 September 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

It is not yet possible to provide a summary plan unless further staff effort is allocated to HEC Plasma. Martin Plummer is fully booked with a mixture of eCSE/EPSRC grant/UK-AMOR CoSeC work in 2021-2022 (with a maximum of 0.05FTE possibly available and currently to be allocated to CCP9 support), but will maintain contact with Ben, Colin and Joseph to build on the links forged by the joint workshop, either as part of UK-AMOR support or more specifically for HEC-PLASMA if the 0.05FTE is more better spent on this project, and if there is useful AMO related work that can be done as a result of the workshop. If effort currently allocated to CCP-PLASMA cannot be carried over to HEC-PLASMA for 2021-2022, then new employment of someone directly related to plasma work is unlikely.

CCPi – Tomographic Imaging

The CCPi was established in 2012 to support the emerging UK tomography community with a toolbox of algorithms to increase the quality and level of information that can be extracted by computed tomography. The size of this community has continued to grow with many academic groups around the UK taking up tomographic imaging and purchasing new lab-based x-ray CT scanners. In 2012 there was an estimated 50,000 CT imaging sources around the world. The size of our community has arisen from ~250 in 2013 to just over 400 in 2018.

Our focus aims to bring together the UK imaging community in maximising the return on investment in imaging software through developing, maintaining, sustaining and prompting the CCPi core imaging toolbox. Primarily we concentrate on advancing the state-of-the-play of the field through: 1) reducing the barrier to access in multi-modality image analysis algorithms, be the data coming from mid-range facilities hosted by the universities or from large scale synchrotron or neutron facilities; 2) improving the accessibility and distribution of the codes; and 3) establishing a national multidisciplinary image analysis focal point for the multidisciplinary community comprising of algorithm developers, material scientists, instrument manufacturers, and instrument scientists.

The CCPi core staff actively involves in three major annual imaging events in the country, each having 50+ attendees. These include an X-Ray user group symposium (ToScA) managed by the National History Museum (NHM) and Royal Microscopical Society (RMS), a technical forum supported by RCaH and DLS; and a “dimensional XCT” conference supported by NPL that is leading to formal BSI/ISO standards.

Summary Report (1 April 2020 – 31 March 2021)

The team is focusing on the development of the Core Imaging Library (CIL) and on the publication of two peer reviewed articles on it. CIL is an open-source library for the processing of tomographic imaging data. CIL provides a variety of iterative regularised reconstruction methods which enable the

reconstruction of highly noisy, incomplete and non-standard tomographic data sets for which conventional methods and proprietary software fail to produce satisfactory reconstructions. However, iterative methods are inherently more complex to utilize but potentially produce high quality results. This is the first attempt in the UK to improve the take-up of iterative reconstruction algorithms by the community through software development and training courses and outreach. This challenge is being achieved by: simplifying the software installation process; improving documentation; releasing sample code and demos; as well as making group specific readers and writers available. During the reporting period there have been four major CIL releases version (20.04, 20.09, 20.11, 21.0) adding functionality. With CIL we aim at reaching more scientists/users by simplifying the installation and distribution and by providing a simple Python interface to a variety of codes and methods that have been developed within CCPi. Additionally, the CCPi software suite covers other areas such as algorithms for pre-processing, quantification, segmentation, and visualisation.

Current development in tomography beamlines in synchrotrons, neutron sources and lab machines is focussed on 4D/5D tomography. This includes 3D imaging of a sample while changing its status (e.g. its temperature) and multi-channel (or spectral) imaging, where sources of different energies can be selectively acquired. This means that there is a push for faster acquisition and energy selective imaging which inherently reduce the signal-to-noise ratio, SNR. For instance, the ISIS/IMAT beamline can resolve up to 2300 different energies of the impinging neutrons, dramatically reducing the statistics of the collected data per energy with regard to the “white beam” acquisition. In X-ray imaging the dose delivered to the sample may change its characteristics. That is why it is necessary to develop methods that are resilient to low SNR data acquisition, which is something the users are going to be interested in. Also, there is a trend in acquiring dynamical data. In this setting we expect to be imaging the sample multiple times; in this case the acquired data will suffer from low SNR. Again methods that are resilient to low SNR are necessary.

In the reporting period the main additions to the CIL software are:

- Redesigned Acquisition Geometry which allows to go beyond standard CT trajectories (circular scans). The main consequence of this is that it is possible to use CIL for laminography, or for center of rotation correction for both parallel and cone beam configurations.
- The development of the Stochastic Primal Dual Hybrid Gradient (SPDHG) algorithm for minimisation, which has the potential to be more memory efficient and faster for reconstruction.
- Added readers for data acquired by NIKON and ZEISS machines, added tools for pre-processing data, e.g. normalisation, transmission to absorption, added plugin for the TIGRE engine.
- Upgrade of all versions of dependencies, mainly to support the latest release of DLS Savu (2.4).

In the reporting period, a large effort was put in the preparation of four publications regarding the CIL, two describing the framework and two using the framework to analyse data from ISIS/IMAT and the Manchester MXIF Colour bay. These papers should be part of the proceedings of the Synergistic Reconstruction Symposium 2019 jointly organised by CCP PETMR and CCPi.

The team has started to support Queen Mary University London in their reconstruction pipeline. Regular meetings are taking place to monitor progress. The team is also collaborating with ISIS IMAT and MantidImaging teams to integrate CIL into their reconstruction pipelines, with regular meetings to monitor progress.

Efforts have been spent to strengthen the compatibility of the CIL with the CCP SyneRBI software SIRF. A joint hackathon of the two CCP's lead to the addition in the CIL of the SPDHG algorithm, published in 2018, which enables both to increase efficiency and to reduce the extreme memory requirements of tomographic applications.

Another important outcome of the collaboration between the CCP's is the use of optimisation algorithms available in CIL to tackle challenges in synergistic medical imaging reconstruction, such as

compensation for motion (e.g. respiratory, cardiac). Specifically, we have employed three different algorithms with regularisation for motion compensated cardiac PET/MR image reconstruction (for which a paper is in preparation).

The team has been busy in the organisation of the XCT-fringe meeting and 2-day Turing summer school at RAL and the fourth annual workshop on Advances in X-Ray Imaging (UoM, 3DImagination, DLS) at DLS; in September ToScA. However these have since been cancelled/postponed due to the pandemic.

Significant Impact (1 April 2020 – 31 March 2021)

Reconstruction of the multispectral imaging from the ISIS/IMAT (neutron) beamline and the Manchester MXIF Colour Bay (x-ray) with iterative reconstruction and regularisation. Thanks to the CIL the regularisation can be applied separately on the spatial and on the spectral dimension with significant improvements in terms of spatial resolution and Bragg edge fitting. Two papers are being written describing the new technique of applying regularisation on spectral and spatial dimensions separately and the comparison with standard techniques.

The new (2018) Stochastic Primal Dual Hybrid Gradient algorithm has been developed and made available to the community. This algorithm can reduce the computational requirements with respect to other algorithms, leading to faster reconstructions.

The CILViewer updates allowed the development of a GUI for the CCPi Digital Volume Correlation code, reducing the barrier for its use. The CIL benefits the CCP SyneRBI (CCP PETMR) community and makes possible to reconstruct with motion compensation, via its algorithms and regularisation, both with PET and MR separated and jointly. This is the subject of another publication being prepared in conjunction with the CCP SyneRBI team/community.

Outline Plans 2021/22 (1 April 2021 – 30 September 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

During the years 2017/2020 a large effort has been put in the new development of the new iterative reconstruction framework CIL, also in collaboration with CCPi flagship. We expect a period of lower development activity, more focused on deployment at university labs and facilities like DLS, ISIS/IMAT and CLF/EPAC. The development and interaction with the CCP SyneRBI is likely to continue and get stronger.

Regarding new development, we will focus on application of iterative reconstruction methods at facilities as well as NRF on XCT, visualisation, remote analysis technology and code efficiency/optimisation.

Additionally, we plan to start to investigate methods based on ML/AI for computed tomography. Currently under preparation together with UoM is a project to study coronavirus with CryoEM data and the use of CIL for reconstruction.

CCP SyneRBI - Synergistic Biomedical Imaging

For medical imaging, the UK is a globally leading country. It has the highest number of medical imaging machines per capita in the world, evenly spread throughout the 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 until 2025, aims at bringing together the best of the UK's imaging expertise to capitalise on the investment in this area. New research shows that the use of MRI intermediate results can improve PET imaging quality and vice versa, and latest scanners can acquire MR and PET data simultaneously. Our CCP is dedicated to exploiting exciting new capabilities that the synergy of MR,

PET and other imaging modalities can deliver. 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.

STFC CoSeC support for this CCP currently focusses on developing the SIRF code base that provides an easy-to-use script-language (Python and Matlab) 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..

Summary Report (1 April 2020 – 31 March 2021)

Despite COVID-19, our work during the reported period mostly progressed according to the job plan, with some delays and with adjustments to avoid face-to-face meetings. We continued our software development and engineering efforts, adding content to our website www.ccpsynerbi.ac.uk, maintaining our steadily growing mailing lists (we now have 100 members on the announcement list, 38 on the developers and 78 on the users’ lists), organising meetings and Hackathons.

On 8 June 2020, we published Release 2.2 of our Open Source software suite SIRF, and moved on to working on SIRF Release 3.0, originally planned for March 2021 (but delayed with 1 month). These releases are based on the PET reconstruction package STIR (Software for Tomographic Image Reconstruction), the MR reconstruction package Gadgetron and NiftyReg for image registration. SIRF is now capable of processing measured data from the Siemens PET-MR scanner, with work-in-progress to support the GE PET-MR system. Major new features of Release 2.2 are the use of GPU for PET reconstruction and conversion of ISMRMRD data to NIFTI format.

On 29 June 2020, we held our 26th Software Framework Meeting followed by our 6th Hackathon, in which we investigated the use of advanced optimization algorithms from Core Imaging Library in SIRF. Our 27th Software Framework Meeting was held on July 6th, followed by our 7th Hackathon, in which STIR PET scanner support was investigated. These meetings were held online only which mainly required adjustments to the format of the hackathon. As opposed to a few concentrated days with in-person meetings, the virtual hackathons were spread over 1 week with regular meeting check-points, communications via Slack, Teams etc. Feedback on this new format was very positive.

On 11 February 2021 we held a joint meeting with UCL developers for XNAT, an open source imaging informatics platform that facilitates common management, productivity, and quality assurance tasks for imaging and associated data developed by the Neuroinformatics Research Group at Washington University. Thanks to its extensibility, XNAT can be used to support a wide range of imaging-based projects. The purpose of this joint meeting was to discuss the integration of the SIRF software for image reconstruction. Follow-up meetings are being planned.

