



EPSRC Service Level Agreement with STFC for Computational Science Support

FY 2018/19 Annual Report
(Covering the period 1 April 2018 – 31 March 2019)

May 2019

Table of Contents

| | |
|--|----|
| Background | 3 |
| CoSeC Project Office | 6 |
| CCP5 – Computer Simulation of Condensed Phases | 6 |
| CCP9 – Electronic Structure of Solids | 8 |
| CCPmag – Computational Multiscale Magnetism..... | 10 |
| CCPNC – NMR Crystallography | 11 |
| CCPQ – Quantum Dynamics in Atomic Molecular and Optical Physics..... | 12 |
| CCP-Plasma – HEC-Plasma Physics | 14 |
| CCPi – Tomographic Imaging | 15 |
| CCP PET-MR - Positron Emission Tomography (PET) and Magnetic Resonance (MR) Imaging..... | 16 |
| CCPBioSim - Biomolecular Simulation at the Life Sciences Interface | 18 |
| MCC – Materials Chemistry Consortium..... | 20 |
| UKCP – UK Car-Parrinello Consortium | 21 |
| UKCOMES - UK Consortium on Mesoscale Engineering Sciences | 22 |
| HEC Plasma Physics | 23 |
| HECBioSim | 23 |
| UK-AMOR | 25 |
| UKTC - The UK Turbulence Consortium | 25 |
| UKCTRF - UK Consortium on Turbulent Reacting Flows..... | 26 |
| Software Outlook..... | 26 |
| Resource Planning and Yearly Costing | 28 |
| Metrics..... | 30 |

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 |
|---------------------|--|---------------------------|-------------------------|
| CCP5 | The Computer Simulation of Condensed Phases. | Prof Neil Allan | Dr Alin Elena |
| CCP9 | Computational Electronic Structure of Condensed Matter | Prof Stewart Clark | Dr Leon Petit |
| CCP-mag | CCP on Computational Magnetism | Prof Julie Staunton | Dr Martin Lueders |
| CCP-NC | NMR Crystallography | Dr Jonathan Yates | Dr Simone Sturniolo |
| CCPQ | Quantum dynamics in Atomic Molecular and Optical Physics | Prof Graham Worth | Dr Martin Plummer |
| CCP-Plasma | The Plasma-CCP Network | Prof Tony Arber | Dr Joseph Parker |
| CCPi | Tomographic Imaging | Prof Phillip Withers | Dr Edoardo Pasca |
| CCP-PETMR | Computational Collaborative Project in Synergistic PET-MR Reconstruction | Prof Kris Thielemans | Dr Evgueni Ovtchinnikov |
| CCPBioSim | Biomolecular simulation at the life sciences interface | Prof Adrian Mulholland | Dr Tom Keal |
| Materials Chemistry | UK Materials Chemistry Consortium | Prof Richard Catlow | Dr Tom Keal |
| HECBioSim | 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 |

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:

| Community | Core support per project (FTEs per annum) |
|-----------|---|
| CCP5 | 3.4 |
| CCP9 | 2.6 |
| CCP-mag | 0.8 |
| CCP-NC | 1.4 |
| CCPQ | 2.0 |

| | |
|---------------------|----------------------------|
| CCP-Plasma | 0.8 |
| CCPi | 1.3 |
| CCP-PETMR | 1.25 |
| CCPBioSim | 1.25 |
| Materials Chemistry | 2.5 (2 from 31 Oct 2018) |
| HECBioSim | 1.0 (0.8 from 31 Oct 2018) |
| UKCP | 1.0 |
| HEC Plasma | 0.2 |
| UKCOMES | 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

Summary Report (1 April 2018 – 31 March 2019)

The major activity of the project office during this reporting period has been to assist with the preparation for the mid-term International Review. This included collating information for the SLA report submitted to the review panel, preparing impact studies for each of the EPSRC funded CCPs and HECs, and assisting with the production of the posters presented at the Review in London. The project office produced the CoSeC annual report for the financial year 1 April 2017 to 31 March 2018. This report was submitted to EPSRC in June and discussed at the SLA Steering Committee meeting in July. An interim report covering the period 1 April to 30 September 2018 was also produced and submitted to the Steering Committee in January 2019.

Ongoing interaction with Technopolis around the production of the commissioned SLA impact report continues with the final report due for release in Summer 2019. Dawn Geatches joined CoSeC in January 2019 in the role of Impact Manager previously occupied by Marion O'Sullivan.

Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

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

Moving forward into 2019/20 the project office will continue with its normal day-to-day activities but will also help to co-ordinate CoSeC involvement in EPSRC's call for CCPs and implement the recommendations made by the panel at the CoSeC international review that took place in October 2018. Activities to expand on the CoSeC impact studies will continue with the introduction and implementation of the CoSeC Impact Awards.

CCP5 – Computer Simulation of Condensed Phases

Summary Report (1 April 2018 – 31 March 2019)

DL_POLY 4.09 was released in September 2018, followed by two bug releases in November 2018 and January 2019. One new important feature added is the two temperature thermostat model. This feature introduces a new way of simulating radiation damage in materials allowing users to explore phenomena not possible before. The two temperature model was applied already to a selection of problems and has been applied to high-energy cascades in, amongst others, iron, zirconia and tungsten. It has also been applied to laser irradiation of gold and tungsten films, as well as swift heavy ion irradiation of germanium, silicon and various metallic structures, mainly by group of Professor Dorothy Duffy at UCL. At the community request new potentials were implemented a number of five allowing faster preparation of simulations and minimizing the possibility of errors during preparatory phases..

An alpha version of DL_POLY was released in October 2018. The new version has the advantages of pure modularisation, thread safe data by design and separation of data and computation. The majority of these changes have the advantage of easier and sustainable development of the software. The Continuous Integrations (CI) system was strengthened with testing increasing from 28 tests to 169 tests. From the user point of view the direct impact is speed, the code is now faster up to 25% for certain problems. Further software developments happened in parallel for addition of new features with Dr Alex Buccheri at the University of Bristol, Dr Oliver Dicks at QMUL and Dr Jacob Wilkins at University of Oxford. Ivan Scivetti's work on Empirical Valence bond implementation and Forward Flux Sampling continues. All these additions will greatly expand the scope of the DL_POLY code for the benefit of the users.

On user accessibility a new set of graphical utilities were released written in Python 3. These applications were received very well by users. A Jupyter notebooks infrastructure was released helping user to create easily workflows for their data analysis pipelines. Some of the GUI applications were appreciated by Michael Seaton who adopted them for DL_MESO.

Work on new DL_FIELD features allowed users to easily setup force field models for pure organic liquid and solutions of a specific concentration or density. The feature is particularly useful if one were to investigate drug molecules-solvent interactions and dissolutions using the newly introduced force field such as Amber GAFF. The bio-inorganic force field model setup has been improved by using simple xyz format (for both organic and inorganic components). Within DL_ANALYSER the introduction of the unique capability to the unpicking of various modes of atomic interactions and to carry out statistical analysis on hydrogen bond interactions, benzene aromatic pi interactions and hydrophobic interactions was completed. These interactions can be described by a standard, chemical-sensitive annotation system for easy references and potential cheminformatics processing. First Py-ChemShell beta (v19.0) released as a milestone version of the Python-based redevelopment and is ready for production calculations on materials systems with new training materials and documents

Cross-fertilisation between CCP5 and other CCPs continued this year by organising two joint workshops. The first was "Simulation and modelling of nuclear materials" at the University of Bristol, 4/5 June 2018 that was organised jointly with CCP9. The workshop was well attended with the key note speaker Blas Uberuaga coming from Los Alamos National Laboratory, Another very successful workshop "Forcefields: Status, challenges & vision" was organised by Leon Petit (CCP9), Alin Elena (CCP5) and Albert Bartok-Partay (UKCP). The aim of this workshop was to bring together experts in classical and ab-initio simulations and experimentalists, to get an overview of the different methodologies used for potential generation and usage and to take stock of the current state of the art. Discussions about the future development strategies in the field and collaborative ways of moving forward in partnership with our sponsoring communities, CCP5, CCP9, UKCP, MCC, CCPBioSim and CECAM-UK Daresbury were held. Fifty academics spent two days at Daresbury discussing Forcefields. New collaborations were forged and a follow-on international workshop is planned. A new DL_Poly Developers meeting took place at Daresbury Laboratory, 13-14 November 2018, this time was a successful hackathon style meeting, with participants from University of Manchester and CCP9.

The following workshops were sponsored by CCP5 6th Computational Molecular Science (CMS), University of Warwick, 27 -29 Mar 2019, New Horizons in Materials Modelling, 4th of January 2019, University of York. The Statistical Mechanics and Thermodynamics Group (SMTG) of the Royal Society of Chemistry Meeting in the University of Manchester, 9th-11th January 2019, Adsorption Summer School, 12 - 15 June 2018, University of Strathclyde, Glasgow, 3rd Conference on Multiscale Modelling of Condensed Phase and Biological Systems - CCPBioSim & CCP5, University of Manchester 21-23 May 2018, Physics by the Lake, 29 July - 10 August 2018 , Hermes Summer School, 19th - 23rd July 2018, Unifying Concepts in Glass Physics , 11th June 2018,

CCP5 Outreach competition winners were announced. Two 1st prizes ex aequo to Anna Akinshina, University of Salford and Anna Sofia Tascini, Imperial College London with the third prize awarded to Simon Bennie, University of Bristol. A follow up competition is organised for 2019.

The CCP5 Lecture tour for 2018 was given by Prof. Fernando Escobedo from Cornell University. The tour had visits at Imperial College, Manchester University, University of Bristol, University of Cambridge and STFC Daresbury Laboratory. The 2019 tour will be given by Prof. Roberto Car from Princeton University. The 2018 tour generated a lot of interest and we are currently in the process of selecting dates and hosts for 2019.

The CCP5 AGM conference organised at University of Manchester, 10-12 September 2018, was attended by over 80 participants, from UK and overseas, from academia and industry. The event

involved a selection of very high quality international speakers. The CCP5 AGM for 2019 will be organised at the London School of Economics, 16th -18th of September 2019.

CCP5 Summer Bursaries concluded with a number of nine students. Reports will be available on the website once all are received. This programme will continue in 2019 also.

Two new members were elected to the CCP5 committee, Dr Kostya Trachenko, QMUL and Dr Misbah Sanwar from Johnson Matthey. Effort is ongoing on keeping the website up to date and new outreach and impact sections were added. Alin Elena, Ilian Todorov and Kostya Tratchenko were successful in securing a grant with QMUL from the EPSRC Impact fund to develop new analysis tools in DL_POLY_4.

Ivan Scivetti and Alin Elena were invited to attend the National Physical Laboratory workshop, Computational Methods for 2D Materials Simulations, 30th of July 2018, where Ivan gave a presentation. Ivan has also been invited to give a talk at Universidad Autonoma de Madrid, in Q2 2019.

