




UK Atomic
Energy
Authority

ExCALIBUR


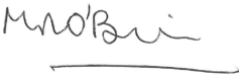
Fusion Modelling System Science Plan

Abstract

This document outlines the Science Plan for research to be commissioned by the Met Office for the "Fusion Modelling System" use case of the SPF "Exascale Computing Algorithms and Infrastructures Benefiting UK Research" programme (ExCALIBUR). Amendments to this document will be tracked in the change log on p.1 and each change will require review and approval by the UKAEA internal ExCALIBUR + E-TASC Project Board Director as well as other associated interfaces as deemed appropriate by the Board.



UKAEA REFERENCE AND APPROVAL SHEET

	Client Reference:		
	UKAEA Reference:	CD/EXCALIBUR-FMS/0001	
	Issue:	1.00	
	Date:	25 October 2019	
Project Name: ExCALIBUR Fusion Modelling System			
	Name and Department	Signature	Date
Prepared By:	Wayne Arter	N/A	21/10/19
	Lucian Anton	N/A	21/10/19
	Debasmita Samaddar	N/A	21/10/19
	Rob Akers	N/A	22/10/19
	MSSC		
Reviewed By:	Rob Akers		22/10/19
	MSSC		
Approved By:	Martin O'Brien		25/10/19
	MSSC		

ExCALIBUR

Fusion Modelling System Science Plan

Purpose

This document outlines the Science Plan for research to be commissioned by the Met Office for the “Fusion Modelling System” use case of the SPF “**Ex**ascale **C**omputing **A**lgorithms and **I**nfrastructures **B**enefiting **U**K **R**esearch” programme (ExCALIBUR). It complements ExCALIBUR Met Office Science Plan and Activities document [1], setting out a vision for the project’s contribution to the overarching ExCALIBUR plan. This document is aimed at stakeholders including the SPF ExCALIBUR Programme Board and Steering Committee, has been informed by consultation with the Met Office, with domain experts from across the UKRI community and from within UKAEA. This Science Plan is intended to complement the ExCALIBUR activities that are funded through the UKRI Research Councils and the Met Office “Weather & Climate Prediction” use case.

The ExCALIBUR programmatic aims described in [1] are common with those of the Fusion Modelling System use case and so are not repeated here.

Fusion Modelling Use Case activities

The aim of the “Fusion Modelling” use case of ExCALIBUR is, via the exploitation of the ExCALIBUR principles outlined in [1], to develop new algorithms, software and related e-Infrastructure that will result in the efficient use of current Petascale and future Exascale supercomputing hardware in order to draw insights from ITER [2] “Big Data” and to guide the design of a UK demonstration nuclear fusion power plant [3] and related fusion technology as we approach the Exascale. The aims of the work are to deliver expertise in, and tools for, reactor interpretation and design, initially with a focus upon the “edge” region of the tokamak [4] plasma where hot plasma comes into contact with the material walls of the machine (see project NEPTUNE below). This challenging, multi-physics, multi-scale intersection between plasma physics and engineering has been identified as critical to the success of commercial fusion energy. The programme is designed to exploit commonality of solutions across the disciplines and domains represented by ExCALIBUR and to foster the development of a UK interdisciplinary community (within our National laboratories and institutes across UKRI and throughout Academia).

Activities will initially be around an ambitious programme to develop a computational model that includes plasma kinetic effects believed to be essential for a first-principles description of the complex dynamics of high temperature fusion plasma. Codenamed NEPTUNE (**NE**utrals & **P**lasma **TU**rbulence **N**umerics for the **Ex**ascale), work will initially focus upon coupling the turbulent plasma periphery to the surrounding neutral gas and partially ionized

impurities that exist between the plasma and plasma facing components, in the presence of an arbitrary tokamak magnetic field and full 3D first wall geometry. Figure 1 shows a schematic of a generic tokamak “poloidal cross section”, highlighting the targeted regions of plasma and machine, namely the main chamber between core plasma, scrape off layer and wall (upper shaded circle) and the so called “plasma exhaust” or “divertor” region (lower shaded circle) where heat and particles come into direct contact with material surfaces. The resulting close coupled model will be constructed using an accurate, high fidelity representation of the geometry described by Computer Aided Design (CAD) systems.