Significant Impact (1 April 2020 – 31 March 2021)

Our Hackathon training sessions demonstrated that SIRF allows biomedical researchers to implement a new reconstruction algorithm and test it on real scanner data within days. Without SIRF, this would take months even for a person with advanced programming skills, which very few biomedical researchers possess. The acceleration of reconstruction algorithms development will translate into faster progress in accuracy of scanners’ images, which will lead to better quality and longevity of life for people affected by cancer, dementia and other serious illnesses.

Outline Plans 2021/22 (1 April 2021 – 31 March 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

In 2021/22 we plan to make further additions to SIRF, in particular, functions to compute gradients and values of MR objective functions, which will be needed by MR iterative reconstruction algorithms. We plan to implement, together with CCPi CIL team, a few generic optimisation algorithms, and tools for working with subsets of acquisition and image data. We will complete and transition work for on motion correction and estimation (CCP PET/MR Flagship) into SIRF. Our SIRF Release 4.0 is planned for the end of 2021. We will start with the integration of our tools into the XNAT platform (with help from UCL) to enable translation towards clinical research. Our integration with the CCPi software continues with several algorithms currently being added in CIL but tested with SIRF. Finally, we plan to investigate integration of other reconstruction packages and installers with precompiled software (conda).

Together with CCPi, we are in the process of organising our next training school in the summer 2021, associated to the 16th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine, also known as Fully3D. Due to COVID-19, this training will be online only, and spread over 3 weeks. See <https://www.ccpsynerbi.ac.uk/SIRFCIL2021>. We are investigating giving attendees access to the STFC Cloud for running real-sized problems, with data and software preloaded and access via JupyterHub with resources managed with Kubernetes.

To prepare for this training school we will hold our 7th (virtual) hackathon in the week of the 19th of April.

CCPBioSim - Biomolecular Simulation at the Life Sciences Interface

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.

From April 2020 onwards CCPBioSim is supported by 1.95 FTE of SLA core effort. Of this 1.5 FTE is allocated to core software development, including 1.0 FTE towards a portfolio of short, community-driven software projects and 0.5 FTE for longer-term maintenance of CCPBioSim software and flagship projects. 0.25 FTE supports coordination of SLA effort, website maintenance, mailing list management and networking activities, while the remaining 0.2 FTE is for infrastructure support including our cloud-based training platform. For further information please see www.ccpbiosim.ac.uk.

Summary Report (1 April 2020 – 31 March 2021)

The P450 multiscale project has achieved most of its objectives, with coarse-grained and atomistic Molecular Dynamics simulations of the membrane-bound P450 system completed, looking at interactions with the drugs warfarin and dapsone. The project to re-implement the protocol was completed in October and publications on the MD and CG aspects are in preparation, while the final QM/MM aspects of the project are being finalised following porting to ARCHER2.

A new 2.5 year project project to incorporate Fluctuating Finite Element Analysis (FFEA) into the Code_Saturne package began in April, with the aim of significantly enhancing the robustness and parallel scaling of the method. The Compatible Discrete Operator version of the code is used for developing the structural mechanics module in Code_Saturne. This part of the code does not handle moving meshes and the reconstruction of the geometric quantities, after mesh deformation has been added to this part of the code and tested. The operators to solve the soft matter equations are added one by one, starting from the computation of a vectorial Laplacian, which will be key to computing the viscous stress.

In December, work began on a project to interface Enlighten (a tool for setting up enzyme simulations) and Py-ChemShell (for running QM/MM simulations). This will provide easy-to-use scripting and graphical workflows to reduce the learning curve for users who are new to multiscale QM/MM modelling. The project is progressing in line with the workplan and is due to complete at the end of May.

FESetup continues to be maintained on behalf of the CCPBioSim user community and supported through the FESetup mailing list. The glycosylated proteins project (with CCP4/CCP-EM) was paused because of the Covid-19 pandemic. Our online training material continues to be available. A new tutorial about equilibration was prepared by Charlie Laughton for the October Training Week and will be added to the online collection.

The CCPBioSim Training Week was held online 2-9 October 2020. The topics covered ranged from the basics of biomolecular simulation to QM/MM calculations and there were talks from CCP5 and CCP-EM. There were 17 sessions including three plenary research seminars by Syma Khalid (University of Southampton), Prem Chapagain (Florida International University, USA) and Viv Kendon (Durham University; Chair of CCP-QC). This was the first time CCPBioSim has run a completely online training event and everything went smoothly including the scaling up of our containerised cloud resources. We would like to thank Georgia Lomas (STFC, SCD) and Jonathan Oldfield (STFC, DI) for their assistance with setting up and running the online sessions.

SlimMD was developed by Sarah Harris (Leeds) and James Gebbie-Rayet (STFC) with contributions from others in the community and is now accessible from the CCPBioSim website. This is a database of molecular dynamics trajectories in standard formats that are small enough to be easily downloaded and manipulated. This will allow undergraduates and other students to work on projects involving visualisation and analysis of simulations without having to generate all the data.

The 4th Manchester Multiscale Conference was held online 29 – 31 March 2021. This event was joint with CCP5 and sponsored by CECAM-Daresbury Node. There were 9 invited speakers, 11 contributed talks and 27 posters. The invited speakers were Maria Fyta (Stuttgart, Germany), Frauke Gräter (Heidelberg, Germany), Syma Khalid (University of Southampton, UK), Christopher Horst Lillig (Greifswald, Germany), Céline Merlet (Toulouse, France), Irina Paci (Victoria, Canada), Lars Pastewka (Freiburg, Germany), Anđela Šarić (UCL, UK), and Rebecca Wade (Heidelberg, Germany). We had 393 registered participants.

During the conference a ChemShell community engagement event was held, with Tom Keal giving a talk on recent progress on the software developments, including the new biomolecular modelling workflow, together with an online Q&A session.

Significant Impact (1 April 2020 – 31 March 2021)

Due to the Covid-19 pandemic the CCPBioSim training week was held online for the first time and was a great success. Holding the training course online significantly expanded the potential audience, with 450 people registered and 187 logged in at any one time. Attendees came from the UK, the rest of Europe, Asia, Africa, North America and South America (36 different countries).

Videos from the training week and our other online events are posted on our YouTube channel (<https://www.youtube.com/channel/UCJLWsk2Hbhf09--POIR67Nw>) which has over 80 subscribers. This allows people who missed the training and/or talks to watch or re-watch them at their own convenience and should improve the accessibility of the learning material.

Another major highlight of the year was the 4th CCP5/CCPBioSim multiscale modelling conference, with 393 registered attendees from around the globe. The fully online event was held as a Zoom webinar for the talks and a series of breakout rooms for the poster session, all of which worked very well.

Outline Plans 2021/22 (1 April 2021 – 31 March 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

The plans for 2021/22 include working on several of CCPBioSim's new short projects. These projects are focused on tools with the potential for high impact for the biomolecular simulation community. For each proposed project there is a form (inspired by the applications for eCSE projects on ARCHER) which includes an overview of the scientific project, a proposed work plan and information on the future outlook/sustainability. These will then be evaluated by the management group and projects selected based on their potential impact for our community.

Our training programme will continue with more online events along with our self-guided online tutorials. There will also be ongoing support and maintenance for existing projects.

CCP WSI+ - Wave Structure Interaction

The CCP is focussed on the scientific area of Wave Structure Interaction (WSI). Primarily this is an engineering topic (both computational and experimental) but crosses into other areas such as chemistry, physics and natural (Earth) sciences. This CCP is a direct continuation of the existing CCP-WSI community, with the "+" representing the fact that the CCP now includes additional communities to directly consider the solid mechanics and other non-fluid aspects of WSI. Key topics revolve around scientific code coupling, high-performance computing, fluid dynamics, solid mechanics and solid body kinematics with a focus on community dissemination of methods, software and results with repeatability in mind. Application areas of particular interest include understanding wave-energy generation devices, flood defences and fundamental understanding of the physics of WSI.

Summary Report (1 April 2020 – 31 March 2021)

Please note: CoSeC support for this CCP begun on the 1st of October 2020, so this report relates specifically to the 6 months from then.

The first iteration of the CCP-WSI community had no CoSeC support associated, therefore a key task for the renewed community was to understand exactly how the 2.0 FTE of support provided to it would be best used and how the community would need to change in order to incorporate this new aspect. While the approved proposal for the CCP includes a clear top-level overview of three work packages for CoSeC a key first task was to refine these and define specific pieces of work within them.

Work packages two and three (high-performance computing and sustaining community output) are based on existing working relationships with STFC staff who were involved in a non-CoSeC capacity with the last CCP-WSI project, therefore they are able to continue in the same vein based on existing activities. The first work package (code coupling) represents a new and targeted approach to support disparate activities across the CCP. In order to collate existing activities and understand what the grand goals of the CCP are (as well as specific activities that should be targeted first), a CCP-wide survey exercise was performed, leading to a document that provides a representative and highly refined high-

level view of what the CCP would like to achieve through CoSeC over the coming 4.5 years, as well as a more immediate set of tangible tasks to help achieve these goals.

Fundamentally it has become clear from this exercise that the overarching goal for work package 1 is to create a modular coupled multi-physics framework to simulate problems related to WSI that incorporates not only the existing solvers produced within the CCP but that also allows integration for other external solvers in order to provide new and complementary physics-solving capability.

An immediate task revealed by this exercise was to produce a working coupled framework that specifically involves a key structural code for the new community called ParaFEM. The primary goal is for a working framework able to simulate practical problems to allow the community to explore its capability for themselves. Past work to integrate ParaFEM with part of the CFD suite OpenFOAM produced an open-source solution referred to as OpenFPCI or Open-source FOAM to ParaFEM Coupling Interface. This represents a monolithic coupled solution (all software in a single executable) with ParaFEM included as part of OpenFOAM as a library. It was therefore decided a key first task was to take the basic premise of OpenFPCI and further develop and existing framework created by STFC which provides a basis for a partitioned approach where each solver remains separate, and a high-performance code coupling library called the Multiscale Universal Interface (MUI) is used to connect them together. At the time of writing, internal processes within STFC to release the partitioned framework as open source have been completed and it is now available on GitHub as parMupSiF or the Parallel Partitioned Multi-Physics Simulation Framework. The OpenFPCI framework has also been deployed on the ARCHER2 and SCARF HPC facilities with upcoming work being to extract the ParaFEM solver from OpenFPCI and integrate this with parMupSiF. Ultimately this new partitioned framework is well suited to the overall goal of the modular multi-physics WSI toolkit defined by the community.