Tom Keal and You Lu supported the CCP5 summer project student Sam Watts at Cardiff University (supervised by Andrew Logsdail) to develop zeolite modelling tutorials for Py-ChemShell. Sam visited Daresbury Laboratory on 31 July for ChemShell training. Mala Alhaji (Cardiff University) visited You Lu and Tom Keal at Daresbury from 23-25 July to learn QM/MM modelling of materials with ChemShell. You Lu hosted summer project student Yikai Qian (Imperial) for two weeks in September to learn ChemShell and create prototype Jupyter notebook tutorials for the code.

Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

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

To provide more scalable electrostatic interaction calculations with local polarizability effects, the Groot DPD electrostatics scheme based on the particle-particle particle-mesh (PPPM) algorithm will be implemented in DL_MESO. DL_FIELD support for xyz for inorganic models by further expending DL_F notation to inorganic models. New refactorised DL_POLY with fully modular structure.

Py-ChemShell will be extended through an interface to the ADF code in collaboration with SCM in Amsterdam, further development of the Aten GUI plugin for graphical setup of materials calculations, support for Z-matrix coordinates and a new testing system. A further release is scheduled for Q1 2020. The CCP5 Summer School moves to University of Durham, organize with the transition. Organise a cross community effort about forcefield parametrisation, initially UKCP and ISIS.

CCP9 – Electronic Structure of Solids

Summary Report (1 April 2018 – 31 March 2019)

Support for the Questaal Project: The validation and verification project compares the accuracy of the full potential lmf code (part of Questaal) systematically with other all-electron and pseudopotential DFT implementations. Work has started on extending this study to the lanthanides elements, specifically comparing Questaal and CASTEP results. The HDF5 implementation developed during the previous ECSE has been ported to the current version of the code.

The Questaal hands-on course was organized May 21-25 at Daresbury Laboratory. Thirty participants from both the UK and further afield took part. The course was organized in collaboration with the code developers from King's College. Lectures by invited speakers as well as tutorials (including one by Jerome Jackson) were given, introducing the students to the capabilities of the code.

The workshop "Forcefields: Status Challenges & Vision" was organized at Daresbury Laboratory, 28-29 January 2019. The aim of this workshop was to bring together experts in classical and ab-initio

simulations and experimentalists, and was organized with CCP5 as a cross-CCP event with co-sponsored by UKCP, MCC, CCPBioSim and CECAM-UK.

Research on correlated electron systems: Leon Petit is involved in a theory-experiment collaboration with Warwick University and Ames laboratory (US), with the goal to probe the interactions that govern the rare-earth/transition-metal permanent magnets and to study the corresponding magnetic ordering temperatures. Work on Fe₃O₄, CeOs₄Sb₁₂, and Eu₂In is currently ongoing. An article investigating the competition of different magnetic interactions in Gd-intermetallic alloys is being finalized. Jerome Jackson is co-author of a paper on magnetic excitations in Mn₃Sn together with collaborators from Hungary and South Korea, and that has recently been published in npj Quantum Materials. A follow-up project to study the helical ordering in Mn₃Sn is ongoing, and a publication is planned on the incommensurate ordering in CrB₂.

The CCP9 Young researchers event and Community meeting was held 18-19 July 2018 at the Park Inn Hotel, York. With over 70 participants, the event figured both technical presentations and application talks from recently appointed academic staff as well as a number of discussions concerning the future direction of CCP9 and its working group.

Widening participation workshop: Advanced Topics in Cs-corrected STEM and Spectroscopy: Theory meets Experiment, Daresbury Laboratory, 29 June 4 July 2018. The workshop successfully brought together experiment and theory, with both hands-on courses in microscopy and KKR based computational modelling. The workshop was arranged by Jerome Jackson, Leon Petit, and Martin Lueders, together with the SuperSTEM team.

Crystal-field parameters in Imf: Collaboration with Salford University is on-going. A meeting was held in June to discuss the progress made both on the experimental and theoretical fronts.

Psi-k: CCP9 provides support for the European electronic structure network Psi-k in the form of finance administration and editing the Psi-k scientific highlight. Six highlights were edited during the reporting period, and submitted to the mailing-list which now reaches more than 3500 people. The CCP9 website was up-dated with a new-mailing list (using listserv e-mail list management service) and that is in line with GDPR regulations.

The current CCP9 Flagship Project - Excitations in Complex Environments: Multiphysics embedding for large scale electronic structure. This project is just past its halfway stage and they have made excellent progress in expanding the capabilities of the code and making it more useful for the CCP9 Community. We ran a ONETEP coding retreat near Peterborough in 2018, which finalised the release of version 5.0 of the code, which is currently available to all CCP9 members. The project has recruited three PDRAs: Dr Joe Prentice at Imperial, who with Robert Charlton has worked on functionality enabling DFT-in-DFT embedding calculations (WP1); Dr James Womack at Southampton, who with Jacek Dziedzic has worked on dramatically increasing the speed and scaling of exact exchange calculations in DFT and TDDFT (WP2); and Dr Jolyon Aarons at Warwick, who has worked on excited state forces (WP3). All three work packages have resulted in code improvements which they plan to make available in upcoming Academic Releases v5.2 and v5.4. Following a successful ONETEP Masterclass in 2017, funded by the CSE Flagship project and by CCP9, they are now running another event in August 2019 at Warwick.

The Questaal Flagship project has run out at the end of September. Brian Cunningham and Myrta Gruening succeeded in including ladder diagrams to improve W in the QSGW cycle in insulators. This largely eliminates QSGW's tendency to overestimate semiconductor bandgaps; now the predicted gaps (when electron-phonon interaction is taken into account) is on the order of 0.1 eV of experiment for a wide range of systems. Swagata Acharya and Francois Jamet have made significant progress in combining QSGW+DMFT. They now have the ability to compute dynamic charge and spin susceptibilities, and they have tested the approach on a number of correlated systems. Where

experiments are available agreement with them seems to be very good. Dimitar Pashov has a working tight-binding transformation of the conventional Questaal basis. His JPO basis, an enhancement of this basis that makes the kinetic energy everywhere continuous, is expected to improve the basis quality and be close to the the best possible basis to solve Schrodinger's equation to a prescribed accuracy for a given rank of hamiltonian.

International collaboration and engagement: Leon Petit visited Arthur Ernst at Halle MPI and Julie Staunton at Warwick University. Jerome Jackson visited collaborators at the University of Budapest, giving an invited talk at each of these institution.

Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

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

The collaboration with Julie Staunton and Ames laboratory on permanent magnets will continue. Assuming the STEM proposal is successful, a postdoctoral research assistant will be hired to work on the project in collaboration with people at SuperSTEM, Belfast University, and King's College.

CCP-mag – Computational Multiscale Magnetism

Summary Report (1 April 2018 – 31 March 2019)

Dr. Kun Cao is now employed full time on this project. He is working on magnetic exchange couplings and also on the magnetic excitations project with the facilities.

As part of the common data format project, Martin Lueders visited Micael Oliveira at the Max Planck Institute for structure and dynamics of matter in Hamburg, where the new implementation of the library was continued and, in particular, the unit testing suite was updated to capture the new functionality. This new implementation makes the library development independent of the actual data to be stored in the files, and hence makes the library more versatile. This activity is done in strong collaboration with ECAM and will benefit the wider electronic structure community.

Kun Cao has started the implementation of the magnetic exchange constants J in the full-potential plane wave pseudopotential method (Quantum Espresso). When finished, the work will be used as a prototype for a planned CASTEP implementation. The calculated J 's can then act as input for atomistic/micromagnetic simulations or as guidelines for experimentalists to fit their data.

Within the project in collaboration with the ISIS neutron Facility (which has leveraged STFC funding), the full 4D calculations for iron have been presented by Martin Lueders at the IoP Magnetism conference in Manchester on the 9th of April. Further calculations, based on Kun Cao's implementation in Quantum Espresso have been completed and will be included in the planned publication. Discussions with Toby Perring showed an increased interest in larger systems, and strongly correlated systems. To this end, further work on parallelising the code to calculate spin susceptibilities has started, which will allow applications to more complex systems, and should eventually allow such calculations to be routine components of experimental studies.

Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

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

In the period 2019/20 we will continue the programme of workshops, hands-on courses and support of the IOP Magnetism conference. We plan to finish the implementation of full potential exchange couplings and interface an ab initio electronic structure code via the library with the VAMPIRE code.

We plan to liaise with Keith Refson to implement TDDFT for spin fluctuations in CASTEP since an implementation has already been done by Kun Cao in a similar code (Quantum Espresso).

The collaboration with ISIS will continue and we will investigate methods beyond the so-called rigid spin approximation and investigate systems beyond elemental transition metals.

The plans have been slightly corrected with respect to the earlier version in order to accommodate the fact that Martin Lueders has left STFC and is no longer contributing to CCP-mag.

CCP-NC – NMR Crystallography

Summary Report (1 April 2018 – 31 March 2019)

Maintenance and development of the Soprano Python library, developed by Simone Sturniolo, is ongoing, leading to version 0.8.1. Recent upgrades and fixes included improving treatment of quadrupolar NMR spectra and support of the machine learning library sklearn for some new functionality. A paper making use of Soprano for the analysis of hydrogen defects in zeolites authored by David McKay, Simone Sturniolo et al. has been published in the Journal of the American Chemical Society.

The effort and milestones required to develop a revamped version of MagresView has been assessed. The first stage of development is complete for `crystcif-parse`, a JavaScript Node module for parsing crystallographic CIF files, making it functional and available on GitHub for public use. Though still barebones, this module provides a functionality until now unavailable in this language. Development has also begun on `crystvis-js`, a crystal structure 3D renderer making use of Three.js and WebGL technology. This is intended as a more efficient replacement for JMol's rendering in MagresView 2.0.

The CCP-NC database website is also ready for public release. This has involved both technical improvements and the clarification of the legal issues involved – licensing of the data and privacy and cookie policy. The process of writing and reviewing these has been carried out for all CCP-NC related websites in light of the entry into force of the General Data Protection Regulation. This work has been carried out with the help of Ben Taylor-Gadd from STFC. The final choice has been to make the database and its contents available under the Public Domain Dedication and License v1.0.

A series of new tools have been developed for computational support to the muon spectroscopy community allowing them to complete their work quicker using simpler interfaces, and are now available on Github under the name of "pymuon-suite". While developed in the framework of a larger cross-topic collaboration with the ISIS Muon Group, many of these tools are of potential interest to the NMR community as well, especially concerning Random Structure Searching for hydrogen in materials and treatment of quantum nuclear effects. For the latter, thanks to the combined efforts of Adam Laverack and Simone Sturniolo, an improved sampling method based on harmonic frequencies has been implemented, and the inclusion of more sampling methods is planned for the future. An Unperturbed Electrostatic Potential method making use of CASTEP's electronic density files has been developed too, and it provides as well a nuclear-independent hyperfine tensor functionality, which is the first example of its kind that we're aware of.

CASTEPconv, a convenient companion tool to CASTEP used to automate convergence calculations, has been refactored to make use of the ASE library, streamlining the code and making future extensions easier. This has been released as CASTEPconv 2.0.