Figure 1: Schematic diagram of a generic tokamak “poloidal cross section” showing the areas of plasma and first wall that will be targeted by project NEPTUNE (shaded circles). Attribution: G. Federici et al. [CC BY 3.0 (<https://creativecommons.org/licenses/by/3.0/>)]†.

Quality control, verification, validation and uncertainty quantification (VVUQ, e.g. via intrusive or ensemble-based methods) will be embedded across all areas of the project to ensure numerical predictions are “actionable”. The initial aim is to develop knowledge and UK capability around how to design world leading Exascale targeted software for the benefit of the UK academic community, UKRI and the UK nuclear supply chain.

Together, the UK plasma and HPC communities have unique expertise required to deliver this project. UKAEA will bring together world-class experts in tokamak edge physics, gyrokinetic theory, and highly scalable algorithms, to address arguably one of the most

† Minor modifications to figure.

important unsolved challenges of fusion research – how to design a plasma “exhaust system” that can reliably restrict power flux reaching the material surfaces to tolerable levels, i.e. no more than $\sim 10\text{MW/m}^2$ in the steady-state. Later in the project, further packages of work will be designed to address other aspects around the “in-silico” design of fusion technology. It is recognised that the project poses a significant human resource and project management challenge, notably to develop and manage a team across many different sites and from many different backgrounds in order to build state-of-the-art software that can reliably and accurately account for a plethora of different complex physical phenomena. Further, the project must maintain an international context and connect to the European EUROfusion programme (in which the UK is a key player) and the US Exascale programme (ECP). Crucially, the emerging software environment must be trusted and “actionable” to guide procurements potentially involving hundreds of millions of pounds (e.g. in the case of the DEMO [5] first wall), and ultimately to ensure the safe operation of a multi-billion pound nuclear plant.

The UK is world-leading in the study of the edge region of tokamak plasmas. The flagship EPSRC/EUROfusion MAST-Upgrade experiment [6] recently commissioned by UKAEA at Culham has been built with the primary goal of testing a novel “Super-X divertor” design for handling the plasma exhaust. Combined with software developed under ExCALIBUR, the result will be a significant step forward in our understanding of how heat and particle flows can be controlled and kept to within material limits inside a reactor.

Project NEPTUNE

(NEutrals & Plasma TURbulence Numerics for the Exascale) - background

The UK plasma community – in UKAEA and several universities - currently makes extensive use of “fluid” codes in its research into the edge-plasma region of fusion devices. “Fluid” implies that the plasma can be treated in one sense like “the atmosphere” in the Met Office’s Dynamical Core, but with the added complication of significant effects due to the electrically charged nature of the plasma. For example, plasma electrons and ions can to an extent be treated as separate fluids with different temperatures, interacting via the electromagnetic field. Plasma fluid codes (such as BOUT++ [7]) are relatively efficient, but there are a large range of effects in the tokamak plasma edge that prevent the plasma from isotropising to a Maxwellian distribution with same temperature ion and electron populations (which would correspond most closely to the atmospheric fluid concept). Firstly, the strong imposed and directed magnetic field leads to anisotropy since charged particles move rapidly in tight, spiral orbits (gyro-orbits) along the closed field lines and only slowly normal to the field. Secondly there may be an inadequate number of collisions (especially in the burning plasma or “reactor” regime) for the species to equilibrate, leaving significant tails in the particle velocity distributions. Such tails may be driven by external heating and/or by collisions with relatively rarefied neutral particle species that are themselves non-Maxwellian. Inclusion of these effects in the fluid models introduces significant modelling uncertainty. Moreover, existing simulations show that attempting to include the required extra physics can significantly add to execution cost overhead or a breakdown of scaling, e.g. when stability/accuracy issues are addressed, a large reduction in allowable timestep can ensue.