Work package 2 has an immediate goal in completing and refining an ongoing task to produce a dynamic load-balancing solution for the OpenFOAM based solvers developed by the CCP. In particular a new tool called parmetisDecomp has been created and deployed on ARCHER2 and is currently being improved to include more complex additions such as a compressed sparse graph routine to allow for graph-based decomposition. This tool produces a better load-balanced domain decomposition for problems solved with OpenFOAM than the default tools it is supplied with. The key benefits are improved parallel scalability, better use of resources on supercomputers like ARCHER2 and an ability to better tackle problems with complex mesh structures.

Within work package 3 a key focus is to build-on and improve the existing code and data repositories available on the external CCP website. This is an important underlying activity for CCP-WSI+ and the wider Fluid Structure Interaction community. It offers a moderated and standardised location for depositing Wave Structure Interaction benchmark problems. The primary benefit for the community is a basis for a standardised approach to verification and validation activities, allowing for direct comparison between different approaches and solutions. A specific task within the original proposal was to perform a software audit. This task is well-underway in conjunction with Plymouth University but has quickly revealed a need to find a way to catalogue software used and developed by the CCP in order to build up a WSI taxonomy of methods and approaches. A prototype of a Software Catalogue has been implemented based on a state-of-the-art approach and is currently being tested within the Software Engineering group within STFC. The goal will be for this to be a platform that the CCP can eventually disseminate coupled software solutions and tackle the tractable but difficult problems around code openness and modelling reproducibility that are only exacerbated when coupled solutions involving multiple solvers are considered.

Significant Impact (1 April 2020 – 31 March 2021)

Impact of the CCP during this period is focused on workshops, delivered virtually due to the global pandemic. In particular the kick-off meeting CCP-WSI+ Focus Group Workshop 3 saw a significant attendance from both industry and academia with order of 100 people in attendance. The CCP has also significantly broadened its reach in terms of membership of the community and outreach with related projects compared to its previous funded period. Active collaborators now include EU H2020 funded projects, a number of UK Fluids Network Special Interest Groups as well as continuing impact into important activities like the Supergen Offshore Renewable Energy Hub. The CCP currently has circa 200 active members, been involved in over 100 journal publications since 2015 and already develops and supports a number of key WSI codes. A highly significant impact that the CCP has both within its community and more widely is the definition and curation of highly refined test cases and associated data for WSI problems that provide a structured approach to validation and verification for the WSI community.

Outline Plans 2021/22 (1 April 2021 – 31 March 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

While CoSeC support for this CCP is only at the 6 month stage at the time of writing this report, significant progress has been made to define the key tasks that will dominate the next 12 month's work.

As the CoSeC support provided is broken into three key work packages, so too can the tasks to be undertaken. While there are overlapping aspects to these tasks that will make good use of the collaborative nature of the CoSeC team, the packages are defined such that they are distinct.

Within work package 1 (coupling), the key task to be undertaken is the implementation and testing of a partitioned coupled framework for use on WSI problems. Good underpinning work has already taken place and a starting point in the existing OpenFPCI framework is provided. Once the ParaFEM structural solver has been converted from its use in OpenFPCI to ParMupFiP, rigorous testing will be performed, and the software made generally accessible to the community with the goal of producing collaborative published results.

Work package 2 (high-performance computing) has a defined set of tasks around performing dynamic load-balancing within existing coupled solvers developed by the CCP over the last 5 years as well as generally improving HPC performance. As specific HPC issues relating to the development of the ParMupSiP coupled framework emerge then these will also fall within the remit of this work package.

Work package 3 underpins all of the external dissemination tools used by the CCP (website, software and data repository, workshops etc.) therefore a key element for this package is to continue to support and develop these. The idea of the creation of a taxonomy of tools and techniques for the WSI community also falls within this package and will be explored as the Software Catalogue concept is grown and integrated with existing resources like the data repository.

CCP Turbulence - Turbulence

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.

Summary Report (1 April 2020 – 31 March 2021)

The main activities of the CCP Turbulence of the past year were for the CoSeC staff to perform some improvements for two UK turbulence flagship software (OpenSBLI and Xcompact3D) to improve portability of both codes. The effort on OpenSBLI mainly concerned the optimisation of the OPS library. CMake and CTest have been introduced to both codes as well as steps have been undertaken to set up a proper environment for continuous integration (CI).

Work to improve the serial performance of Xcompact3d has also been performed, these being a requirement for porting to any pre-exascale system where vectorisation will be important to achieve good scalability. Some new features such as the Synthetic Eddy Method for generation of turbulent fluctuations have also been introduced to the code to improve the initialisation procedure of a turbulent field, by reducing the time to development, and therefore the computing time. This can be adapted in the future for the imposition of the boundary conditions.

Work has also started to port the code for acceleration in order to make the framework suitable for heterogeneous architectures typical of the next generation of HPC systems. The choice of OpenACC for the acceleration was dictated by features of the programming model to target a variety of parallel accelerators and therefore the effort would be useful for both GPU porting, typical of many HPC pre-exascale system around the world, as well as the many CPU cores of the new Archer2. This work has mainly been associated with Xcompact3D in order to extend the capabilities of the code from the pure MPI/distributed memory parallelisation strategy currently adopted by the code.

The Gitlab CI/CD service has also been created for the OPS library in order to improve the workflow for the code development. This is a very important milestone for the project since it will make easier the adoption of the OPS framework by codes in the turbulence modelling community. After creating the framework, effort has been devoted into maintaining a virtual machine to run the service and adding more testing items. (e.g., intel compiler support).

A mini-app for solving the 3-D Burgers equation has been created to demonstrate the usage of the OPS library and its tridiagonal solver. Good performance is observed on GPU using the NVIDIA Tesla V100 card.

Changes have also been made to make the OPS Python translator compatible to Python 3 since Python 2 is no longer supported.

Significant Impact (1 April 2020 – 31 March 2021)

The change in the building for XCompact3d from manually editing the Makefile to CMake has improved the portability of the code making the installation more flexible. Some more testing work is necessary as well as some inclusion in the CMake build of more options for the FFT library. The first part of this work is now complete and has been released on GitHub. This work is part of the initiative-2 (I2) in WP2 to make Xcompact3D more accessible and easier to run.

This work has now been released on Github <https://github.com/rfj82982/Incompact3d/tree/dev-stfc>.

A new version of the code with several improvements of the code's serial performance as well as the release of new features like the Synthetic Eddy Method (SEM) for the generation of turbulent fields is now available.

Work under WP1 T1 has also started by porting to OpenACC the tridiagonal solvers of XCompact3D. This work is a first attempt to port the code on GPU. The Gitlab CI/CD service is an important effort to improve the workflow of the OPS code development. The CMake system has made the code easier to compile and more portable, which is of importance for the codes based on the library, e.g., OpenSBLI.

The work lays down the foundation for sustainable software development of WP1 (particularly T3), and potentially other WPs.

The mini-app solving the Burgers' equation will serve as an example for other codes to use the OPS tridiagonal solver..

Outline Plans 2021/22 (1 April 2021 – 31 March 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

Xcompact3d: The main target is to continue the work of GPU porting of the code. The first task (second/third quarter of 2021) will be to finish the porting of the code for a single GPU acceleration. While the current work is mainly based on OpenACC, the plan is to include also OpenMP to compare and contrast. The result of this work will be used as the base for the multi-GPUs acceleration that will be pursued in the second part of the year (Q4 2021/Q1 2022).

OpenSBLI: The main target for next year is to make significant progress in developing a Python translator for creating implicit solvers based on OPS and OpenSBLI and start to implement the multi-diagonal solvers (T1 & T2 WP1). This is the foundation for extending OPS to other CFD codes, e.g., Xcompact3d and help in porting them onto emerging computer platforms.

CCP NTH - Nuclear Thermal Hydraulics

The next-generation nuclear reactors under development to be deployed in the next decades are aimed at achieving inherent safety using technologies such as passive cooling. Such systems require a significantly advanced thermal-hydraulics approach to deal with much higher temperature and pressure systems and/or non-conventional coolants such as liquid metal and molten salts. The traditional methodology is insufficient to deal with the new challenges to be encountered. The aim of CCP-NTH is to support a community of researchers and engineers to develop and maintain computational methods and software packages to modernise the nuclear thermal hydraulics tools to meet the demands imposed through the development of advanced next generation nuclear reactor systems to be employed in the coming decades. CCP-NTH aspires to achieve a number of specific objectives arranged in two work packages: WP1, Community building and networking and WP2, Methodology and code development and maintenance, which includes high fidelity modelling and simulation and robust (reliable, affordable and user-friendly) CFD.

Summary Report (1 April 2020 – 31 March 2021)

The community code, CHAPSim, was used by several research groups and there were both many repeated code sections and particular codes meeting different needs. When CCP-NTH started in April 2020, we collected different version of CHAPSim from those research groups. We filtered repeated functions, made unique function universally useable for different applications, reformatted the code with a standard Fortran 95 format, created a user-friendly input/output interface, added some python scripts for basic data plotting and visualisation and added comment section and variable description sections to generate a user-friendly and readable Doxygen file. To enrich the application of CHAPSim1.0, we added subroutines/functions to enable the code to deal with numerical simulation of thermal-hydraulics of liquid metal (i.e. liquid sodium, liquid lead, liquid bismuth and liquid lead-bismuth eutectic), extending the application of the code from only supercritical fluids to liquid metal for thermal-hydraulics. After all the work done, a unified version of CHAPSim was launched in GitHub for existing users (<https://github.com/WeiWangSTFC/CHAPSim>). This is called CHAPSim1.0. This work made a solid foundation for future code development and was required for effective user support in the long term. Now, most existing users of CHAPSim have moved to this launched CHAPSim1.0 for their research. A CHAPSim users' forum has been created in SLACK to promote communications between users and we also provide informal support for users' questions in this forum. Two users' meetings have been

held to support users' application of CHAPSim to their research. The unified code and regular users' meetings will promote this CFD solver to more users and enhance the communication of users from different research institutes.