Meta-GGA implementation in CASTEP reached the stage that a large part of the standard functionality of CASTEP is accessible, including molecular dynamics, geometry optimisation and NMR parameters. Code release has been impeded by the fact that fundamental problems have been identified within the

flagship functional, SCAN, which was confirmed by the original developers. A revised SCAN functional has been developed.

Two scientific papers have been submitted, to Phys. Rev. B and J. Chem. Phys., which are currently under review. The latter is expected to be accepted imminently as the required revisions were very minor. A new geometry optimisation in CASTEP preconditioner is currently benchmarked using difficult, real-life problems collected from CCP-NC members, and the results will be included in a paper.

Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

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

By 2019, we plan for the CCP-NC database service to be available to the public and fully functional. It will therefore be possible to undertake projects that make use of it, as well as re-direct effort towards developing other software. One avenue that can be pursued is the automation of a pipeline that computes NMR parameters for known crystalline structures in existing databases and uploads them to the CCP-NC one. There are also plans to rewrite MagresView into an updated 2.0 version which would make full use of the modern HTML5 and JavaScript ecosystem to provide a more maintainable code and better performance. Finally, the new opportunities opened by the implementation of SCAN in combination with dispersion correction schemes in CASTEP will require extensive testing and exploration, and may open up spaces for new collaborations.

CCPQ – Quantum Dynamics in Atomic Molecular and Optical Physics

Summary Report (1 April 2018 – 31 March 2019)

CoSeC CCPQ support work for 2018-19 followed the planned milestones with some changes due to circumstances and new milestones added. Martin Plummer took part in steering panel planning meetings for the new R-matrix UK_AMOR HEC, as well as its official inaugural members' scientific meeting in September, and (with AGS) in training days and workshops for the R-MADAM Software Flagship project. CoSeC general administration supported the international AttoFel conference for attosecond and free electron laser science (UCL) in July, handling the registration and web-based administration.

Following the announcement of the closure of CCPForge, we worked with colleagues to help ensure a smooth transition of repositories to various GitLab sites. The full diverse set of R-matrix codes was moved to the Belfast based site which initially housed the RMT code (mirrored to CCPForge). While each code suite would have its own administrator, CoSeC (MP) will assist as overall administrator if/as required. The SEG group's support for the Oxford/Bath TNT package mainly concerned the move of the (large and complex, with both full code and automated test suite trees) repository from CCPForge to Oxford. The various projects linked to the ANVIL continuous integration tool have had log-in credentials changed from CCPForge. STFC graduate Tom Dack's 6-month assignment to SEG involved testing of new parallelism in the QUANTICS code, as well as transfer of links to GitLab and detailed testing of QUANTICS on ANVIL. ANVIL QUANTICS will be one subject of a forthcoming ECAM workshop on quantum dynamics. TD also set up the R-matrix RMT code on ANVIL with follow-up work to be continued by SEG as required.

As part of general support, CCPQ input was provided for the CoSeC international Review and the website has been kept up to date, including a new page detailing recent software-related publications and links to the new version of the molecular collisions package MOLSCAT (JM Hutson, Durham, and co-workers). Martin Plummer has begun initial negotiations to organize a joint meeting with other CCPs

and HECs (CCP-PLASMA/HEC-PLASMA) and related industries (Quantemol Ltd) on applications and use of AMO physics calculations.

The atomic R-matrix double-continuum code project has made solid progress, with both a code and a CPC publication to be delivered in the near future following some final optimization work. The coding structures and techniques were presented at the UK-AMOR HEC inaugural meeting in September, along with other CoSeC support developments, and will also, subject to acceptance, be presented at the International Conference on Electronic, Atomic and Photonic Collisions (ICPEAC 2019) in July.

The UKRMol(+) memory optimization project (Andrew Sunderland) continued with major work expanding the re-indexing optimization from combinations of bound and continuum atomic orbitals to the more complicated combinations of bound and continuum molecular orbitals. The revised code used 1/13 of the memory of the original code in a test case (the atomic re-indexing used 1/8 of the memory in a large test case, saving ~2.8GBytes). Final debugging of the Fortran 2003 code to achieve exact reproduction of results has concluded and the optimized code is being added to the main trunk of the UKRMol(+) repository. Also, the 'heavy-particle' version of the R-Matrix code PFARM has been included as a part of UCL's new RMat_REACT suite, with the UCL group now using the code. Martin Plummer is working with the group to extend the range of multichannel calculations available (with the complication compared to the atomic case of recognising body-fixed and space fixed reference frames). Plans for further work to be carried out in tandem with HEC UK-AMOR support were made, involving adapting the TIMEDELn parallel framework to handle efficient studies of bound states: work will begin in earnest in the next reporting year. AGS has continued overseeing PRACE work on GPU-PFARM, and is also advising QUB about GPU-BLAS routines in RMT.

Martin Plummer was involved in antimatter-based discussions at UCL, Nottingham, in particular regarding use of properties of binomial coefficients to check accuracy of integral transformations needed for efficient calculation of rearrangement collision parameters, and with MM Law (Aberdeen) and student Lara van Lydon on antihydrogen calculations involving H₂ collisions: MP noted some similarities between LvL's preliminary work with antiproton-protonium 'bound' states to studies of Ps-, and to earlier positronic resonance studies by Martin Plummer and EAG Armour which should assist the Aberdeen work.

Oultine Plans 2019/20 (1 April 2019 – 31 March 2020)

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

In 2019-20 CoSeC support will continue for the main CCPQ code packages. Direct coding support (and associated scientific collaboration) will be given to the various atomic and molecular R-matrix codes: RMT, UKRMol+ and the 'in-house' PRMAT (PFARM) packages, plus CCPQ support will be combined with additional HEC CoSeC funding to support further development of the Rmat_React ultracold chemistry resonant collisions package. The atomic double-continuum code should in this year be fully evaluated and passed on to users (with support) in the UK-AMOR HEC consortium. The support will help to maintain these packages as internationally leading scientific codes in their areas, and with the assistance of UK-AMOR, extend the codes' impact. The Software Engineering Group's involvement will continue to provide general support for TNT and the Quantics package, and for RMT and UKRMol(+) as required by CCPQ, and appropriate CoSeC support will also be given to other members of the diverse CCPQ community with Steering Panel approval. The types of specific support will follow on from the 2018-2019 work with review throughout the year. We will support funding applications for and related to CCPQ as they arise, including the CCPQ renewal call when it is announced, and maintain and promote the website and mailing lists. CoSeC's general support (currently Damian Jones) will provide expertise in setting up registration and payment for workshops. We will ensure that the support provided includes interactions with various university PhD students and PDRAs to provide effective (on the job) training.

Plasma Summary Report (1 April 2018 – 31 March 2019)

The project is proceeding as planned with 1FTE from Joseph Parker divided between the codes GS2 and BOUT++.

Minimizing plasma turbulence in tokamaks (nuclear fusion devices) is vital for achieving fusion, but simulations of fusion plasmas are computationally very expensive. GS2 is a plasma turbulence code which has been developed since the 1990's and is already highly optimized; however further optimizations are required for simulations to include all physical effects and to resolve necessary space or time scales. Moreover, code modernization is necessary for GS2 to remain a relevant and accessible tool for fusion research.

In this reporting period, we have assessed the performance of the new “split domains” data layout implement in the previous reporting cycle. We studied the scaling behaviour of different variants of this data layout in one part of the code, the calculation of the nonlinear term (a Fourier convolution). We found the expected scaling behaviour, and in doing so, uncovered performance bugs that slowed both the new and the existing implementation. With these fixed, the new layout is faster than the old layout for all relevant core counts, and is ~10% faster at the optimal core count. This work complements ongoing work on the field solve, undertaken by OeRC and managed by STFC.

In this period, we also undertook a major overhaul of GS2's organization. GS2 now has a “Governing Committee” of six members (from STFC, CCFE and the Universities of Oxford, York and Maryland) to make organizational decisions. We migrated repository hosting from Sourceforge to Bitbucket, and version control from svn to git. We also implemented a “gitflow” workflow, with protected stable and development branches, and code review and continuous integration (with SESC's Anvil Jenkins server) for pull requests. These changes have had a strong positive impact on the quality of code, the user experience, and the level of community engagement. We have made a release of changes since the migration (v8.0.1), and have accepted a further 700 commits to the development branch.

Achieving fusion also requires an understanding of plasma behaviour at the tokomak edge, for which BOUT++ is a widely-used code. In this reporting period, we further improved the standardized looping method developed in the previous cycle. We implemented a new “Region” class which integrated with this looping method to give a great deal of flexibility in the exact domain to be iterated over in many operations. We also increased the number of places in the code which could take advantage of the vectorization and OpenMP parallelization. We wrapped this approach to looping in a standard macro, making it easy for developers to fully exploit these improvements and make future changes to the method. This change was included in the v4.2.0 release. Overall, test problems in that release are up to 40% faster than in the previous release.

Simulations of the tokomak edge require inversion of elliptic operators to obtain the eletrostatic potential from the vorticity and density. The operator is often assumed to have no component parallel to the magnetic field lines, and is inverted as a 2D operator on uncoupled perpendicular planes. Although often justified, this approximation is unsuitable when the coordinate system is strongly non-orthogonal, as in the tokomak edge. In collaboration with CCFE, we have extended BOUT++'s inversion routines to include parallel coupling, and also to relax the Boussinesq approximation that the density must be constant. This will allow CCFE to undertake much more realistic, fully three dimensional simulations of turbulence in their tokomk plasmas.

Plasma Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

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

GS2 solves the gyrokinetic equations, which describe plasma turbulence on relatively short timescales. To model turbulence in tokomaks on timescales comparable with the plasma confinement time, GS2 is coupled to the transport solver Trinity. After our optimization work on GS2, it is likely that GS2 will no longer be the performance bottleneck in GS2/Trinity turbulence simulations. We anticipate that our optimization work on GS2 will continue into 2019, but thereafter, we plan to focus on optimizing the coupling between GS2 and Trinity, and on improving the convergence of Trinity's iterative method.

By 2019, BOUT++ users – particularly those at CCFE with whom we work closely – will be focussing on more realistic simulations of turbulence in the tokomak edge. These simulations will require new functionality in BOUT++, such as including kinetic corrections and non-local effects, and coupling to models to neutrals and plasma impurities. We will implement the new features that are necessary to support these simulations, and ensure that these do not have a negative effect on code performance and scalability. GS2 solves the gyrokinetic equations, which describe plasma turbulence on relatively short timescales. To model turbulence in tokomaks on timescales comparable with the plasma confinement time, GS2 is coupled to the transport solver Trinity. After our optimization work on GS2, it is likely that GS2 will no longer be the performance bottleneck in GS2/Trinity turbulence simulations. We anticipate that our optimization work on GS2 will continue into 2019, but thereafter, we plan to focus on optimizing the coupling between GS2 and Trinity, and on improving the convergence of Trinity's iterative method.

CCPi – Tomographic Imaging

Summary Report (1 April 2018 – 31 March 2019)

The team is focusing on the development of the Core Imaging Library (CIL), a toolbox of algorithms and tools for tomographic imaging. This includes algorithms for pre-processing, reconstruction, quantification, segmentation, and visualisation. This is the first attempt in the UK to improve the take-up of iterative reconstruction algorithms by the community. 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 CIL version 0.10.0 was released.