Unfortunately, there are significant challenges with moving to a straightforward particle-based model (e.g. PIC [8], which could in principle address these problems) - the strength of the electric field in the plasma edge exaggerates the effects of noise when sampling charged particles.

The next most widely studied level of treatment of velocity space, namely “gyrokinetics” [9], averages over gyro-orbits so that velocity space is treated as a 2D problem (perpendicular and parallel to the local magnetic field). This leads to 5D gyrokinetic models, which in principle are detailed enough and accurate enough to model reality. Ideas around the discretisation of the velocity (phase) space dependence include techniques suitable for dealing with infinite coordinates, such as moment-based methods using truncated series and/or mapped finite elements perhaps combined with spherical harmonics. Unfortunately, existing 5D gyrokinetic models are currently at or beyond the limit of current HPC capability in terms of scalability.

Kinetic levels of complexity are nonetheless going to be necessary (at least locally) for modelling the burning plasma regime, due to the inherent uncertainty in the fluid codes. The plasma in a fusion reactor may well behave significantly differently to plasma in existing devices because it will in general contain two main ionic species (Deuterium and Tritium), neutral fuel particles and ionised Helium ash (or alpha particles), as well as impurity ions originating from the wall. Further, the plasma will be hotter, reducing collisions so that yet more complicated terms need to be added to the fluid approximations. These additional contributions to the fluid models will require tuning in much the same way as meteorological micro-physical effects, but in advance preferably of suitable experimental data from ITER or other reactors. Kinetic code results will inevitably be needed for this “tuning” exercise (i.e. providing kinetic closures to the fluid codes).

Evidently the interpretation of data from ITER and consequently the design of a DEMO fusion power plant will require a hierarchy of models, from those that can be deployed upon Exascale hardware down to the surrogate models that will be deployed at scale upon the high throughput computing platforms of the ITER era (e.g. for uncertainty quantification and parametric optimisation). There are also likely to be additional demands upon code execution speed from DEMO operation, for example there will inevitably be a safety-related need for surrogate models that can be deployed as part of future real-time systems to help control the temperature of the first wall. These future needs will be embedded within the heart of the project requirements throughout the development lifecycle of the NEPTUNE software stack.

The project itself has been designed in stages, so that early work will make significant contributions to existing software capability and to the development of standalone mini-apps [10] that capture the scaling characteristics of the full models being targeted in the long term. The overall aim, which will guide the direction of the project and choice of sub-tasks (most of which will be defrayed across UKRI and the Universities), is to build a hierarchy of models that are capable of representing edge plasma behavior to within a specific level of uncertainty, with the option of at least partial incorporation into the development of software across the EUROfusion work programme, notably the TSVV (Theory, Simulation, Verification and Validation) activities.

High level aims of the ExCALIBUR Fusion use case, aligned with the Met Office Weather & Climate prediction system use case include:

- To apply the principles of ExCALIBUR to deliver the benefits outlined above.
- To develop and deliver cross-cutting research that aligns with the UKRI Research Council contribution (i.e. to help deliver a UK interdisciplinary team that can address aspects of Exascale software design that lie in common across the represented use cases).
- To help train the software engineers, architecture specialists, computer scientists/algorithms specialists and data scientists of the ITER and Exascale era.

UKAEA will work with the SPF ExCALIBUR Programme Board and Steering Committee to ensure alignment and a close working relationship with the UKRI funded research activities, which for example could include participation in workshops and knowledge exchange activities, participation in stakeholder engagement exercises etc.