High order accuracy for spatial discretisation is desirable for a Direct Numerical Simulation solver in order to capture subtle characteristics from turbulence and heat transfer with limited numerical dissipation. The CHAPSim1.0 is of 2nd order accuracy. Thus, one of the main tasks for the development of CHAPSim is to increase its numerical accuracy up to 6th order accuracy. Another main task for the development of CHAPSim is to increase its parallel capacity. The domain parallelisation in CHAPSim 1.0 is only in one direction, which has severely prevented its application for large problems (e.g. thermal-developing liquid-metal flow with a high Reynold's number) and usage in HPC systems above Tier 2. In order to make complete usage of the latest advanced HPC systems (i.e. ARCHER2) and to improve its scalability for pre-exascale machines, updating CHAPSim to multiple-direction parallelisation and improving its parallel performance is urgent. After a comprehensive literature review on high-order numerical schemes and strategies on multiple dimensional parallelisations and the current available open-source DNS tools, we have chosen the 'compact scheme' for an up to 6th order numerical discretisation, and the '2DECOMP&FFT' as the library for multiple dimensional parallelisations. As these two subtasks are highly coupled, we planned to carry out these two subtasks at the same time. CHAPSim with higher-order accuracy and multiple dimensional parallelisations is called CHAPSim2.0. Until now, we have added the basic operation modules (1st and 2nd derivation and interpolation on both uniform and stretching grids) to build up governing equations. The new FFT lib (2DECOMP&FFT) is under test to make it compatible with CHAPSim2.0. A higher-order multiple dimensional parallelised CHAPSim (i.e. CHAPSim 2.0) will be launched this year. CHAPSim2.0 would provide the community with a powerful tool for thermal-hydraulic research on more areas (for instance, to capture subtle features with higher-order schemes) with much more improved parallelisation capability in HPC (e.g. ARCHER2). The development of CHAPSim2.0 is a big step to meet current and evolving scientific challenges and hardware capabilities.

Due to the COVID-19 pandemic, the planned training and outreach activities have been carried out via Zoom meetings. An annual technical meeting was organized to make CCP-NTH visible to the whole community and it provided an opportunity for the community to communicate their works.

Significant Impact (1 April 2020 – 31 March 2021)

CCP-NTH started in April 2020. One of the most important community-building work was the first annual technical meeting for the CCP-NTH via Zoom. More than 50 researchers from universities, research institutes and industrial companies attended it. We introduced and updated our project objectives, work packages, progress and plans. Fifteen researchers from academia and industry from the UK and abroad (including STFC, Imperial College London, Universities of Leeds, Sheffield, Manchester, Cambridge and Brunel, EDF, Rolls-Royce, Moltex Energy, UKAEA, Penn State University) have presented their most recent work. The attendees contributed to a general discussion and suggestions about this project and expectations from this project. This event enhanced the connection within the NTH community in an exchange of research ideas and discussion of collaborations.

One of the most important coding work carried out during the reporting period was the release of the unified working version of the CFD solver CHAPSim1.0. This work makes a solid foundation for future code development and was required for effective user support in the long term. Another coding work consists of building a uniform and user-friendly input and output data interface for users, which will make new users quickly proficient with the code to only focus on their own research. This builds up parts of the doctoral training for PhD students. The unified code and user-friendly interfaces also promote this CFD solver to more users and enhance the communication of users from different research institutes.

The development of CHAPSim2.0 is still ongoing. Once it is released, it will provide the community a massively parallel high accuracy (from 2nd order accuracy of CHAPSim1.0 to 6th order accuracy of CHAPSIM2.0) DNS thermo-fluids solver..

Outline Plans 2021/22 (1 April 2021 – 31 March 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

The main community building and networking activities included a one-day special topic workshop, a one-day cross-CCP workshop and a one-day annual technical meeting. The main training was the annual CHAPSim users' meeting.

The code-development will focus on the development of CHAPSim2.0. It is designed to have a higher order of accuracy with multiple dimensional parallelisation for the core of the code, and to be able to handle conjugate heat transfer for one of the dedicated physics..

CCP QC – Quantum Computing

CCP-QC aims 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 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. We want to 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. We will interface with the National Quantum Computing Centre (NQCC) in developing early applications of quantum computing and provide training in quantum computing for researchers who are expert in computational science but lack quantum computing knowledge.

CCP-QC will 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, as well as promoting cross-CCP networking to share knowledge on early applications of quantum computing, enabling the widest possible early adoption of quantum enhanced computational science. We will also interface with industry and other partners who are key players in developing quantum computing technology in the UK, and internationally.

Finally CCP-QC will provide information on early quantum computing applications in academic research to the wider community through this 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..

Summary Report (1 April 2020 – 31 March 2021)

The field of quantum chemistry calculations on quantum computers (QC) is less than 15 years old, yet it has experienced explosive growth in the last couple of years with more than 20,000 new publications a year. The recent US announcement of a \$1.2 billion investment over a 5-year period to promote quantum information science and UK's £1 billion investment through NQTP into commercialisation of quantum technology gave an enormous impetus to this development. This quantum gold rush has led to quantum startups mushrooming and developing smart methods for the first wave of noisy intermediate scale quantum (NISQ) computers.

A working group was setup to coordinate the work on Quantum Computing (QC) and Quantum Chemistry. Members came from the Materials Chemistry Consortium (MCC) community and Quantum

Computing community. Vlad Sokhan from STFC was the person delivering the work and has spent three months doing a thorough review of various methods used in QC to solve the Schrodinger equation for small molecules, within Hartree Fock confines. A report and a presentation was prepared for the MCC community and QC executive committee. A new working group was setup to direct the work for QC and Crystallography. In addition, Alin Elena, has setup mailing lists and a website for the newly formed CCP-QC and coordinated three executive meetings where directions are setup plus the coordination of the two working groups. Two meetings with the National Quantum Computing Centre took place on possible cooperations at CoSeC and CCP-QC level.

Working group 2, QC and Crystallography, has regular meetings up to now refining the problem that can be realistically targeted with a quantum computer at this time, in this case the D-Wave system. A candidate to conduct the work is identified, Adam Callison from Imperial College, who will be tasked with the implementation on QC asap.

A minicolloquium "Integrating Quantum Computers in Condensed Matter Physics Simulations." organised by Francois Jamet (NPL), Liviu Chioncel (University of Ausburg), Martin Kiffner (University of Oxford) and Alin Elena (Daresbury) proposal was submitted to European Physics Society Condensed Matter biannual conference and was accepted. Due to covid the workshop will happen in August 2022. The workshop includes invited speakers from Google and IBM. We plan to run a small version at NPL this autumn, restrictions allowing.

Significant Impact (1 April 2020 – 31 March 2021)

Activities were severely restricted by COVID-19 but we have acceptance of the workshop to the EPS.

Outline Plans 2021/22 (1 April 2021 – 31 March 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

To investigate possible applications and quantum advantage to optimisation problems in Crystallography.

MCC – Materials Chemistry Consortium

The Materials Chemistry Consortium exploits high end computing in a broad programme of work modelling and predicting the structures, properties and reactivities of materials. The consortium is a broadly based but coherent grouping comprising 36 university groups, with the emphasis on modelling at the atomic and molecular level but with growing links to models at larger length and time scales. Founded in 1994, the current scientific programme is built around seven related themes: catalysis, energy storage and generation, surface and interfacial phenomena, nano- and defect structures, soft matter, biomaterials, environmental materials. The Consortium has an active programme of code support, development and optimisation, tapping into the ecosystem of UK based software development initiatives including CoSeC.

CoSeC supports the consortium across the range of techniques used by its members, embracing both force-field methods employing static and dynamical simulation methodologies and electronic structure methods with a strong emphasis in recent years on Density Functional Theory (DFT) techniques employing both periodic boundary conditions and embedded cluster implementations. The four main codes supported by CoSeC are ChemShell, CRYSTAL, DL_FIELD and DL_POLY.

Summary Report (1 April 2020 – 31 March 2021)

ExCALIBUR: CoSeC support from Ian Bush, Alin Elena and Tom Keal has been an integral part of the formation of the Materials and Molecular Modelling Exascale Design and Development Working Group under the EPSRC ExCALIBUR programme. The CoSeC team are on the steering group for the project and have leadership roles on the three exascale challenges identified (large scale calculations, complex workflows and I/O). The team gave presentations on these themes at the project kick-off workshop held online in May, and are investigating relevant libraries and frameworks to address them with the aim of producing blueprints for future work under each theme. Meetings have also been held with NVidia and Intel to discuss shared objectives. The team are now identifying major application areas in discussion with the community to target the next phase of the ExCALIBUR programme, and participating in a cross-cutting bid on the theme of coupling technologies from atomistic modelling to continuum.

ChemShell/DL-FIND: In this period the majority of the effort has targeted the release of a new beta Py-ChemShell version with improved materials modelling features, which we now regard as suitable for production materials calculations and are encouraging the MCC ChemShell user community to switch. Significant efforts have been made this year to port the Py-ChemShell code to the new ARCHER2 platform (including switching standard compilers from Intel to GNU) and the Tier 2 YOUNG machine, along with the legacy Tcl-ChemShell program and associated GAMESS-UK QM program. 0.1 FTE of ChemShell effort was committed to supervision of the ChemShell work package in the MCC flagship software development project "SAINT", which has now successfully concluded with the implementation of a periodic QM/MM scheme to enable the modelling of metallic systems, which cannot be accurately described using existing finite cluster models.

DL_POLY/DL_FIELD: The effort spent in this period addressed a number of issues related to critical stability and performance of DL_POLY in its path to a new refactored release as version 4.10 and a final wrap-up version of 4.09.5. The new version incorporates new defects-detecting functionality, developed in collaboration with Prof. Trachenko at QMUL. The most notable HPC methodology improvement is the single sided halo Verlet Neighbour List strategy leading to over 15% speedup for critical regime runs and diminishing to 5% for very large systems on very large processor counts. Numerous memory optimisations and fixes have also taken place. In DL_FIELD auto-detection of atom types clay minerals has now been accomplished. This will enable users to set-up bio-inorganic interface models using simple xyz formats. Further effort, as an outcome of SAINT and BioMat ambitions to integrate with ChemShell, has gone in improving stability and memory optimisation of DL_FIELD.

SAINT: Calculations on Group II oxide surfaces with probe molecules (CO, H₂O and NH₃) were running successfully at the end of the project. These provide initial data for the service and for a possible publication on the project's capabilities.

CRYSTAL: CRYSTAL17 v1.0.2 was tested and installed on the ARCHER2 4 cabinet system for use by consortium members.