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 w.r.t. the “white beam” acquisition.

The main additions to the CIL are addressing the challenges introduced by low SNR data, and spectral (multichannel) tomography. The CIL now includes an implementation of a new reconstruction framework for iterative methods; ‘the regularisation toolkit’ for de-noising and regularisation during iterative reconstruction; and ‘TomoPhantom’ a software for generation of synthetic phantoms for CT to be used as benchmark of reconstruction algorithms.

In 4D tomography multiple 3D images of a sample are acquired whilst the sample is subjected to external forces or changes in status, e.g. temperature. The post-processing of the reconstructed data is a fundamental step for characterisation of samples and materials. To address this type of data analysis two packages are under active development, although not yet released. These are ‘Digital Volume Correlation’ DVC package which quantifies the displacements (stresses) between the sets of

3D images, and the 'Simpleflex' segmentation tool. The segmentation of 3D data is the basis for any quantification of the acquired/reconstructed data.

Effort has been spent to enable the use of the CIL reconstruction framework with the CCP PET/MR software SIRF. This allows the use of novel reconstruction algorithms in the PET/MR community.

Last but not least, 2 publications have been successfully submitted to SoftwareX detailing 'TomoPhantom' and the 'regularisation toolkit'. Additionally 2 conference presentations feature CIL/SIRF data.

In the reporting period, the core support effort (1.07 FTE) was provided by Edoardo Pasca, Gemma Fardell (from February 2019) and Tomas Kulhanek (October 2018 to February 2019) with external contribution by Martin Turner (University of Manchester, SCD Visiting Scientist).

Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

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

The Core team will focus on releasing the CIL with a 6 monthly stable release plan. The main additions to the CIL are expected to be with the Digital Volume Correlation (Brian Bay, USA) and development of algorithms and methods for rich tomography, such as multichannel CT. The team will investigate methods of machine learning for tomography.

The Core team will continue to give support to the Flagship team in the image reconstruction, and support the community (DLS, ISIS/IMAT) with their software. Distribution of the CIL software will be on conda channel and virtual machines running on the SCD Cloud based on DAaaS will be provided.

The core team will still provide support to the community via the website, mailing list and by organising conference and meetings.

CCP PETMR - Positron Emission Tomography (PET) and Magnetic Resonance (MR) Imaging

Summary Report (1 April 2018 – 31 March 2019)

Our work during the reported period progressed according to the job plan: software development and engineering efforts, adding content to our website www.ccppetmr.ac.uk, maintaining mailing lists (we now have 88 members on the CCP-PETMR announcement list, 18 on the developers and 60 on the users lists), organising working group and executive meetings, organising a series of well-attended Developers Days to present and discuss progress in our software framework development, and organising the first two CCP PETMR Hackathons (also attended by members of CCPi)

In the reported period we published two releases of our Open source software suite SIRF (Synergistic Image Reconstruction Framework): Release 1.0 (3 April 2018) and Release 1.1.1 (17 May 2018). These releases are based on the PET reconstruction package STIR (Software for Tomographic Image Reconstruction) and the MR reconstruction package Gadgetron and are capable of processing measured data from the Siemens PET-MR scanner, with work-in-progress to support the GE PET-MR system. We are working towards releasing SIRF 2.0 in April 2019. Highlights in SIRF 2.0 including the addition of the image registration capability (work led by the CCP PETMR Flagship project) and integration with the CCPi Core Imaging Library (CIL). CIL contains advanced regularisation techniques and recent algorithms for optimisation of non-smooth objective functions. CIL was written to largely adapt SIRF terminology and (Python) class structure, with some aliases now introduced into SIRF. This joint effort of CCP PETMR and CCPi was facilitated by joint CoSeC staff between the 2 projects. The

SIRF 2.0 SuperBuild will contain an option to easily install CIL in conjunction with SIRF for research to be able to exploit benefits of both software packages with very little effort.

SIRF distribution includes source code, installation instructions and scripts, test scripts, demo scripts and several layers of documentation. An Oracle Virtual Machine (VM) is also provided that has all the necessary software (except, for licensing reasons, Matlab) pre-installed for a quick start in any operating system that supports VMs. We now have a docker image and Terraform scripts for deploying on Azure available as well. All this software is available for free download on our public website www.ccppetmr.ac.uk and via github.com/CCPPETMR.

We currently use Travis for Continuous Integration testing. Work on adding Anvil is in progress.

As mentioned in our previous report, we have obtained official agreement from GE Healthcare to provide open source software for reconstructing of their PET data, directly from the raw data from the scanner, i.e. disclosing the previously confidential file format. While this capability is not yet available in SIRF, work on STIR at UCL and Leeds is progressing, with support from CoSeC staff.

Researchers at the PTB (Germany) have acquired a test-suite of MR data to test capabilities for Siemens MR scanners. Some issues were uncovered related to repeated sequences and orientations in the previous reporting period. These have now been addressed.

Our major highlights of this reporting period were our first two CCP PET-MR Hackathons. The first one was held at RAL 26-27 Aug 2018, and it was attended by 14 developers, including 2 CoSeC staff and 2 RAs from Manchester who are the main (non-CoSeC) CCPi developers. The hackathon was organised around 3 themes: “implementation of extra reconstruction algorithms”, “adding geometrical info on position to SIRF images”, “progressing GE Signa support in STIR”. Participants were divided in the 3 groups, with occasional cross-checks and CoSeC staff and the PI floating between groups. Noteworthy was also the good interaction with CCPi which resulted in increased alignment between the CCP PET-MR SIRF and CCPi CIL codes. Although the hackathon did not produce any “ready” code, a large leap was made in these 3 themes, and feedback from the participants was excellent. Moreover, we have since had continued contributions from the participants on the hackathon themes, with some contributions already merged into the STIR and SIRF master branches.

The second Hackathon was held at St Thomas’ Hospital in London 17-19 Dec 2018, and it was attended by 12 developers from UCL, KCL, Manchester and STFC, including 2 CoSeC staff. The Hackathon started with a half-day training for new SIRF users, and then proceeded to working on 3 themes: (i) Porting of a synergistic algorithm to SIRF, (ii) Resolving various issues with processing mMR data and (iii) Working on SIRF registration functionality. Attendees’ feedback was overall very positive. During the hackathon we created a new repository for user contributions using SIRF (<https://github.com/CCPPETMR/SIRF-Contribs/>) where part (i) of the output of the hackathon is now available (other parts were merged directly into SIRF itself). This serves as an example for non-experts how to implement a synergistic algorithm in SIRF (Python) and test it on real data.

We had some difficulties renewing the website certificates for ccppetmr.ac.uk and other sites due to change procedures. Plans are in place to prevent this happening again.

During this reporting period, CoSeC staff contributed to 6 conference submissions (all accepted) by CCP PETMR. One particular example is a submission led by CCPi to Fully3D 2019, the 15th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine. This submission is about using the CCPi Regularisation Toolkit but thanks to the integration between CIL and SIRF it also contained examples on PET data. We are now in the final stages of drafting a paper on SIRF. The drafting was led by CoSeC staff but contains contributions from many other CCP PETMR collaborators.

The progress of our development of SIRF over the reported period opens up significant opportunity for the user community to adopt or test the codes in a real PET-MR system. For the first time ever the PET-MR community has access to a software system that facilitates end-to-end PET-MR imaging methods testing, from pre-processing to reconstruction to post-processing, all under one software framework. This will accelerate the progress in the development of efficient reconstruction algorithms, which will translate into higher quality of PET-MR scanners' software and, eventually - and most importantly - to better quality and longevity of life for people affected by cancer, dementia and other serious illnesses.

Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

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

In 2019/20 we plan to continue our software development effort aiming at joint motion estimation, list mode reconstruction without conversion to sinograms, time-of-flight reconstruction and motion-guided reconstruction. We plan to have some of these features in SIRF Release 3.0, to be published in the last quarter of 2019. We will continue our engagement with the CCP PET-MR community by maintaining our website and mailing lists, organizing meetings, developers' days and other events. We will continue to organize and support training courses and developers' workshops and assist in new proposal writing. Our embedding within Institute of Nuclear Medicine at UCL Hospital for up to two days a week will continue, and similar arrangements with KCL will be sought. Last but not least, we will seek the continuation of EPSRC support and look for other sources of funding for our PETMR research and development efforts beyond 2020.

We plan to submit our paper "SIRF: Synergistic Image Reconstruction Framework" to the special issue of Computer Physics dedicated to its 50-years anniversary in Q2 2019.

The following papers will be presented at conferences in 2019:

- Richard Brown, Sam Ellis, Edoardo Pasca, Evgueni Ovtchinnikov, Ashley Gillman, Camila Munoz, James Bland, Abolfazl Mehranian, Claudia Prieto, Andrew J. Reader, Kris Thielemans, "SIRF: A research tool for rapid prototyping of PET-MR reconstruction", PSMR 2019, Muenchen, Germany, April 2019.
- Mayer J, Thielemans K, Brown R, Ovtchinnikov E, Atkinson D, Marsden P, Schaffter T, Kolbitsch C, "Flexible numerical simulation framework for dynamic PET-MRI", International Society for Magnetic Resonance in Medicine, Montreal, Canada, 11 May 2019 - 16 May 2019. 6235.
- Daniil Kazantsev, Edoardo Pasca, Mark Basham, Martin Turner, Matthias J. Ehrhardt, Kris Thielemans, Benjamin A. Thomas, Evgueni Ovtchinnikov, Philip J. Withers, Alun W. Ashton, "Versatile regularisation toolkit for iterative image reconstruction with proximal splitting algorithms", Fully3D 2019, Philadelphia, USA (Joint work with CCPi).

CCPBioSim - Biomolecular Simulation at the Life Sciences Interface

Summary Report (1 April 2018 – 31 March 2019)

Sarah Fegan continues the FESetup code review looking for ways to make the task of setting up and running alchemical free energy simulations more robust and easier to fit into larger workflows. Support for Amber18 is being added and a number of bugs have been fixed. Alongside this she has been working on a benchmark set of relative binding free energy calculations for testing free energy workflows and simulation protocols. The benchmark set includes 8 different proteins each with between 10 and 40 ligands. Sarah has also been working with the BioSimSpace developers on testing and making plans for future work. She attended the BioSimSpace co-location week in August in Edinburgh.

Work continues on developing multiscale modelling workflows for cytochrome P450 systems. Some of the files from the previous flagship work have been located, and we are filling in the gaps to make a fully reproducible workflow. This workflow should reduce the effort for setting up and running studies of drug binding, reactions and interactions in membrane-bound proteins with potential impact in drug discovery.

In collaboration with Charlie Laughton (University of Nottingham), Christopher Woods (University of Bristol) and James Gebbie-Rayet (HECBioSim), we have been developing training materials to be delivered through Jupyter notebooks deployed with a cloud-based infrastructure. The cloud-based training will help workshops to run without having to install special scientific programs on local machines making it easier to run training events around the country and online, and may potentially provide a best practice solution for other CCPs to follow.