As part of this Science Plan, four work packages will initially be commissioned by the Met Office, i.e. FM-WP1/2/3/4. As outlined in the Met Office Science Plan [1], the Met Office will manage the cross-cutting theme work packages XC-WP1/2 (see below). The four initial Fusion Modelling System work packages are:

1. Numerical representation;
2. Plasma multiphysics model;
3. Neutral gas & impurity model;
4. Code structure and coordination,

as shown in table 1 (and in [1]). All work packages will be built around modern best practice in co-design and there will be significant interaction between the work packages. Each work package will be supported by Research Software Engineers (RSEs) (including engagement with the Society of RSE [11] where necessary).

PSRE Use Cases				
Weather & Climate prediction system:	WC-WP1: Component model co-design		WC-WP2: System co-design	WC-WP3: System integration
Fusion Modelling system:	FM-WP1: Numerical representation	FM-WP2: Plasma multiphysics model	FM-WP3: Neutral gas & Impurity model	FM-WP4: Code structure & coordination
Cross-Cutting Themes				
XC-WP1: Common approaches & solutions		XC-WP2: Emerging technologies		

Table 1: SPF ExCALIBUR Met Office commissioned Work Packages.

FM-WP1 will focus upon the suitability of available numerical algorithms (or the development of new algorithms) for Exascale targeted plasma modelling. Elements of this work package together with FM-WP4 will also surround the “coupling technology” that will be required to connect the edge/pedestal region of the plasma (addressed by FM-WP2) and the neutral gas/impurity model (FM-WP3).

FM-WP2 and **FM-WP3** will concentrate upon development of the two close coupled models of the NEPTUNE programme, specifically FM-WP2 around the inclusion of kinetic effects into existing and new edge plasma models, and FM-WP3 of particle based models for describing the region outside and just inside the plasma (neutral atoms/molecules and partially ionized impurities). Initial exploratory work will be carried out using existing codes and via the development of mini-apps, for example to expose options for exascale targeted hardware (GPUs, ARM technology etc.).

FM-WP4 will focus upon the design of data structures to interface between the different models and generally ensure best practice in scientific software engineering (being responsible for Quality Assurance, co-design, integration etc.). Aligned with Met Office work package WC-WP1, this programme of work will also explore the use of a “separation of concerns” methodology for the Fusion use case critical components (e.g. by exploring the use of Kokkos [12] and Domain Specific Language (DSL) [13] technologies).

Scoping work for the Met Office led cross-cutting themes (XC-WP1 and 2) will be defined in collaboration with the other ExCALIBUR partners early in year 2 with a view towards starting work later that year. The specifics of the cross-cutting themes are here assumed to be covered by the two overarching themes of “Common Approaches and Solutions” and “Emerging Technologies” (discussed further below).

UKAEA Research Plans

Numerical representation (Fusion modelling, work package **FM-WP1**)

The ideal numerical algorithms for forming the Exascale edge plasma codes of the future will have (but will not be restricted to) the following properties:

1. Accurate solution of hyperbolic problems.
2. Ability to deliver efficient and accurate solutions of corresponding elliptic problems.
3. Accurate modelling of highly anisotropic dynamics.
4. Accurate representation of first wall geometry (face normals to within 0.1°), and correspondingly of complex magnetic field geometries.
5. Accurate representation of velocity (phase) space.
6. Preservation of conservation properties of the underlying equations.
7. Scalability to likely Exascale architectures:
 - a. interaction between models of different dimensionality,
 - b. interaction between particle and fluid models,
 - c. dynamic construction of surrogates.

8. Performance portability to allow rapid deployment upon emerging hardware.

The choice of geometrical representation and numerical scheme will have an impact upon almost all areas of the project. Options will be identified by means of literature and code surveys and consultation with UKRI and industry experts. Research will be commissioned where necessary to eliminate unsuitable choices as early as possible. Relatively small development tasks will initially be undertaken to test remaining candidate methods for accuracy, stability and HPC scalability potential. This task, along with FM-WP4 will have a prioritised start to provide initial inputs into FM-WP2 and FM-WP3 developments and work to be defrayed in year 2. Tasks FM-WP1-3 will incorporate numerical, finite element and other plasma physics libraries of suitable quality and exascale applicability.