Significant Impact (1 April 2020 – 31 March 2021)

A particular highlight of this period has been the setting up of the Materials and Molecular Modelling Exascale Design and Development Working Group under the EPSRC ExCALIBUR programme. CoSeC have been heavily involved in the scoping, setup and steering of the working group with the aim of preparing the MCC community for future exascale systems. The steering group have carried out a survey of the community and ongoing discussions to identify community needs with the aim of maximising the impact of the project and preparing for later phases, and sought further ExCALIBUR funding for knowledge integration activities and cross-cutting coupling technologies.

Outline Plans 2021/22 (1 April 2021 – 31 March 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

ExCALIBUR: the phase 1 project will conclude by end of July, with the aim of completing the python API to Crystal, demonstrating complex workflows involving Py-ChemShell, and finishing the benchmarking of high performance I/O approaches, together with roadmaps for follow-on work subject to phase 2 funding.

MCC benchmarking: Ian Bush will develop a set of Benchmarks over the year to compare different codes used by the MCC community. This will not simply be to take standard benchmarks and run them on machines of interest. Rather with help from the community scientific cases of interest will be identified, and then it will be attempted to solve those problems with different codes on Archer2 and Tier 2 systems, and the performance compared. The important point is that each code will be required to solve the problem to the same accuracy. Where appropriate code developments will also be supported to enable best performance on the architectures of interest.

ChemShell development will continue to focus on the new Py-ChemShell code. An excited state interface in Py-ChemShell will be developed following the original implantation in Tcl-ChemShell, with support for multiple electronic state optimisation. Py-ChemShell will be maintained on ARCHER2 and Tier 2 systems such as UCL's Young facility and a new release of Py-ChemShell prepared including the completed developments from the SAINT project. The legacy Tcl-ChemShell and GAMESS-UK packages will also continue to be maintained for the benefit of the community. The current set of materials modelling tutorials in Py-ChemShell will be reviewed and revised.

DL_POLY/DL_FIELD: A new DL_POLY_5 was released in February 2021, including features and fixes from Year 20/21. A new frozen-frozen Ewald interactions routine will be implemented. Scalability performance of short-range interactions investigated and work-plan to make domain decomposition independent on cutoff designed, seeking further funding through the eCSeE route. Compilation and running on ARCHER2 are a priority. A new DL_FIELD version 4.8 will be released, including new features implemented in Year 20/21. The software stability will be greatly improved and enable multiple software calls within ChemShell for carrying out QM/MM optimisation. The software development will continue to focus on workflow and improve ease of use, especially in setting up complex system model. In particular, the consolidation of seamless processes to solvate a wide range of system models including inorganic-organic mixed system components.

UKCP – UK Car-Parrinello Consortium

The United Kingdom Car-Parrinello Consortium (UKCP) is a group of researchers across the UK who develop 'first principles' quantum mechanical techniques for studying atomistic systems and apply them to a very wide variety of systems. The UKCP consortium is one of the longest-running High-End Computing Consortia in the UK, and has been funded almost continuously by EPSRC since the 1990s. The 1 FTE of core support is currently focused on supporting the CASTEP code, one of the UK flagship first principles codes.

Summary Report (1 April 2020 – 31 March 2021)

This reporting period has been most heavily affected by COVID-19 related disruption. In particular the availability of Dominik Jochym was reduced to two fifths of usual effort, except during September/October 2020 and March 2021. The priority during this time was to maintain the working environment for CASTEP development on the Bitbucket platform and to ensure the continued availability of academic CASTEP licenses following the "soft launch" of the new free-of-charge worldwide license in August 2019. This work was critical for the continued development of CASTEP functionality by the UKCP community. The transition from Mercurial to Git as the version control tool used for CASTEP has highlighted that a change to developer workflow is needed. An updated workflow document is to be drafted and circulated to the CASTEP developer community for comment before

adoption. The influx of new international CASTEP users has brought requests for video-based tutorial material in how to get the software installed, particularly on Microsoft Windows. Dominik is investigating options for developing such material using the OBS (Open Broadcasting Software) package.

Dominik was heavily involved with the launch of STFC's electronic licensing portal, with CASTEP as the first available "product". This service uses UCL-B's E-lucid platform and is available at <https://licenses.stfc.ac.uk/>. There were challenges addressed during testing, particularly to ease the user experience while maintaining the requirements of the academic CASTEP licence. On 11th March the service was used to launch the release of academic CASTEP version 20. Dominik will be providing guidance for other CoSeC funded projects to make use of the electronic licensing service.

As of March 2021, effort to the project is back to planned levels. The major code development planned for 2020/21, van der Waals DFT functionality in CASTEP, has been moved to the 2021/22 period. The YAMBO interface for CASTEP and ORCID authenticated downloads on the CASTEP website have started development, with delivery delayed to 2021/22. The commercial release of CASTEP 21 (outside of CoSeC) was delayed to Q2 2021 so the academic release has been delayed accordingly.

Significant Impact (1 April 2020 – 31 March 2021)

In this period an additional 350 groups worldwide have been granted access to CASTEP free-of-charge for academic use. Prior to August 2019, the only option for researchers outside of the UK was to purchase a commercial license for CASTEP through the Materials Studio product. The increased availability of the software is expected to lead to a substantial increase in CASTEP usage and citations, with applications spanning materials design, pharmaceuticals and fundamental science.

The use of STFC's electronic licensing platform for CASTEP allowed for over 100 groups to obtain access to CASTEP v20 during March 2021.

Outline Plans 2020/21 (1 April 2020 – 31 March 2021)

(Detailed plans and milestones are included in the Technical Appendix for reference)

In direct support of the UKCP and wider CASTEP community, the annual CASTEP software release management, teaching workshop and code developer's workshop will be carried out. Maintenance of the software development and testing platform will allow for the continued growth of the CASTEP package and the software standards required and expected by the commercial partners and academic community. It is expected that at least some of the workshops will be held virtually. Workshop materials will be re-evaluated and adapted as needed for online presentation.

The electronic license management of CASTEP will be in its first extended period of production. While operating at scale, the workflow will be monitored and refined as needed. Lessons learned will be communicated to the other CoSeC project leads through internal presentation and discussion. Guidance will be offered for projects wishing to use the service.

The following are previous objectives delayed because of COVID-19 disruption.

- The implementation of "van der Waals DFT" will be prototyped within CASTEP. This major new functionality promises to improve the underlying approximations used in CASTEP, with particular impact in the field of pharmaceuticals and organic electronics.
- An interface between CASTEP and YAMBO will be produced. This will give CASTEP users access to advanced spectroscopic simulations offered by YAMBO.
- To complement the electronic licensing of CASTEP, a download section on the CASTEP website will be made available using ORCID to authenticate. At this time only the lead licence holder for a

group can download CASTEP through the STFC licensing portal and individual users are given access on an ad hoc basis.

UKCOMES - UK Consortium on Mesoscale Engineering Sciences

The United Kingdom Consortium On Mesoscale Engineering Sciences (UKCOMES) – founded in 2013 – is a group of researchers across the UK who develop and apply mesoscopic modelling techniques to explore systems of scientific and industrial interest at scales between atomistic and continuum-based levels. Several modelling techniques are applied in this consortium, but the most frequently used and studied are Dissipative Particle Dynamics (DPD), a particle-based method similar to atomistic molecular dynamics but capable of modelling larger systems with correct fluid behaviour, and the Lattice Boltzmann Equation (LBE) method, a particle-based statistical technique capable of modelling fluid flows with complex geometries and interactions between multiple fluids and phases.

The 0.6 FTE of core support per year is focussed on developing DL_MESO and MPLB, the consortium's community codes for DPD and LBE simulations, by adding new functionality, optimising for various computing architectures and improving interoperability with other codes and libraries. Both activities allow for a wider range of systems to be modelled with available computing resources, including the UK's national supercomputer ARCHER 2. Development, porting and optimisation support is also provided for other LBE codes used by the community, particularly particularly MPLB and HemeLB, which is based on sparse system geometries.

Summary Report (1 April 2020 – 31 March 2021)

A constraint solver (RATTLE) has been added to DL_MESO's DPD code to enable it to model bonds with fixed lengths. Since oscillations of flexible bonds between particles in molecules are typically the reason why larger timesteps tend not to be used in DPD simulations, using fixed-length constraints instead should enable mesoscale modelling of molecular systems over considerably longer times. The Langevin barostat in the same code has also been changed from using an iterative approach to a more sophisticated piston force integration scheme, which will guarantee numerical stability for constant pressure simulations when larger changes in pressure and surface tension across e.g. membranes are required.

Changes are being made in the MPLB code to exchange information between different mesh blocks. Multi-block meshes can greatly help to reduce memory consumption for many geometrically complex flow problems, such as flows that occur in oil and natural gas pipelines and in city piping networks. To prepare for a release by Q1 2022, the code performance has been investigated based on previous milestones. Based on this investigation, the code has been changed and significantly improved performance is now observed on GPU-based computers compared to CPU computers. The facilities for multi-block meshes will be further iterated based on the new improvements in code performance. Large Eddy Simulation modelling based on LBE has been investigated, and it is found that the adaptive mesh refinement capability is important for achieving efficient simulations. It might therefore be good to look into the backend code of the HiLeMMS project based on the AMReX library in the future.

Significant Impact (1 April 2020 – 31 March 2021)

The inclusion of fixed-length constraints for DPD simulations in DL_MESO will greatly expand both the range of materials that can be modelled at the mesoscale and the timescales over which those simulations can be applied. The development on multi-block meshes for MPLB can greatly help to reduce memory consumption for many applications, (e.g., T-junctions). The performance improvements

on GPUs will accelerate modelling research for academic members of the community and enable reasonable 3D simulations with affordable resources (e.g., using a high-end workstation).

Outline Plans 2021/22 (1 April 2021 – 31 March 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

Various models for suspended particles in LBE simulations will be added to allow simulations of solid particles moving through fluids and potential coupling of LBE with particle-based simulations (e.g. molecular dynamics, DPD). A new version of DL_MESO (numbered 2.8) will be released with new functionalities added since the previous release to provide new DPD and LBE modelling capabilities, as well as improved documentation and user tutorials.

The main plan of the MPLB code is to improve the stream-collision scheme and optimise the memory usage, as well as improving performance based on a recent investigation. A new version of the code will also be released with various implemented new capabilities, enhanced documents, and tutorials..

HEC Plasma Physics

HECPlasma and CCPPlasma resources have been combined into a single workplan. Please see the report above under CCPPlasma.

HECBioSim - High-End Computing for Biomolecular Simulation

HECBioSim exists to bring High-End Computing for biomolecular simulation to a wider community, including users from industry and experimental bioscientists, and to engage physical and computer scientists in biological applications. The Consortium works closely with CCPBioSim.