Tom Keal led the organisation of the 3rd CCPBioSim/CCP5 Multiscale Modelling Conference at the University of Manchester on 21-23 May 2018. The event was held again at the Manchester Conference Centre, and eight invited speakers attended: Sarah Harris (UK, Leeds), James Kermode (UK, Warwick), Misbah Sarwar (UK, Johnson Matthey), Tomas Kubar (Germany, KIT), Florian Muller-Plathe (Germany, Darmstadt), Ali Karimi (Germany, Continental), Rommie Amaro (US, UCSD), Berend Smit (US, Berkeley). The event was well attended again with 73 people registering from the UK and overseas, helping to improve the international visibility of the CCP projects.

Tom Keal led a biomolecular QM/MM training course at St Andrews in June as part of the annual ScotCHEM symposium, with participants learning how to run Cytochrome P450 simulations using ChemShell. The workshop was sold out with 30 participants.

In November, Tom Keal and Sarah Fegan gave talks on biomolecular modelling at the Simulations for Experimentalists and Industrialists workshop held at Diamond Light Source. Discussions at the meeting have resulted in a new project in collaboration with CCP4 and CCP-EM to study the conformations and dynamics of glycosylated proteins.

Along with Sarah Harris (University of Leeds), Charlie Laughton and Christopher Woods, Sarah Fegan ran a training workshop 'Getting Started with Biomolecular Simulations' at the University of Leeds 10-11 January 2019. The workshop was fully booked with 38 people attending.

In February Tom Keal led a two day training course on QM/MM modelling with ChemShell at the University of Bristol..

Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

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

FESetup will continue to be maintained and developed for the benefit of the CCPBioSim community, with new functionality in the form of integration with the LOMAP package for automating the creation of morph pair lists and ease of use improvements concentrating on documentation and installation. Multiscale workflow development will continue with redevelopment of the full CG->AA->QM/MM workflow including benchmarking of QM/MM calculations with ChemShell. Support will be given to the BioSimSpace flagship software development project including development of a BioSimSpace metadynamics node for binding kinetics calculations based on GROMACS/Plumed. A new joint project "Structure and Dynamics" with CCP-EM/CCP4 is planned which will use MD simulation to refine low resolution structural information from experimental data focusing on glycosylated proteins. The 4th multiscale modelling conference will be planned to be held in 2020 in Manchester and support given to the CCPBioSim training workshop programme.

Summary Report (1 April 2018 – 31 March 2019)

CRYSTAL: Leonardo Bernasconi left STFC in November 2018, to be partially replaced by a new hire in May 2019, and discussions with the MCC Chair about how best to make use of the remaining resource are ongoing. The eCSE project on hybrid OpenMP/MPI parallelism is complete, allowing CRYSTAL calculations to run more efficiently on HPC platforms. Work is also in progress on the application of CRYSTAL to very large systems, in collaboration with Dr Ian Bush (RSE Fellow at Oxford). The modification in the integral calculation routines for the evaluation of exchange response integrals is also proceeding steadily. In particular, the origin in the time-dependent Hartree-Fock excitation energy discontinuities in pseudo-periodic systems as a function of cell size has been clarified, and the code now gives correct results and no discontinuous behaviour. In addition, we have tested some of the new capabilities of CRYSTAL17, including the more robust SCF solvers and the inclusion of van der Waals forces in simulations of molecular crystals, as from the plans. We published 5 papers during the reporting period, and a further paper is currently under consideration. We have been involved in teaching at the CCP5 Summer School on Molecular Simulation (Lancaster, July 2018) and at the MSSC 2018 CRYSTAL Summer School (Imperial College London, 2018).

ChemShell: David Gunn has developed a Py-ChemShell interface to the ORCA QM code for embedded cluster calculations, including support for effective core potentials (ECPs) and both implicit and explicit QM/MM interfaces. ORCA is a flexible, efficient and easy-to-use general purpose tool for quantum chemistry with an emphasis on spectroscopic properties. The ORCA interface has been validated against the GAMESS-UK and NWChem implementations, and is now available as an additional option for users. Tom Keal continues to supervise the development of a periodic QM/MM scheme in Py-ChemShell targeting metal systems as part of the SAINT project which is progressing to plan, and took part in the MCC software workshop in Lincoln giving an introduction to QM/MM calculations with ChemShell. The first beta version of Py-ChemShell was released in March 2019, and the redeveloped program is now considered ready for production calculations on materials systems and will be installed as a module on ARCHER and THOMAS. This represents a major milestone in the redevelopment of ChemShell and will enable MCC users to make full use of the code in their applications.

DL_POLY/DL_FIELD: The projects' allocated effort during this period has progressed to plan. The Two Temperature Thermostat functionality, a major development effort from 2010 S4F call in collaboration with QMUL (Trachenko) and UCL (Duffy) has been refactored, verified, adapted and pushed for a major release. This will enable sophisticated simulations of radiation damage. The release also contained numerous new features, default changes and bug fixes. This would have not been possible without the help of consortia collaborators and the large refactoring efforts at the much larger team of developers at Daresbury. DL_FIELD progress has also been excellent with numerous new features. The project was showcased at the consortium conference in Lincoln, where it demonstrated its strength and capability, leading to a collaborative proposal by the University of Lincoln that includes both simulation (SCD partnering) and experiment (ISIS partnering). Both DL_POLY and DL_FIELD projects have successfully acquired interest by MCC partners at ISIS and Harwell's Catalysis Hub for organic/inorganic and zeolites modelling (extra funding support has become available from STFC facilities).

Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

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

We will continue to follow the SLA plans outlined in the MCC renewal bid under the constraints of the reduced funding awarded, subject to ongoing review with the MCC committee.

John Purton has joined the ChemShell development team following the departure of David Gunn at the end of 18/19. Our planned targets for Py-ChemShell include the development of an interface to DL_MONTE for Monte Carlo calculations and extension of the LSDalton interface to embedded cluster QM/MM calculations. 0.1 FTE continues to be committed to supervision of the ChemShell work packages in the MCC flagship software development project "SAINT". The remaining effort will be used to maintain Py-ChemShell packages on ARCHER and the Tier 2 platform THOMAS.

For DL_POLY, the development work will concentrate on supervision and testing of the new code implementations to ensure software robustness and interoperability with other projects such as DL_FIELD and ChemShell. The allocations for DL_POLY and DL_FIELD will be temporarily reduced due to successful bidding to MCC related work at ISIS and Harwell Catalysis Hub. Allocations to these codes will increase appropriately at the end of these projects.

Support for the SAINT project will focus on improving the backend software to generate surface models from bulk unit cells and run CRYSTAL simulations. CRYSTAL targets will be finalised following recruitment of a successor to Leonardo Bernasconi, who left STFC in November 2018.

UKCP – UK Car-Parrinello Consortium

Summary Report (1 April 2018 – 31 March 2019)

The annual CASTEP core developer workshop "Codefest", was once again a success. This year Dominik Jochym provided support for a migration to the Bitbucket development platform. As this task was with most of the CASTEP development team present, it proved to be a smooth transition. Dominik has also produced and maintains a workflow document for CASTEP development that includes code contribution guidelines. The work on Bitbucket and the workflow document has been shared with a number of CoSeC staff and Dominik has provided discussion and advice. This provides a base for CASTEP code development as the CCPForge service comes to an end and continued use and adoption of current software development best practice.

The joint UKCP/CCP-NC CASTEP user workshop was again a success with approximately 60 participants with some international attendees. Alongside Albert Bartok-Partay (UKCP/CCP-NC) and Simone Sturniolo (CCP-NC) we provided lectures, tutorial support and direct discussion with delegates.

Use of the Jenkins Continuous Integration platform on the SESC Build Service (to be migrated to the Anvil front end) continues to be of great benefit to the quality of CASTEP development. The ability to test CASTEP throughout development on a wide variety of operating systems and build environments is invaluable. Dominik Jochym has worked with Eli Chadwick (STFC-RAL) on the extension of the Anvil service to deploy tests on STFC's SCARF cluster. A working prototype for the SCARF interface has been produced and tested against CASTEP. This will form the starting point required for a planned interface to Archer and Tier 2 machines.

The STFC deployment of the e-Lucid license management platform has been delayed through a change in personnel. Dominik Jochym is in contact with those who will continue this project and it is planned that CASTEP will be a pilot use for this service. Progress has been made on new academic CASTEP licencing terms, including a worldwide free-of-charge license for academic use. Dominik Jochym has provided support, input and advice to Cambridge Enterprise (CASTEP IP owners) and the CASTEP Development Group in the draft of the new license terms. STFC is placed (subject to finalised contract) to continue the administration of CASTEP licenses and source code distribution on behalf of Cambridge Enterprise.

Code development projects around CASTEP's perturbation theory approach to Raman, Non Linear Optical coefficients, Genetic Algorithm tool and van der Waals DFT are in progress with agreed delivery of these tasks moved to the 2019/20 period.

The activity in support of UKCP is synergistic with the STFC funded PACE project (Proper Analysis of Coherent Excitations). As part of the PACE project Rebecca Fair (supervised by Dominik Jochym) has produced an initial Python version of the existing CASTEP Perl tools for electronic/vibrational dispersion and density of states. It is planned to release this with academic CASTEP in Q1 2019, alongside the Perl tools. Amongst PACE and other projects with the ISIS facility, a further 3 people will be joining Dominik Jochym's team at RAL. These roles will focus on software development projects to bring together simulation (from a number of atomistic modelling codes) and analysis of neutron experiments across two groups in ISIS.

Albert Bartok-Partay has developed machinery for CASTEP to compute "meta-GGA" approximations. This initial implementation of this work is complete and is the first of its kind in providing a consistent treatment of pseudopotential construction. This improved level of approximation provides a more accurate treatment of the physics and chemistry in DFT simulations and has potential impact on all areas of the field.

Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

(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. New code development will continue to extend the capabilities of the exchange correlation functionals done in 2018. Work will continue on the CASTEP licensing platform, code development environment and continuous integration. New functionality will be investigated around the calculation of phonon lifetimes (Boltzmann transport) and an interface with the YAMBO code (electronic excitations). Release of the Genetic Algorithm tool will be carried out.

UKCOMES - UK Consortium on Mesoscale Engineering Sciences

Summary Report (1 April 2018 – 31 March 2019)

The work has been proceeding mostly as planned with support from Michael Seaton and Jianping Meng.

Developments on the LBE code (MPLB) based on the Oxford Parallel library for Structured-mesh solvers (OPS) have been made by Jianping Meng, including the implementation of a first-order body force term for shallow water equations and the three-dimensional MPLB code (both completed). Code can be automatically generated and used for a wide range of hardware, including Intel Xeon Phi and GPUs via the support of various threading models. MPI can also be used to connect multiple accelerators, thus enabling large-scale numerical simulations for the UKCOMES community. This code is in use as a prototype for the EPSRC-funded High-Level Mesoscale Modelling System (HiLEMMS) project, involving the development of a high-level coding abstraction for LBE simulations for automatic code generation.