Options, which do not preclude consideration of others, have been tentatively identified for initial investigation as follows:

1. Spectral/hp element [14], combined with Discontinuous Galerkin [15], to meet 1,3 and possibly 5 above.
2. Multigrid methods for 2.
3. Nekmesh for 4.
4. For Exascale (7):
 - a. matrix-based approaches, hierarchical geometric structures,
 - b. kinetic enslavement, multi-index Monte-Carlo methods,
 - c. physics based Neural Network approaches.
5. MUSCLE 2 etc. as referenced in [16], MUI [17], ADIOS [18] etc. for code coupling 6,7.

Plasma multiphysics model (Fusion modelling, work package **FM-WP2**)

This work package will begin by identifying a referent in conjunction with potential users, whereby “referent” is meant a model that establishes the maximum detail and complexity of plasma that the software could ever be reasonably expected to model (beyond exascale). Critical features of the referent will be identified and prioritised for implementation as part of the project. This may involve replacing a kinetic model by moment-based or fluid models. Input from the European Boundary Code (EBC) development (funded by EUROfusion wherein UKAEA is a core partner) will be important to this process. In addition, the speed of model execution will be a consideration as indicated earlier. A provisional sequence of developments is as follows (and will be tuned as part of the initial NEPTUNE requirements capture exercise funded in year 1):

1. 2D model of anisotropic heat transport.
2. 2D elliptic solver in complex geometry.
3. 1D fluid solver with simplified physics but with UQ and realistic boundary conditions.
4. Spatially 1D plasma model incorporating velocity space effects.
5. Spatially 1D multispecies plasma model.
6. Spatially 2D plasma model incorporating velocity space effects.
7. Interaction between models of different dimensionality.
8. Spatially 3D plasma kinetic models.

Users from the wider fusion community will be engaged in the later stages to help define and add new surrogate models through a wide engagement programme and co-design (e.g. to treat boundary sheaths) and/or other important physical effects (e.g. radiation, charge exchange recombination, etc.) and compare with existing codes and experiments.

Neutral gas & Impurity model (Fusion modelling, work package FM-WP3)

This work package will begin by identifying a referent (as above) in conjunction with potential users and experts in atomic physics. Critical features of the referent will be identified and prioritised for implementation as part of the project. This may involve deploying a particle-based method, a moment-based model or a fluid model (or even a combination thereof). Input from the European EBC programme will be important to this process, as will the speed of model execution as indicated earlier. A careful assessment of existing codes currently in use will be necessary, as will cross validation with the established models (notably B2-EIRENE [19]). A provisional sequence of developments is as follows:

1. 2D particle-based model of neutral gas & impurities with critical physics.
2. 2D moment-based model of neutral gas & impurities.
3. Interaction with 2D plasma model when available.
4. 3D model of neutral gas & impurities.
5. Interaction with 3D plasma model.
6. Staged introduction of additional neutral gas/impurity physics.

As per FM-WP2, users from the wider fusion community will become involved in the later stages of the project to add new physics and compare with existing codes and experiments, providing detailed and rigorous verification and validation (V&V). This is likely to require the provision of databases for different ionisation and excitation reactions, both in the plasma volume and at surfaces.

Code structure and coordination (Fusion modelling, work package FM-WP4)

The most important aim of this work package will be to drive user engagement and ensure that the software is fit for its defined purpose, first by requirements capture, then by defining suitably flexible code structures and related e-Infrastructure for users, ultimately supporting uptake of the new code(s). In order to achieve this aim, this work package will coordinate across the other tasks FM-WP1-3, to ensure that outcomes are compatible and of sufficiently high quality. There will be management and coordination tasks that will grow as the project matures, connecting with the EUROfusion E-TASC (TSVV) programme, the EPSRC T.P. Turbulence Programme [20] and the US ECP programme etc.