HECBioSim is supported by 1.0 FTE of SLA core effort, which provides support for scientists applying for time on ARCHER, primarily through maintenance of the HECBioSim web portal. It includes help on preparation of applications e.g. with the HECtime resource calculator, and on reporting the outcomes of approved projects. The SLA post also works on a variety of codes for biomolecular simulation and analysis appropriate to High End Computing. For further information please see www.hecbiosim.ac.uk.

Summary Report (1 April 2020 – 31 March 2021)

Due to the emergence of the SARS-CoV-2 virus and its affliction of a pandemic on the human population (Covid-19) a decision was made by the HECBioSim management group to temporarily suspend the work plans planned for this reporting period to focus on establishing and supporting a global Covid-19 initiative.

Work in this area initially focused on establishing lines of communication and knowledge exchange, so a new area of the HECBioSim website was created and populated with information available at the time <https://www.hecbiosim.ac.uk/covid-19>. We focused on populating this new web resource with up to date information obtained from computation/experimentally, reliable statistics and news as it emerged and the portal quickly became a go-to resource for those doing simulations to find out the latest or links to structures, stats and information.

We then had all the ingredients to start building deeper collaboration, we put out a call for researchers to send information on their active research <https://www.hecbiosim.ac.uk/covid-19-projects> or offers of expertise/resources or hands available within their groups <https://www.hecbiosim.ac.uk/covid-19-resources>. This really allowed our chair (Syma Khalid) to start matchmaking groups offering, with groups look for help or actively working on research to form deeper collaborations. We have one such success story written up as a highlight here <https://www.hecbiosim.ac.uk/covid-19-alliance>.

Another facet of the Covid-19 initiative was coordinating access to HPC resource that was rapidly made available by many HPC facilities around the UK for use for Covid-19 specific research. With this came a high volume of requests for support from the community with using this newly accessible resource, we ran this on an ad-hoc basis to cope with the volume and breadth of requests along with the rapidly evolving landscape. Work in this area included software compilation of new versions of existing codes and codes that were not currently available due to being commercial (made temporarily free) or that were missing from some HPC machines. A large volume of debugging and benchmarking work supporting groups with problems running their CV-19 simulations and getting the correct performance from the hardware they had access to. Some groups had not used some architectures or even HPC outside of their local lab, some assistance was provided in the form of example submit scripts and specific advice for their simulation systems.

Some urgent work has been performed on the HECBioSim website due to a number of lengthy down time events caused by installation and maintenance activities in the data centre. This work included migration to a microservices hosting model where the site has been ported to a docker container and is now hosted out of the STFC Scientific Computing Department cloud. To do this required updating the whole underlying software subsystem since this site was first put together 7 years ago, the result is this is now on the latest generation of all software packages with support lifetimes running into the mid 2020s. The hardware is much more suited to running websites, and we now have a new security model with auto renewing SSL security certificates. In the long term, this should require much less maintenance since containers can be dynamically pulled with latest software versions.

There have been major advancements in how the biosim consortia are going to deliver training courses throughout the pandemic. Due to the nature of the pandemic preventing face to face teaching, we are now targeting online delivery of training and have seen a large uptick in the volume of researchers completing our training available online. We are now running two separate training clusters, one smaller cluster with all of the biosim consortia training catalogue available 24/7 for up to 20 concurrent users. A separate cluster is brought online for instructor lead workshops, this cluster has only the selection of courses being delivered and is now fully scalable to allow upto 500 concurrent users to take part. This required work on provisioning a new cluster on the STFC cloud for the 24/7 cluster and a refresh of all containers to the latest versions of software. Work on the live workshop infrastructure required a full redesign of the architecture to cope with the huge loads such scaling would place on the facility. This required benchmark studies to identify the pinchpoints and then work to solve the load balancing issues found. We now have the blueprints to run extremely large online training sessions by combining this technology with zoom.

Work on benchmarking systems around the UK has progressed well. The data has been gathered for ARCHER, ARCHER2, Bede, JADE, JADE2, ISAMBARD, THOMAS and a number of DIRAC machines. We now have performance data, optimised and run scripts for all these machines. There is a new section of the HECBioSim website dedicated to providing access to these HPC machines and this includes the benchmarks mentioned above in addition to example scripts on how to best run the codes for each machine, advice for PIs on how to increase their chances of a successful application for resources and links to documentation and information on how to get support for each machine. These can all be found under the "access HPC" menu tab on the hecbiosim website, these services should serve to enhance the ability of users to get more performance from all of these architectures and will enable PIs to make successful applications to each of the different HPC resources backed by our benchmarks.

A significant amount of work has gone into supporting the consortium in preparing for the emergence of new Tier2 resources. This has entailed software compilation on JADE2 and Bede and preparing software modules. We now fully maintain and support all HECBioSim remit software on JADE, JADE2 and Bede. Significant support was given to users in the HECBioSim early testing programme on these

machines in getting their simulations up and running on these new architectures, some of which requiring very specialised debugging.

Work on Longbow that was initially planned for earlier in the reporting year was significantly delayed by the urgent work on the Covid initiative. It has been decided that the changes that have been implemented will be rolled into the upcoming 2021.0 release of Longbow due to both releases being major releases, so it is better to make such a release once. The new features include better support for generic software, better support for GPU based machines, more control over openMP and MPI, a deeper chemshell integration and full support for Tier2s out of the box. There will be a new security model compatible with the enhanced security now required by ARCHER and other machines, and Longbow will switch to more modern configuration interface. All of these changes will make the use of all of the different HPC machines available to our consortium members very simple.

Significant Impact (1 April 2020 – 31 March 2021)

There have been two projects undertaken that have led to significant impact in this reporting period.

The first has been the extremely successful covid-19 <https://www.hecbiosim.ac.uk/covid-19> initiative that our community launched at the time that the SARS-CoV-2 novel virus emerged and caused the Covid-19 pandemic. This initiative really focused on the research that was being done in the biomolecular simulation arena. CoSeC (via James Gebbie-Rayet) was able to provide much needed expertise and effort to rapidly bring this resource into a position where it became an internationally recognised initiative. This work entailed crafting a new web portal, keeping very up to date with research being done (not waiting for published work), new structures emerging from experiments, new offers of compute from computing centres and offers of hands and expertise from labs as well as a multitude of other interesting links to resources as they emerged. This allowed the wider consortium and community to organise itself and start to form working groups to tackle specific problems and avoid duplication. A success story of this in action can be found here <https://www.hecbiosim.ac.uk/covid-19-alliance>.

The second piece of work is on the infrastructure for training workshops. It is one thing to deliver a training workshop to 30-40 students but quite another to deliver for 180+. CoSeC (via James Gebbie-Rayet) has enabled this new capability to deliver instructor led workshops at an unprecedented scale. In the old days, we would spend the first hour fiddling around with individual computer problems when running workshops where users bring their own laptops. With this new capability, only a web browser is required so this issue is gone. This fits the new conditions imposed upon us by the pandemic, forcing everybody to dial in from home and has enabled us to innovate our way to delivering training at a larger scale than ever before. We do all this by constructing all of the complex infrastructure and software required in the cloud, this means everything is built and tested by experts and eliminates issues with individual computers. This has recently been tested with 187 delegates at the CCPBioSim training week, and we are already planning on more large training events like this later in the year!

Outline Plans 2021/22 (1 April 2021 – 31 March 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

The plans for the next working year are a work in progress, we are in the process of identifying opportunities for new projects to add to our planning schedule for next year. It is certain that ARCHER2, Bede and JADE2 will be newly in production and there will inevitably be a high volume of support in this initial phase. We are planning some training events surrounding the use of UK HPC facilities, particularly surrounding what is available, what the differences are and how to best to make use of the facilities.

UK-AMOR - Atomic, Molecular and Optical Physics R-matrix Codes

The UK-AMOR consortium exists to develop and exploit the UK Atomic, Molecular and Optical Physics R-matrix codes on ARCHER. The R-Matrix suite is a set of programs for electron (positron) -atom and -molecule scattering, (ultrafast) laser pulse interactions and related problems. UK-AMOR's scientific goals are to: study attosecond physics in atomic and molecular systems; compute fusion-relevant atomic data; study electron-molecule collisions with relevance to biological radiation damage, fusion energy and technology. The consortium allows high-end use of the CCPQ codes RMATRIX I, PRMAT (RMATRXII, PFARM), UKRMol(+) and RMAT_REACT, and RMT. Major code development projects of UK-AMOR include the development of a novel R-matrix treatment of ultracold chemistry, RMAT_REACT, which will be used to study key systems (atom-atom and molecule-molecule) in parallel with experiments being performed in the UK and abroad, and algorithmic improvements made to existing codes to improve functionality and scalability.

The consortium is the focus for UK and international activity for users of the UK-AMOR codes suite. It is committed to: making its codes widely available; interacting closely with CCPQ; engaging with industrial partners and international collaborators through workshops and other activities; broadening the user community for the AMOR methodology and codes; providing training for the next generation of computational scientists in the UK.

UK-AMOR has received some administrative CoSeC support from CCPQ funding (Martin Plummer) and an average of 0.2FTE per year direct CoSEC support for code developments..

Summary Report (1 April 2020 – 31 March 2021)

UK-AMOR has been linked to CCPQ, which sadly was not renewed in EPSRC's CCP Call. This led to the removal of CCPQ CoSeC support apart from a small residual carried-over amount for administration, searching for new funding and assistance with necessary winding up of CCPQ scientific CoSeC support activities. The most practical way of using this residual support was to combine it with the UK-AMOR support. The main coding highlight for this reporting period was the completion of Andrew Sunderland's memory optimization work for the UK-AMOR code UKRMol+ serial and OpenMP-threaded GBTOlib library as requested by the developers. This highly optimized code is now applicable to a wide/full range of electron-molecule applications and is part of the UKRMol+ GIT repository (final debugging and testing of the optimization extended to the new MPI parallel version of the library was delayed due to COVID-related disruption). Andrew has also been obtained 3 person months work from PRACE to work on PFARM (part of the Prace 6iP Applications benchmark suite: <https://repository.prace-ri.eu/git/UEABS/ueabs>) on several international Tier 0 platforms between now and October 2021.

Another highlight is the award of an ARCHER2 eCSE project to work on UK-AMOR's RMT code for laser atom/molecule interactions (MP will be the main worker, the PI is Dr Andrew Brown, QUB and UK-AMOR). This work, commencing in October, will improve the parallel efficiency of RMT for the large-scale jobs enabled by the introduction of arbitrarily polarisation of the laser (and future planned new functionality).