An improved form of cascaded lattice Boltzmann collisions was implemented in DL_MESO by Michael Seaton. This makes use of matrix multiplications to transform distribution functions into central moments and back, improving the reliability of the previous method (which directly calculated distribution functions from collided central moments) and also allowing an extension to an additional three-dimensional lattice scheme (D3Q19). Improvements have also been made to the various boundary conditions available, extending them to cover concave edges and corners at the outside of a simulation box: these were implemented in DL_MESO in preparation for release in Q4 2018. Additional changes to the LBE code in DL_MESO for its release included new options to reduce the number of output files produced by a

simulation and to write checkpoint files, making it easier for users to analyse its results and to restart previous simulations.

Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

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

Newly developed viscoelastic models for lattice Boltzmann simulations will be added to DL_MESO's LBE code to compliment previously developed non-Newtonian rheological models and expand the types of fluid that can be modelled with this code. The same code will also be prepared for re-engineering along the lines of the High-Level Mesoscale Modelling System (HiLeMMS) and previous work to optimise DL_MESO for Intel Xeon Phi accelerators. The OPS-based code MPLB will be extended by adding turbulence models for higher speed problems and implement a model for thermal problems.

HEC Plasma Physics

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

HECBioSim

Summary Report (1 April 2018 – 31 March 2019)

Longbow is developed under CoSeC in collaboration with Nottingham University. Longbow functions to seamlessly offload simulations to powerful super computers in a manner that appears to the user as the simulation had run on their own PC. Longbow handles all file transfer, scheduling, monitoring and is capable of doing complex multi-machine submission patterns. Currently Longbow has been downloaded 6,211 times split between our website and the command-line installer pip. Two new versions of Longbow have been released in this working period so far, v1.5.1 and v1.5.2.

Version 1.5.1, released in April, fixed a large number of user reported bugs and the addition of parameters to support GPU features on newer generations of HPC facilities. The API was simplified based on feedback from developers of other codes that had incorporated Longbow. The new python chemshell integration was released along with a reorganised and refreshed set of quick start examples.

Version 1.5.2, released in September, saw fixes for four user reported bugs and documentation corrections. The delivery of the update mode is a major step forward for usability, it will enable users to run "fire and disconnect" type submission and instead of being persistently connected for monitoring, users can now do one-time updates to sync their files and progress. A large number of enhancements have also been made including, removal of support for deprecated python versions, customised replica job naming schemes, use of custom scripts, new documentation for examples, refactoring of error handling, and the inclusion of documentation in the version control and CI process.

HECBioSim is partnering up with EPSRC tier 2 facility JADE to offer allocations of compute time to consortium applicant's in-line with what is already offered for ARCHER. During the pre-production state for JADE compilation and benchmarks have been run of a number of GPU ports of popular biosimulation codes. These benchmarks will be used in a similar manner to the ARCHER ones to advise consortium resource panels and their applicants towards sensible resourcing levels for projects. As part of this work, a number of technical issues with libraries and system utilities were investigated and we worked with the system administration team to resolve issues preventing compilation and full utilisation issues. Popular codes have been compiled on JADE and Linux module files created and added to the central module system, these codes will be supported going forwards with updates and problems fixed by HECBioSim.

The JADE benchmarks and ARCHER benchmarks have been collected together with the benchmark suite input files and a dedicated page on the HECBioSim website has now been produced. This page highlights our benchmarking results and also provides download links per each software package so users can benchmark other machines or their local machines to get performance comparisons or to assess project resourcing levels. This can be found here www.hecbiosim.ac.uk/benchmarks, the 2018 versions of codes to date have also been added to those pages for both machines.

Other work on the HECBioSim website so far has included patching the server and content management framework to the latest version. The forums have been removed due to poor utilisation and enabled better GDPR compliance as we no longer have to manage user registered accounts. The project application form, guidance material, documentation, successful projects and FAQs pages for the JADE project applications have been created and are now live and accepting applications.

Work on a project to create a well supported Chemshell-Longbow bridge has been completed. For this project a number of methodologies were tested, however it was decided that modifications to the Longbow core code was the best way to achieve this for a number of technical reasons that became apparent as the project developed. The outcome of this project is that Chemshell users can now use Longbow to launch large arrays of Chemshell jobs on any HPC resource they have access to. This is an interesting development since Chemshell cannot be run interactively on some HPC facilities, so this allows them to do setup on their local PC but then seamlessly hand over the heavy compute to a big machine. This should make the barrier to utilising Chemshell on HPC lower, with the impending inclusion of biosimulation methods into Chemshell in the next release this will be of great benefit to both the biosimulation and materials chemistry communities alike. This feature will be included with the v1.5.3 release of Longbow in the next work year.

Work to assist the Charlie Laughton group at Nottingham University in deploying their internet of simulations platform "TioS" central hub to Microsoft Azure cloud was completed. The work involved a colocation visit to the Laughton lab where it was quickly realised that the TioS software originally developed by a small private company was a Microsoft Windows only utility and had used outdated Azure protocols that no longer worked. This was adapted to be deployed with the newer Microsoft Azure CLI tools which is newer and cross platform. The service was deployed to the Azure cloud, and a paper is in preparation as to how this works.

Work has started on a new cloud based training platform to deliver an elastic scaling Jupyter enabled training platform utilising the Oracle cloud. Due to the successful format of the self-contained training courses based around content in a version control repository. It became apparent that it was a simple technological leap to deliver a fully scalable and interactive cloud based platform that can server interactive Jupyter notebooks to any computer with a web browser (without local IT getting involved). Work on this project to date has included a refactor of the container infrastructure originally developed at Bristol University. The original containers and infrastructure relied heavily on all software being installed into a single large container, which caused long pull times and delays in provisioning resources for workshop delegates. It also made it prohibitively expensive on time to change the training programme. A new micro-services approach has now been applied where a HECBioSim base container has been developed which includes only a basic Linux system along with common tools and utilities and the JupyterHub server. Then each training course then has its own container which builds ontop of this base container, thus cutting down the container size from 50GB to around 1.5GB and fixing the issues of long delays. This is much more simple since adding new workshops, simply builds on the core container meaning it no longer is expensive on developers time solving complex software dependency conflicts etc. This platform will be available in the next work year since we are currently having procurement issues of the cloud resource.

Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

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

Plans for the 2019/20 period currently include ongoing maintenance of the HECBioSim web infrastructure and ongoing benchmarking of codes on ARCHER and tier 2 regional resources. Projects with Richard Henchman in Manchester University are currently in discussions with one currently looking interfacing his entropy based analysis package with Longbow to increase its usability with other MD software packages and another to look at interfacing the entropy tool with biosimspace. A project with Charlie laughton at Nottingham university to produce a Jupyter based training package for Longbow for users and developers is currently in the planning phase. Some project scoping with Syma Khalid at Southampton will also take place to identify suitable projects in the new year.

UK-AMOR

Summary Report (1 April 2018 – 31 March 2019)

The main CoSeC support given to UK-AMOR in this period was Martin Plummer's support from CCPQ SLA/CoSeC funding, as described in the CCPQ report. This included management meetings and discussions and work on the RMAT_REACT interface to the PFARM code. Direct support for UK-AMOR at an average of 0.2FTE per year began in Q4 2018-Q1 2019 as Andrew Sunderland prepared for detailed RMAT_REACT work as described in the plans. As agreed with CCPQ and UK-AMOR steering members, the completion of his UKRMol(+) code memory optimization work for CCPQ was prioritized (this work will have direct impact on UK-AMOR UKRMol(+) work) and 0.1FTE has been reallocated to the year 2018-2019. UK-AMOR itself has made a flourishing start with a successful kick-off workshop and full use of its allocated ARCHER resources.

Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

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

The plans for direct UK-AMOR CoSeC support for 2019-2020 will follow on directly from the 2018-2019 work, continuing the outlined development of the bound state (and resonance) characterization package (the completed package would then be used for ultracold molecules work and fed back into the electron-molecule package UKRMol(+)). This is a result of the limited time (0.3FTE) allocated for direct UK-AMOR support. Martin Plummer will continue to support UK-AMOR separately as part of CCPQ support plans.

UKTC - The UK Turbulence Consortium

Summary Report (1 April 2018 – 31 March 2019)

Work on this project began in June 2018 but initial work was assigned to the grant funding and not to CoSeC. Effort assigned to CoSeC will be back-loaded to the end of the financial year and began in January 2019.

In this short three month period Xiaojun Gu installed the latest version (5.3) of Code_Saturne on ARCHER and retrieved user subroutines of early studies of Taylor-Green Vortex. He also conformed the user subroutines to the current version of Code_Saturne. Over the next 12-18 months we aim to be able to benchmark different UKTC codes, which is something that was flagged up in last year's Review of th CoSeC programme. We are currently trying to build up a benchmark test case that other codes can apply and this will allow code developers to identify bugs and correct. Through our work they will be able to assess and improve their codes.

Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

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

Investigate the feasibility of implementing Lagrangian transport modules in Code_Saturne into another UKTC code.

UKCTRF - UK Consortium on Turbulent Reacting Flows

Summary Report (1 April 2018 – 31 March 2019)

Work on this project began in January 2019 and focussed on work around the HAMISH code, which is a CFD solver for turbulent reacting flows using the adaptive mesh refinement (AMR) technology. In the reporting period Jian Fang delivered a stable HAMISH code to researchers in Cambridge University and Newcastle University, and finished converting the HAMISH code from f77 to f90 format. He also began development of bit operation-based Morton code algorithms for the code and started to create a new version, written with modern Fortran language and dynamic memory allocation.

Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

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

Integrating development from the Software for the Future project into HAMISH and comparing with SENGA+.

Software Outlook

Summary Report (1 April 2018 – 31 March 2019)

During this reporting period, Software Outlook has been working on the hybrid programming, maths/physics libraries and software audit work packages. We also completed the outstanding tasks from the mixed-precision and code coupling work packages. Software Outlook now has its own dedicated website <https://www.softwareoutlook.ac.uk>.

Hybrid Programming: A number of frameworks for porting codes to GPUs being identified and their attributes collated. Some CCPs already have experience with some of these frameworks and information was gathered from them to find out the pros and cons based on their experience. An initial benchmark was identified for use to compare the porting process and the code performance prior to comparison with more complex codes. Benchmarking runs have been performed for the initial benchmark and are in the process of being turned into a technical report. Some of the frameworks were difficult to configure to work with the GPUs: lessons learnt from this will be turned into training packages. Unfortunately, staff illness means that we haven't started work on the more complex codes but this project is scheduled to continue into 19/20. We spoke about this work at the CCP-PETMR Software Meeting in Leeds on 20 September.

Maths/Physics Libraries: Software Outlook identified an extensive list of Fast Fourier Transform (FFT) libraries and their attributes, which is available at <https://www.softwareoutlook.ac.uk/?q=fft>. Working with CCP-PETMR, we identified a benchmark for performing useful comparisons. CCP-PETMR wish to change FFT library due to licensing requirements but we are ensuring that the libraries we compare are wider-ranging than this requirement to ensure that the results are applicable to a wide range of CCPs and HEC consortia. As well as libraries aimed at C/C++, we are also benchmarking those aimed at Fortran codes. Due to staff illness, part of this project had to be performed by another member of the Software Outlook Team to try to get it back on track. The benchmark runs for both the C/C++ and Fortran runs have now been performed and need compiling into technical reports. We spoke about this work at the CCP-PETMR Software Meeting in Leeds on 20 September.