Coordination tasks will include (but are not restricted to):

1. Allocation of resource between tasks and setting project priorities.
2. Ensuring a consistent choice of definitions (ontology) of objects or equivalently classes.
3. Definition of common interfaces to components for data input and output.
4. Design of suitably flexible data structures for common use by all developers.

5. Establishment, promotion and support of good scientific software engineering practice.
6. Evaluation and deployment of performance portability tools and DSLs targeting Exascale-relevant architectures.
7. Integration of the developed software into a VVUQ framework (exploiting common approaches developed under XC-WP1 and XC-WP2).
8. Coordination of a benchmarking framework for correctness testing and performance evaluation of the developed software stack.

Along with FM-WP1, this work package will be prioritised for an early start, as good scientific software engineering practice needs to be agreed quickly, and well documented interfaces to components need to be available early to ensure that best practice design is embedded from the start.

Common approaches and solutions (Cross-cutting themes, **XC-WP1**)

The examples listed in [1] are very well aligned with the needs of the Fusion use case comprising initially NEPTUNE, especially the methodologies surrounding a “separation of concerns”, coupling technologies, the use of mixed precision arithmetic and fault tolerance. In addition, project NEPTUNE will benefit from an exploration of the convergence of HPC and AI (e.g. the use of Neural Network PDE solvers [21]) and parallel in time methods [22]. These areas of common ground and opportunities for interdisciplinary working therein shall be explored and agreed by partners as outlined in the Met Office Science Case.

Emerging Technologies (Cross-cutting themes, **XC-WP2**)

The examples listed for XC-WP2 in the Met Office Science Case [1] are again highly aligned with the Fusion use case. Other technologies that are deserving of consideration include High Bandwidth Memory (HBM [23]), novel exascale targeted IO/storage technology (e.g. H2020 project SAGE II [24] led by Seagate wherein UKAEA is a partner), next generation accelerator technology (Nvidia Volta Next, A64fx) and ARM systems for efficient scaling of Flops/Watt.

References

- [1] Strategic Priorities Fund: ExCALIBUR Met Office Science Plan and Activities
- [2] <https://www.iter.org>
- [3] <https://www.gov.uk/government/news/uk-to-take-a-big-step-to-fusion-electricity>
- [4] <https://en.wikipedia.org/wiki/Tokamak>
- [5] <https://www.euro-fusion.org/programme/demo/>
- [6] <https://reuters.screenocean.com/record/1366443>
- [7] https://bout-dev.readthedocs.io/en/latest/user_docs/introduction.html
- [8] <https://en.wikipedia.org/wiki/Particle-in-cell>
- [9] <https://en.wikipedia.org/wiki/Gyrokinetics>
- [10] <http://www.cs.sandia.gov/~maherou/docs/MantevoOverview.pdf>
- [11] <https://society-rse.org>
- [12] <https://cfwebprod.sandia.gov/cfdocs/CompResearch/docs/Kokkos-Multi-CoE.pdf>
- [13] https://en.wikipedia.org/wiki/Domain-specific_language

- [14] https://en.wikipedia.org/wiki/Spectral_element_method
- [15] <https://www.sciencedirect.com/topics/engineering/discontinuous-galerkin>
- [16] <https://royalsocietypublishing.org/doi/10.1098/rsta.2018.0144>
- [17] <https://www.sciencedirect.com/science/article/pii/S0021999115003228>
- [18] <https://www.olcf.ornl.gov/center-projects/adios/>
- [19] <https://www.tandfonline.com/doi/abs/10.13182/FST47-172>
- [20] <https://www.york.ac.uk/physics/news/departmentalnews/plasma-fusion/major-grant-award-supports-fusion-energy-research/>
- [21] <https://arxiv.org/abs/1906.01200>
- [22] <https://parallel-in-time.org>
- [23] https://en.wikipedia.org/wiki/High_Bandwidth_Memory
- [24] <http://sagestorage.eu/about/overview>