Martin Plummer (as planned co-I) supported Dr MM Law (Aberdeen) in an ambitious New Horizons proposal which would extend the RMAT_REACT framework and code-base to treat antimatter rearrangement collisions (antimatter chemistry) relevant to low-energy experiments planned by the ALPHA project and others at CERN when work recommences after the May 2021 restart. While this was not successful it provides a basis for a responsive mode application for a new 3-year project. This new application will now be underway in Q2 2021.

While the second half of the reporting year has mainly been consolidating development work relating to the eCSE project and the AquA-DIP grant, a major highlight has been the revival and reorganisation of the joint UK-AMOR/CCPQ and HEC/CCP-PLASMA workshop, 'Atomic and Molecular Data Needs for Plasma Applications', postponed from April 2020 (with Jonathan Tennyson joining the organisers). This three day online workshop, to take place on 13-15 April 2021 with ~160 registered participants, will have 23 talks, 3 informal discussion sessions, a demonstration by Quantemol Ltd and a poster session (a lightning introduction session and a more detailed as-live breakout rooms for each poster). The posters and talks will be preserved on the CCPQ website.

Finally, Dr Andrew Brown (QUB) has been awarded a Research Software Engineer fellowship to unify diverse AMO codes (concentrating on the various 'R-matrix' packages) into a multipurpose effective single resource across the range of AMO physics. His proposal was supported by UK-AMOR generally and by CoSeC support (practical commitment from MP) in particular and will be a major continuation and expansion of the CoSeC/SLA code support and curation work given over the years for CCPQ and UK-AMOR.

Significant Impact (1 April 2020 – 31 March 2021)

Impact of CoSeC support work in this period will take longer than the reporting year to appear.

The main immediate impact is the new eCSE funding for work on UK-AMOR's RMT code. [The earlier CCPQ/UK-AMOR work on adapting PFARM for RMT_REACT will have impact in the form of new collision calculations in the near future, while longer term impact is expected from the work on the bound-state code.]. The earlier work on PFARM (within CoSeC and as part of PRACE) has resulted in new PRACE funding for AGS and the PFARM code being part of the Prace 6iP Applications benchmark suite.

CoSeC support enabled Martin Plummer to be involved with a successful 3-year EPSRC grant application (the 'AQuADIP' project) which commenced in May this year, with PDRAs based at QUB and UCL, linking RMT with a semi-classical approach. A result of this grant in this reporting period was the international workshop 'Quantum Battles in Attoscience, co-funded by CECAM, which switched from being an in-person event to an online event (organised by Prof Carla Faria, UCL and colleagues with some help from QUB) and was extremely successful. It had much wider international participation (>300 attendees) than the in-person event would have had, as well as a highly successful disciplined approach to discussing opposing viewpoints and approaches to aspects of attosecond science (with follow-up articles on this approach to organisation in Physics World). The workshop is also an excellent template for future workshops in a post-COVID environment.

In particular, the 'Atomic and Molecular Data Needs for Plasma Applications' workshop, with ~160 participants, should have strong impact in strengthening and forming links and collaborations between the AMO and Plasma physics communities.

Andrew Brown's well-deserved RSE fellowship will involve work that follows on from general CoSeC/SLA support over past years and the present to develop and unify the AMO (R-matrix in particular) code packages for wider use. AB told MP that the fellowship was "... a recognition of the value of AMO software in general, the R-matrix codes specifically, and their potential impact over the long term. You have been such a key part in maintaining the community and the codebase. I am excited to work with you over the next five years to create something that will deliver benefit for the R-matrix community and AMO physics more generally."

Thanks to Andrew Sunderland's CCPQ support and follow-up PRACE work, the PFARM benchmark suite, available for public download from the Unified European Accelerator Benchmark Suite repository, has been restructured with the CPU and GPU versions streamlined (now both versions are included in the one source code download and configuration). Scripts to automate benchmarking of the atomic and molecular datasets have been published and the accompanying documentation has been revised and expanded. Testing of the benchmark by third-parties is currently underway, shortly leading to a new release of the UEABS.

Outline Plans 2021/22 (1 April 2021 – 31 March 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

In 2021-2022, UK-AMOR CoSeC support is limited to 0.2FTE for Martin Plummer, without any additional CCPQ funding for administrative work. Detailed milestones will be worked out following UK-AMOR management meetings in Q2 2021-2022, however Martin will carry on with RMAT_REACT coding work (the bound state code and/or follow-up work following the departure of UK-AMOR PDRA E Spinlove) as far as this is practical, or support other R-matrix work with which Martin is involved, and with general support for consortium meetings, the renewal proposal and pro-active work to find other UK-AMOR related new funding (Martin is separately involved with an eCSE and an ongoing physics grant with the UK-AMOR code RMT for laser atom/molecule interactions). Andrew Brown's RSE Fellowship is a major boost for continued general curation and user-friendly expansion of the UK AMO codebase.

UKTC - The UK Turbulence Consortium

Understanding, predicting and controlling turbulent flows is of central importance and a limiting factor to a vast range of industries: naval, aeronautical, automotive, power generation, process, pharmaceutical, meteorological and environmental. For example, energy-efficient fluid mixing is a major limiting factor in the chemical industry which is the second largest industrial consumer of energy worldwide, but also in the pharmaceutical and process industries. How can one mix best with as little power loss as possible? The cost of pumping oil and gas through pipelines is directly proportional to the frictional losses due to turbulence. Polymer additives can modify the turbulence in these pipelines to reduce frictional losses and speed up the flow by one or two orders of magnitude with the same power input. How can one maximise this effect which remains poorly understood? Moving vehicles often generate noise, and in the case of airplanes, it can have a severe detrimental impact close to airports. What kind of airbrake could generate enough drag with as little aeroacoustic losses as possible? Worldwide ocean shipping consumes about 2.1 billion barrels of oil per year while the airline industry uses 1.5 billion barrels per year. How can turbulent boundary layers be controlled to significantly reduce friction drag which will simultaneously reduce costs and emissions? Simulating and understanding turbulent flows to answer the previous questions is one of the most challenging problems in science. Many of the environmental and energy-related issues we face today cannot possibly be tackled without a better understanding of turbulence. The overarching objective of the UK Turbulence Consortium (UKTC) is to facilitate world-class turbulence research using national High-End Computing (HEC) resources. This involves performing numerical experiments with turbulence-resolving computational approaches. Such simulations are ab initio calculations based on first principles and produce data to answer basic questions regarding the physics and modelling of turbulent flows found across a range of engineering, physiological and geophysical applications. The consortium serves as a forum to communicate research and HEC expertise within the UK turbulence community, and to help UK science remain internationally leading in this aspect of HEC-based research.

Summary Report (1 April 2020 – 30 September 2020)

The planned work is to implement velocity-slip and temperature-jump in OpenSBLI. It is a Python-based open-source DNS code to solve the compressible Navier-Stokes (NS) equations with application to

shock-boundary layer interactions (SBLI). It automatically generating C code that performs the finite difference approximation to obtain a solution of the NS equations. This C code is then targeted with the OPS library towards specific hardware backends, such as MPI/OpenMP for execution on CPUs, and CUDA/OpenCL for execution on GPUs. The code relies heavily on external libraries and only works on Python V2.7 only. While learning Python, I have gathered the correct version of all external libraries and installed them in a local machine. OpenSBLI can generate the C code in the local machine for HPC.

Since April 2020, due to Covid-19 special leave, no effort has been booked against UKTC. No further progress has made in this period.

Significant Impact (1 April 2020 – 30 September 2020)

Due to the drop in staff effort caused by Covid-19 there is no significant impact to report for this period.

Outline Plans 2021/22 (1 April 2021 – 31 March 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

Due to the drop in staff effort caused by Covid-19 the 2020/21 plans have been deferred to 2021/22.

OpenSBLI currently uses traditional no-slip boundary conditions whereas it is known that high-speed flows can involve both slip and temperature jump. Using kinetic theory through the moment method, we will incorporate high-order components to test and validate impact on results.

As we perform simulations on large-scale facilities, such as ARCHER2, the need for in situ visualisation becomes a key consideration. Previous work considered Code_Saturne as an exemplar to understand how we can develop an in situ capability in other UKTC codes. We will now explore what is done in nek5000 to get a better understanding of best practice in delivering this capability.

UKCTRF - UK Consortium on Turbulent Reacting Flows

The UKCTRF will offer a step change in the fundamental understanding of turbulent reacting flows using advanced numerical simulations utilising developments in HPC, and translate physical insights into high-fidelity models for engineering simulations through collaborative and complementary research.

Summary Report (1 April 2020 – 31 March 2021)

Work on this project began in January 2019 and focussed on work around the HAMISH code and adaptive mesh refinement (AMR). In the reporting period, Jian Fang has worked with partners in Cambridge and Newcastle and contributed to validating the HAMISH code in a series of 1-D and 2-D test cases. In addition, using the TGV benchmark case up to 10243 mesh, the HAMISH code has demonstrated fourth-order spatial accuracy and fifth-order of dissipation confirming it is well suited for DNS/LES of turbulent reacting flows. Jian Fang also worked on the development of an AMR library based on modern Fortran, in which dynamic memory allocation, parallel IO, and binary operations are enabled. The library has been benchmarked with the TGV test case and was shown to have fourth-order accuracy.

Significant Impact (1 April 2020 – 31 March 2021)

The successful benchmarking of HAMISH has confirmed its potential for LES/DNS applications relevant to reacting flow simulations. Additional work relating to boundary conditions is on-going at Cambridge.

A new AMR code (PAMR) abstracted from HAMISH has been developed using FORTRAN language, which supports dynamic memory allocation, parallel HDF IO, and binary operations. It can be used as a standalone library and can be applied to other PDE solvers. The cell-based parallel AMR library could be useful to the broader CFD, MHD and other scientific communities.

Outline Plans 2021/22 (1 April 2021 – 31 March 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

In 2021/2022, the plan is to assess SENGAs using parallel I/O and to benchmark HAMISH on problems that test the NSCBC boundary conditions. Further testing and validation studies will then be undertaken for 3D reacting flows.

Software Outlook

Software Outlook focuses on software technologies that are vitally important to the development and optimisation of the world-leading scientific software produced by the CCPs. This includes evaluation of new programming techniques that are essential for the timely and cost-effective exploitation of current and near-future High Performance Computing systems, demonstrating how specific software technologies can be applied to existing applications, and providing guidance on the principles of best practice within software engineering.