Maths/Physics Libraries: Due to the requirement to perform additional work with regards the Software Audit, the start of the linear solvers work was delayed. Additionally, the staff illness and future staffing plans, meant that the FFT work was prioritised ahead of the linear solvers work. General information and guidance on classes on linear solvers was added to the Software Outlook website as well as links to the Netlib webpage that provides an exhaustive list of libraries and their properties. Work identifying the linear solver needs for a 1D BOUT++ example started but had to be put on hold. This project is scheduled to continue in 19/20.

Software Audit of CCP/HEC Codes: The software audit targeted the flagship CCP/HEC codes with the aim to establish whether they are “user-friendly” and “fit-for-future”. The Software Audit was sent out to all CCP and HEC Consortia chairs in mid-December. The submission deadline was extended following requests from the community. As a result, a full analysis had not been completed by the end of March but the work package suggestions were used to inform our revised 2018-19 Work Plan. The audit report, which covered 42 codes, was presented at the CoSeC Steering Committee in July, and we were asked to produce a second version of the report that anonymised the codes and research communities. The committee gave approval for the delay of the linear solvers work. This new version of the report is being used by the EPSRC as part of its development of a new e-Infrastructure Road Map.

Outline Plans 2019/20 (1 April 2019 – 31 March 2020)

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

Hybrid Programming with GPUs and CPUs: Hybrid GPU/CPU possibilities are now widely available and frameworks/libraries for using them have rapidly evolved. Many CCPs and HECs are now wanting to exploit these next generation architectures and wish to appropriately port their simulation codes. In the 2018/19, Software Outlook started to investigate different frameworks and will have completed an initial investigation into the suitability of different frameworks using a simplified framework (provided by CCP-PETMR), which will allow us to narrow down which frameworks are of interest. We are not expecting the comparative study using example CCP/HEC codes to have been completed by the end of March 2019 and, hence, to truly cement this work, it will need further work in 2019/20. Effort: 0.5FTE.

Physics and Mathematics Libraries: Physics and mathematics problems lie at the heart of most simulation codes and, hence, there are normally parts of the code where external physics and maths libraries can be used. These parts of the code can quickly become a bottleneck if the wrong library or method is used. For example, in BOUT++ (CCP-Plasma), the code is using time-steps that are much smaller than they need to be because the convergence of the underlying iterative linear algebra method is so poor. By improving the convergence through more effective preconditioners, this will also enable the use of a larger time-step with the potential of dramatically improving the execution times and energy consumption values for these computational simulations.

As part of Software Outlook’s 2018/19 work, Fast Fourier Transform libraries and Preconditioning libraries for use in 1D examples from BOUT++ (SD1D), will have been benchmarked. In 2019/20, Software Outlook aim to take this work to the next level by performing further investigations into 3D examples that span the CCPs/HECs. Effort: 0.5FTE

Code Coupling at Scale: Since the initial code coupling work was performed by Software Outlook, code coupling tools have evolved along with their documentation. Historically, considerable effort has been given to in-house development of coupling libraries etc. This work package will aim to assess/demonstrate the practical applicability of existing code coupling tools (PLE, MUI, etc.) in coupling two independent codes. It also allow us to provide recommendations for, and improvements to, available code coupling tools to maximise their value across a range of HPC applications. Effort: 0.3FTE (+ effort in 20/21)

General Consultancy Work for CCPs/HECs: As Software Outlook’s visibility increases, it is being asked to perform small pieces of consultancy work for the CCPs/HECs. As the need arises, we will help provide this consultancy work for the CCPs. Effort: 0.2FTE

Resource Planning and Yearly Costing

Level of funding and resource for the current SLA

During the current SLA cycle, the total funding agreed corresponds to 113 FTE over the 5 years. This figure includes the effect of the £200k drop in funding from 2018/19 until the end of the SLA imposed by EPSRC’s own budget constraints and made known to STFC at the time of the SLA renewal in 2016. In addition to delivering the full 113 FTEs of effort for the current five year SLA cycle, we have agreed that during this cycle we will make up for a shortfall in effort delivered that arose for several projects during the previous SLA cycle (up to 2015/16), as indicated in the table below. The table also reflects the outcome of the EPSRC call for HEC consortia in 2017, ie an increase of the number of consortia supported by the SLA – from 5 to 8. A call for CCPs will take place in 2019, and in the following we have not included the likely but at this point unknown changes in the SLA effort distribution that will ensue.

In addition to delivering the full 113 FTEs of effort for the current five year SLA cycle, we have agreed that during this cycle we will make up for a shortfall in effort delivered that arose for several projects during the previous SLA cycle (up to 2015/16), as indicated in the table below. This historical shortfall was due to staff sickness and difficulties in recruiting suitable staff.

Planned Resource Profiles

In this current planning process, for each project we allow the amount of effort to vary from year to year and to be re-adjusted yearly as long as each project recovers its full support over the five years of the SLA cycle.

The following table shows the predicted spend profiling for the next five years. For 2016/17, 2017/18 and 2018/19 two figures are included – in **black**, the planned effort used to calculate the annual cost, and in **red** actual effort provided. For the HEC consortia yearly allocation, the figure in brackets shows the allocation before the renewal in 2017. The figures in **pink** are the proposed figures and totals for the remainder of the the current CoSeC SLA cycle.

| | Agreed Carry Over | Yearly Alloc | 2016-17 | 2017-18 | 2018-19 | 2019-20 | 2020-21 | Total |
|----------------|-------------------|--------------|--------------|--------------|--------------|---------|---------|---------------|
| Project Office | 0.00 | 2.25 | 2.23 2.25 | 2.23 2.25 | 1.79 1.45 | 1.50 | 0.89 | 11.25 8.64 |
| CCP5 | 0.42 | 3.20 | 1.95 2.82 | 3.19 3.56 | 3.68 3.68 | 4.38 | 3.22 | 16.42 |
| CCP9 | 0.83 | 2.40 | 2.64 2.71 | 2.28 2.70 | 2.70 2.90 | 2.70 | 2.51 | 12.83 |
| CCPmag | 0.47 | 0.74 | 0.64 0.83 | 0.70 0.40 | 1.24 1.23 | 0.72 | 0.87 | 4.17 |
| CCPNC | 1.80 | 1.30 | 1.62 1.50 | 2.00 2.00 | 1.79 1.80 | 0.64 | 2.25 | 8.30 |
| CCPQ | 0.00 | 1.86 | 1.69 1.86 | 1.80 1.86 | 1.92 2.00 | 2.00 | 1.89 | 9.30 |

| | | | | | | | | |
|-------------------|-------------|--------------------------------|------------------------------|------------------------------|------------------------------|--------------|--------------|--------------------------------|
| CCP Plasma | 0.48 | 0.75 | 0.79 1.00 | 0.79 0.80 | 0.85 0.80 | 0.80 | 1.00 | 4.23 |
| CCPi | 0.33 | 1.20 | 0.82 1.45 | 1.01 1.50 | 1.13 1.35 | 1.70 | 1.67 | 6.33 |
| CCPPetMR | 0.33 | 1.20 | 1.09 1.15 | 1.53 1.50 | 1.44 1.40 | 1.30 | 0.97 | 6.33 |
| CCPBioSim | 0.00 | 1.20 | 1.19 1.16 | 0.79 1.20 | 1.23 1.20 | 1.40 | 1.39 | 6.00 |
| MCC | 0.00 | (2.50) 2.00 | 2.53 2.50 | 2.32 2.50 | 1.95 2.15 | 1.45 | 2.90 | (12.50) 11.15 |
| UKCP | 0.42 | 1.00 | 1.36 1.42 | 0.97 1.00 | 0.98 1.00 | 0.90 | 1.21 | 5.42 |
| UKCOMES | 0.00 | (1.00) 0.60 | 0.93 1.00 | 1.16 1.00 | 0.75 0.60 | 0.60 | 0.43 | (5.00) 3.87 |
| HEC Plasma | 0.00 | 0.20 | 0.20 0.20 | 0.20 0.20 | 0.21 0.20 | 0.20 | 0.19 | 1.00 |
| HECBioSim | 0.50 | (1.00) 0.80 | 0.98 1.00 | 1.01 1.00 | 1.02 1.00 | 1.00 | 0.99 | (5.50) 5.00 |
| UK Turbulence | 0.00 | (0.00) 0.40 | 0.00 0.00 | 0.00 0.00 | 0.10 0.30 | 0.40 | 0.60 | (0.00) 1.10 |
| UKCTRF | 0.00 | (0.00) 0.50 | 0.00 0.00 | 0.00 0.00 | 0.14 0.13 | 0.50 | 0.49 | (0.00) 1.13 |
| UKAMOR | 0.00 | (0.00) 0.20 | 0.00 0.00 | 0.00 0.00 | 0.09 0.20 | 0.30 | 0.21 | (0.00) 0.60 |
| Software Outlook | 0.10 | 2.00 | 2.31 2.00 | 1.64 2.00 | 1.54 1.50 | 1.50 | 1.61 | 10.10 8.60 |
| FTE TOTALS | 5.68 | (23.80) 23.80 | 22.97 24.85 | 23.62 25.47 | 24.55 24.89 | 23.99 | 25.29 | 124.68 120.42 |

Planned vs Actual Resource

While the resource planning is now done over a 5-year period, the costing of the programme needs to be calculated on a yearly basis based on the planned effort. Deviations between planned and actual effort are calculated and used to reprofile the effort and costing for future years. For 2016/17, 2017/18 and 2018/19, the difference between delivered and planned effort for each project has been calculated and redistributed over future years so that by the end of the SLA cycle each project will spend the amount of FTEs awarded for this SLA plus any historical underspend that a project may have incurred.

2019/20 Planned Resource and Costing

Of the planned 23.99 FTEs for 2019/20, 20.33 FTEs have been costed via the SLA, while 3.66 FTEs will be provided based on funding received previously (i.e. carry-over from previous SLA or adjustment due to deviation between planned and actual effort). The detailed cost breakdown for 2018/19 for the 20.33 FTEs is as follows:

| | | | | | | |
|--|--|--|--|--|--|--|
| | | | | | | |
| | | | | | | |
| | | | | | | |

| | | | | | | |
|--|--|--|--|--|--|--|
| | | | | | | |
| | | | | | | |
| | | | | | | |

Please note that recurrent expenditure includes:

- Travel and Subsistence – this includes T&S for our staff attending management/ project meetings. Includes T&S for visitors.
- Consumables – includes items such as telephone charges, training, photo repro, registration fees and workshop costs.
- HW/SW – this includes capital purchase of desktop/ laptop systems, software maintenance and applications packages.
- Participation in Supercomputing, production of annual report and general support activities around the SLA including the web.
- Hardware and software maintenance costs.
- Computing Infrastructure covers maintenance of servers, software licenses and printing consumables. Staff effort supports hardware and software maintenance and systems management of desktop systems and file servers and support of visualisation and grid infrastructure.

CoSeC Starters and Leavers

During this reporting period the following staff left the CoSeC programme:

- Marion O’Sullivan, Impact Manager – replaced by Dawn Geatches
- Martin Luders – relocated to University of Hamburg, Germany
- Steven Lamerton – relocated to industry, The Netherlands
- Thomas Kulhanek – relocated to industry, Czech Republic
- Leonardo Bernasconi – relocated to Pittsburgh Supercomputing Centre, USA

New members of staff joining the CoSeC programme included:

- Dawn Geatches, Impact Manager – replacing Marion O’Sullivan
- Andrey Brukhno – CCP5
- Gemma Fardell – CCPi
- Kakali Sen – CCPBioSim
- Sebastian Metz – CCPBioSim
- Xiaojun Gu – UKTC
- Jian Fang – UKCTRF
- Philippe Gambron – Software Outlook

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 | 600 | 3.20 |
| CCP9 (Questaal) | 1981 | 450 | 2.40 |
| CCPPlasma (GS2 / BOUT++) | 2007 | 150 (with HECPlasma) | 0.75 |
| CCP-NC (MagresView) | 2011 | 60 | 1.30 |

| | | | |
|--------------------------------|------|-----------------------|------|
| CCPQ (R-Matrix, TNT, Quantics) | 2011 | 150 | 1.86 |
| CCPBioSim (FESetup) | 2011 | 345 (with HECBioSim) | 1.20 |
| CCPi (CCPi CIL) | 2012 | 380 | 1.20 |
| CCP-Mag (KKR) | 2015 | 44 | 0.74 |
| CCP PETMR (SIRF) | 2015 | 80 | 1.20 |
| UKCP (CASTEP) | 1990 | 150 | 1.00 |
| MCC (CRYSTAL, Chemshell) | 1994 | 464 | 2.00 |
| HECBioSim (Longbow) | 2013 | 345 (with CCP BioSim) | 0.80 |
| UKCOMES (DL_MESO) | 2013 | 150 | 0.60 |
| HEC-Plasma (GS2, BOUT++) | 2013 | 150 (with CCP Plasma) | 0.20 |
| 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.

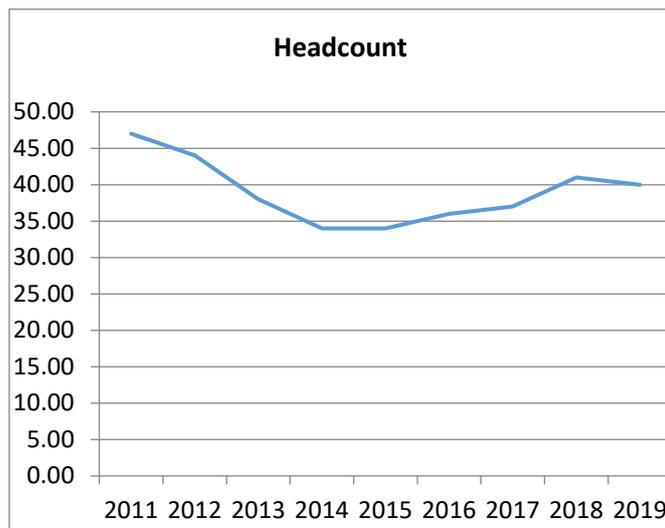
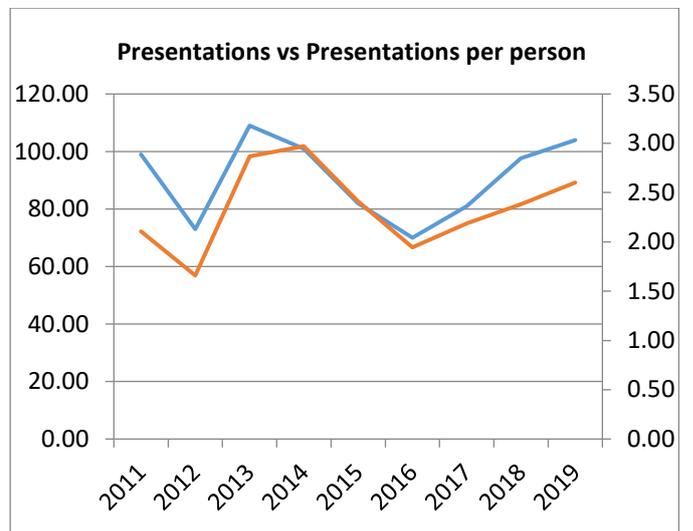
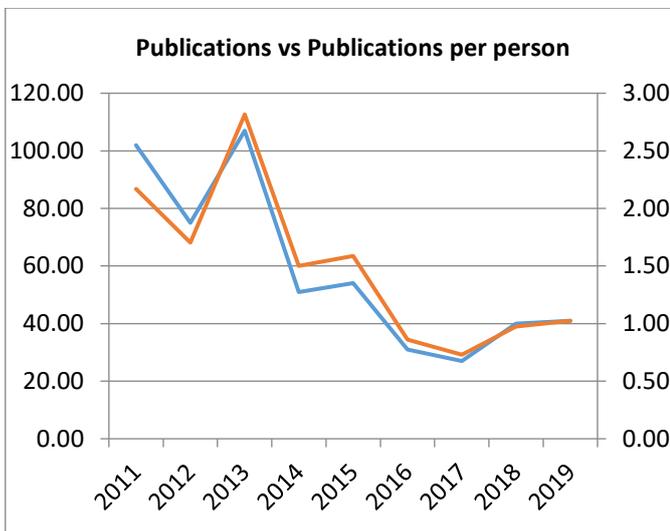
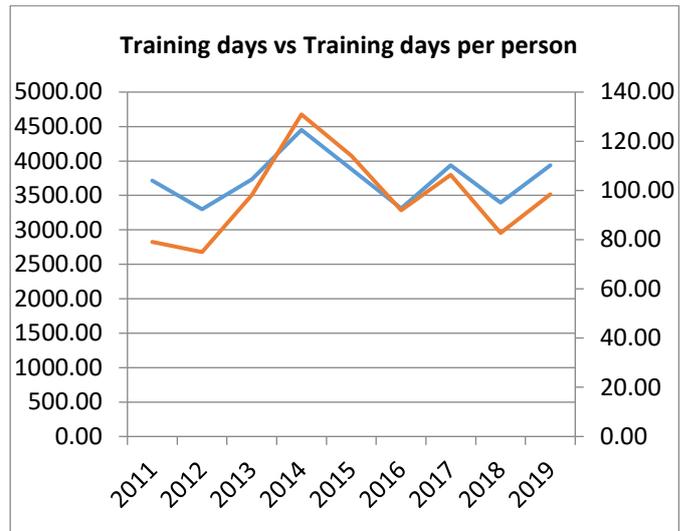
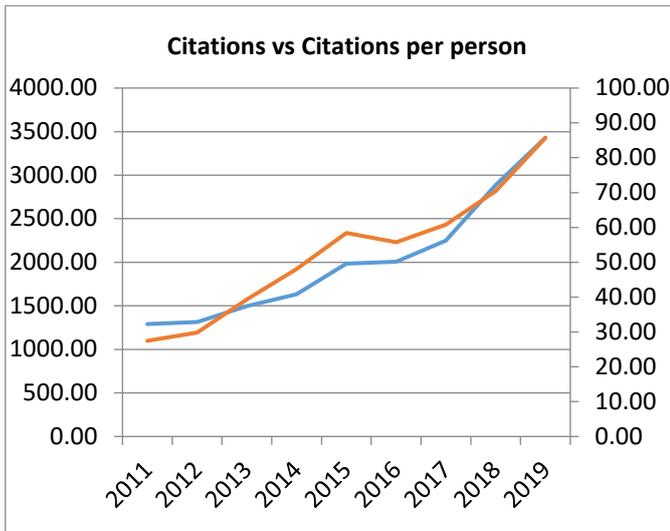
In the metrics graphs below, the blue line refers to the total number while the red line refers to the number per scientific/technical staff member. This is useful because metrics are included for any person irrespective of whether they are partially or fully funded by the SLA. For information, a graph of the scientific/technical staff headcount has also been included.

The metrics show that the programme is in very good health. The number of code citations is in steep growth, edging above 85 citations per person (not per FTE). The trend for training days delivered is also very positive, with nearly 4,000 training days delivered in the reporting year.

After a steady fall in the number of publications from 2013 where the average was around 2.5 publications per person, the reasons for which were explained in previous reports and Steering Committee meetings, we have started to pick up again since 2017 and we now average around 1 publication per scientific/technical staff member.

The 2018 SLA review panel report stresses strongly the importance of SLA staff's involvement in research activities, and it would be therefore important to increase the number of publications. The current mismatch between the demand from an increasingly large user base in terms of software maintenance, user support and training (as evidence by the metrics above) and the amount of SLA resource is squeezing out the time for research activities. This is a cause for concern that the SLA Steering Committee and EPSRC should consider

Likewise, after a fall in the number of presentations between 2013 and 2016 the CoSeC programme now averages around 2.5 presentations per person, with a total of 104 presentations throughout the programme in financial year 2018/19.



Metrics breakdown by area

| | FTE | Publications | Presentations | Training Days | Citations |
|------------------------------------|--------------|--------------|---------------|----------------|----------------|
| CCP5 (including DL_ codes) | 3.20 | 13.00 | 43.16 | 2401.50 | 608.00 |
| CCP9 | 2.40 | 1.00 | 4.00 | 240.00 | 250.00 |
| CCPmag | 0.74 | 0.00 | 2.00 | 0.00 | 7.00 |
| CCP-NC | 1.30 | 6.50 | 2.50 | 150.00 | 480.00 |
| CCPQ | 1.86 | 0.00 | 3.50 | 55.00 | 85.00 |
| CCP Plasma/HEC Plasma | 0.95 | 0.00 | 3.00 | 6.00 | 83.00 |
| CCPi | 1.20 | 4.00 | 2.00 | 69.00 | 0.00 |
| CCP PET-MR | 1.20 | 1.00 | 2.00 | 40.00 | 0.00 |
| CCP BioSim/HEC BioSim | 2.00 | 1.00 | 3.68 | 195.00 | 11.00 |
| MCC (including ChemShell, CRYSTAL) | 2.00 | 6.50 | 22.66 | 533.50 | 562.00 |
| UKCP (including CASTEP) | 1.00 | 0.50 | 3.50 | 150.00 | 1317.00 |
| UKCOMES | 0.60 | 0.50 | 4.50 | 98.00 | 20.00 |
| UKTC | 0.40 | 2.50 | 0.00 | 0.00 | 0.00 |
| UKCTRF | 0.50 | 2.50 | 5.00 | 0.00 | 3.00 |
| UKAMOR | 0.20 | 0.00 | 0.50 | 0.00 | 0.00 |
| Software Outlook | 2.00 | 2.00 | 2.00 | 0.00 | 3.00 |
| Totals | 21.55 | 41.00 | 104.00 | 3938.00 | 3429.00 |

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