Summary Report (1 April 2020 – 31 March 2021)

Software Outlook has enjoyed a lot of activity during the past year. Code containerisation strategies are currently of high interest to many of the CCPs but, as has been experienced by several, it is not always straightforward to use them effectively or to know which technology to use. Software Outlook has written a report that provides best practice guidance, a review of both the commonly known technologies and the lesser known ones, a review of container management systems and guidance on building and running Docker and Singularity containers. We have presented this work to some of the CCPs and are available for further presentations. This guidance will allow CCPs/HECs to use code containers effectively and guide them to choose the correct technology for their aims of use.

Continuing on the theme of best practices, Software Outlook has also been developing best practice guides on (i) Version Control Systems, (ii) Continuous Integration, Deployment and Delivery, (iii) Documentation Tools and Best Practices, and (iv) Code Testing. The latter is currently in the last stages of being delivered but the others are available via the Software Outlook website. By following best practice guidelines, CCPs and HECs should be able to use their time more effectively and increase the quality of their outputs. Some members of CoSeC have reported that they are already seeing the benefits of using the advice from the guides.

As computer simulations are targeted to more and more complex problems, simulation codes are needing to be coupled together to capture different features and components: there are a number of code coupling libraries available to do this. Whilst each code being coupled might exhibit great HPC scalability, it can be easy to lose this scalability when codes are coupled together but it is important that this loss is as limited as possible. Software Outlook has identified a number of code coupling libraries and has started to implement them within a testing framework. As well as comparing performance in terms of HPC scalability, we have also compared their ease of use and the types of coupling that they are best suited to. We have seen that there is a large difference in user experience when trying to incorporate them within our framework. The outputs from this work will guide the CCPs/HECs in their choice of coupling library and, since multiple CCPs are actively considering coupling, this work is being done once instead of multiple times, which is better value within CoSeC.

In addition to coupling of codes, simulations are becoming ever more detailed and the linear algebraic systems that need solving are becoming harder and harder to solve due to time and memory limitations. Software Outlook has been exploring different preconditioning strategies available and is in the process of populating the website with this information to help guide the CCP/HEC consortia. Effective preconditioners can speed-up CCP/HEC codes by several orders of magnitude in terms of compute time, which will allow more simulations to be performed or more complex simulations to compute. Effective preconditioners can also result in increased accuracy within the simulations. As a result of the work we had already done within Software Outlook, we were successful in bidding for an ExCALIBUR Project to investigate preconditioners for use within UKAEA's NEPTUNE Project. Rather than duplicate the effort of comparing preconditioners for use within BOUT++ and other CCP codes, the Software Outlook work has moved towards a knowledge gathering activity to help guide the CCPs. The results from the NEPTUNE preconditioning project will be communicated to the CCP communities.

In the past decade, GPGPUs have moved from being a fairly niche item within a compute node to being widely available. At the same time, a number of frameworks have evolved that can be used within software to allow the GPGPUs to be utilised. The different frameworks range widely in terms of ease of code porting and how effectively the GPGPUs are used. Does a relatively fast to port framework mean that the GPGPUs will not be used as well? Is it worth spending the time to use CUDA? We have been using a CCPi code to explore the different frameworks so that CCPs/HECs are able to choose the best framework for their situation, which will allow the CCP/HEC to use its resources more effectively. The report/training document is in the final stages of development and will be published in April 2021. Given this is a problem for a number of the CCPs/HECs, Software Outlook is ensuring that the same work package is not being replicated multiple times within CoSeC.

Significant Impact (1 April 2020 – 31 March 2021)

It is hard to identify the most significant impact because the different work packages should all have large impacts. However, in the last year, the most significant impact of the work has been the best practice guidance reports on version control systems, code containerisation, code documentation, testing, and continuous integration, deployment and delivery. The true impact of this work is only just starting.

Outline Plans 2021/22 (1 April 2021 – 31 March 2022)

(Detailed plans and milestones are included in the Technical Appendix for reference)

Software Outlook will focus on four main areas of work: a software audit to allow CoSeC and, more generally, UKRI to understand the changing nature of CCP/HEC codes and the needs of the communities; evaluation of a wide range of quantum computing software frameworks and simulators to identify whether any of these would be useful for some of the CCPs and HECs, as well as the curation and evaluation of a selection of quantum computing algorithms; an investigation into the applicability and usefulness of recently developed reduced precision AI chips, which follows on from previous reduced precision work and also extends the GPU framework analysis from 2020/21 work; provision of best practice guidance on software maintenance and development when underlying base code dependencies are updated. Time has also been included to support CCPs and HECs as they incorporate the guidance that Software Outlook has developed during 20/21 and 21/22. Software Outlook received many more work package requests than funding allows but, with this work plan, we aim to support as many communities as we possibly can within the constraints that we have.

Resource Planning and Yearly Costing

For this report the information on resource planning and yearly costing will be provided in a separate document to be circulated separately by CoSeC Director Barbara Montanari.

Size of Communities Supported

The following table shows the community size of each of the supported CCPs and HECs in terms of the number of members of each community.

CCP/HEC (Main supported code)	Start Date	Community Size (Members)	Allocation (FTEs per year)
CCP5 (DL_POLY etc.)	1980	1050	0.00
CCP9 (Questaal)	1981	450	3.00
CCP-Mag (KKR)	2015	44	0.00
CCP-NC (MagresView)	2011	60	1.40
CCPQ (R-Matrix, TNT, Quantics)	2011	150	0.00
CCPPlasma (GS2 / BOUT++)	2007	150 (with HECPlasma)	0.00
CCPi (CCPi CIL)	2012	380	1.50
CCP SyneRBI (SIRF)	2015	80	1.80
CCPBioSim (FESetup)	2011	402 (with HECBioSim)	1.95
CCP-QC	2019	New Community	0.60
CCP-WSI+	2019	169	2.00
CCP Turbulence	2019	New Community	1.75
CCP NTH	2019	New Community	1.00
MCC (CRYSTAL, Chemshell)	1994	464	2.00
UKCP (CASTEP)	1990	150	1.00
UKCOMES (DL_MESO)	2013	150	0.60
HEC-Plasma (GS2, BOUT++)	2013	150 (with CCP Plasma)	0.20
HECBioSim (Longbow)	2013	402 (with CCP BioSim)	0.80
UK-AMOR (R-Matrix)	2018	40	0.20
UKTC (Code_Saturne)	2018	47	0.40
UKCTRF (SENGA+)	2019	47	0.50

Metrics

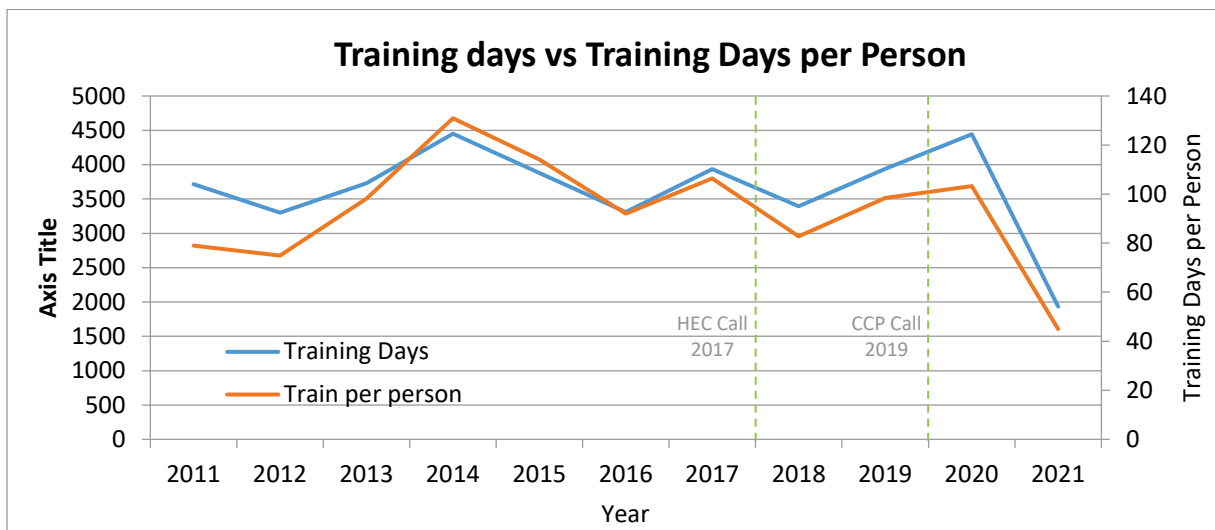


The metrics currently used for this programme are defined as:



- Number of citations in peer-reviewed journals of a publication about software supported by SLA-funded staff. Please note that not all software packages we support have a citeable publication.
- Number of training days delivered by SLA-funded staff. This metric measures outputs, ie how many people were trained and over how many days.
- Number of publications in peer-reviewed journals authored by SLA funded staff.
- Number of scientific/technical presentations at external events delivered by SLA-funded staff.

The metrics are per financial year, with the exception of the citations metric, which is for calendar year (for ease of data collection). Please note that for an interim report the data collection for the current year is partial for all metrics except the citations.

In the metrics graphs below, the blue line refers to the total number while the orange line refers to the number per scientific/technical staff member. This is useful because metrics are included for any person

irrespective of whether they are partially or fully funded by the SLA. For information, a graph of the scientific/technical staff headcount has also been included.





Metrics breakdown by area

	FTE	Publications	Presentations	Training Days	Citations
CCP5 (including DL_codes)	0.00	0	0	0	0
CCP9	3.00	4	0	250	289
CCP-NC	1.40	3	1	0	62
CCPQ	0.00	1	0.5	17.5	58.5
CCP Plasma/HEC Plasma	0.00	0	0	0	87
CCPi	1.50	0	2	68	0
CCP SyneRBI	1.80	1	0	20	6
CCP BioSim/HEC BioSim	2.75	3	2	1354	19
CCP-QC	0.60	0	1	0	0
CCP-WSI+	2.00	1	2	0	77
CCP Turbulence	1.75	0	1	40	0
CCP NTH	1.00	1	2	52	0

MCC (including ChemShell, CRYSTAL)	2.00	5	5	7	595
UKCP (including CASTEP)	1.00	1	0	0	1218
UKCOMES	0.60	0	5	79	18
UK-AMOR	0.20	1	0.5	17.5	58.5
UKTC	0.40	0.5	1	0	0
UKCTRF	0.50	0.5	0	0	35
Software Outlook	1.50	4	2	29	3
Totals		26	25	1934	2526

NOTE: fractional figures arise due to cross-project activities where the metric is shared between multiple projects.