

## Enabling validated exascale nuclear science

Andrew Davis<sup>1,\*</sup>, Aleksander Dubas<sup>1</sup>, and Ruben Otin<sup>1</sup>

<sup>1</sup>United Kingdom Atomic Energy Authority, Culham Science Centre, Abingdon, Oxfordshire, United Kingdom, OX14 3DB

**Abstract.** The field of fusion energy is about to enter the ITER era, for the first time we will have access to a device capable of producing 500 MW of fusion power, with plasmas lasting more than 300 seconds and with core temperatures in excess of 100-200 Million K. Engineering simulation for fusion, sits in an awkward position, a mixture of commercial and licensed tools are used, often with email driven transfer of data. In order to address the engineering simulation challenges of the future, the community must address simulation in a much more tightly coupled ecosystem, with a set of tools that can scale to take advantage of current petascale and upcoming exascale systems to address the design challenges of the ITER era.

### 1 Introduction

The goal of achieving magnetically confined controlled thermonuclear fusion has been the active research of several organisations, universities and collaborations since the early 1950's. Fusion is ostensibly in the delivery era evidenced by the rapid explosion of the number of startup companies and several independent national efforts. However, several of the planned startup and international efforts face an uncertain regulatory environment; stringent fission regulation being overly restrictive, but fusion is also not quite as uncertain as the high energy physics (HEP) safety environment where differences between models can impact results significantly. Several key materials that one would likely use to build a large scale reactor relevant facility have not been pre-qualified for the fusion reactor environment, the fission environment being too different in terms of radiation damage, gas production, and much lower operational temperature.

Given the lack of pre-qualified materials, the lack of availability of high intensity (greater than  $10^{15}$  neutrons per second) 14.08 MeV neutron sources, and the current absence of a reactor relevant test facility, then if we are to deliver fusion power on a short (pre 2050) timescale; we must qualify our reactor designs almost entirely *insilico*. This *insilico* design must be done with a lack of validation data in key areas, notably we lack any data for a representative multi-load case and several key physics processes have unknown data associated with them. The various load conditions experienced by components near the first wall is complex; there are neutron fluxes in excess of  $10^{15}$  neutrons  $\text{cm}^{-2}\text{s}^{-1}$ , photon heat loads from nuclear reactions in the core of the plasma, fast (greater than 20 keV) ions, 10's of Tesla magnetic fields, synchrotron radiation loads, mechanical loading, and several others. The components near the first wall of the tokamak can expect to experience temperatures in

---


\*e-mail: [andrew.davis@ukaea.uk](mailto:andrew.davis@ukaea.uk)

excess of 800-1000 K depending upon internal cooling rates. If we consider the range of heat fluxes seen for example on the divertor, fusion bears similarity to the materials challenges seen in advanced air breathing rocket engine designs, and in Formula 1 brake disc design in terms of durability requirements from very thin composite materials. Both of these fields have taken to complex geometric setups heavily laden with simulation driven development.

In the following sections we discuss the multiphysical nature of fusion related engineering challenges, outlining the current state of the art and a vision for the way forward in *insilico* fusion reactor and component qualification. Then we consider the related aspects of multi-scale simulation, necessary for accounting for the evolution of materials and their properties as they are exposed to fusion environments. Multiphysics and multiscale simulations are then combined into a fully-functioning digital representation, a digital twin.

## 2 Modelling Challenges

We can frame the modelling challenge for predictive modelling of fusion power plants around a number of key components that encompass geometric complexity, physical complexity, and criticality to device operations. Thus we can form a set of challenge problems to guide our modelling solution(s) around three particularly critical components; the Divertor, the Breeding Blanket, and the Magnet system, as shown in Figure 1.



**Figure 1.** ITER divertor and associated materials (upper) [9], example blanket module [10] (lower left) and superconducting conductor (lower right)

The divertor, an example here being the ITER divertor 1, is a set of components that lie (usually) at the bottom of the tokamak (for some designs so called double null there is a divertor at the top also), its purpose is to remove heat and helium ash from the plasma. The

divertor is subject to very high areal heat loads, in excess of  $10 \text{ MW m}^{-2}$  of charged particle loads, additionally there will be a significant (greater than  $10 \text{ MW m}^{-3}$ ) nuclear heat load, and magnetic fields of 10 Tesla in addition to others. The particular modelling challenge for the divertor comes from the large flow volumes within the divertor which feed an intricate and narrow web of cooling channels within divertor structures, stressing the scale of the problem for modelling turbulence and heat transfer.

In addition to the massive scale disparity between different parts of the divertor, there are multiple types of physics occurring, including thermo-mechanical, fluid turbulence, radiation transport, radio-chemistry, neutron activation, amongst others. Thus depending upon the quantity of interest, many physics must be solved for at any one time. There are similar challenges when modelling the magnet system and the breeding blanket. The breeding blanket has a number of specific challenges also, particularly one of the key materials produced in fusion, namely tritium. Tritium is both used as one of the fuels of fusion (and thus is consumed), but also tritium is bred inside of the blanket via neutron capture reactions. The tritium atom is small, and diffuses quickly through the small interstitial gaps between atoms in a lattice. The particularly complex nature arises in certain blanket designs from the many drivers of tritium diffusion; thermal gradients, chemistry, porosity, and production rates. That is, of course in addition to the other physical drivers mentioned earlier.

### 3 Multi-Physics Approach

Our traditional engineering analysis codes used within the fusion domain are not actually born of the requirements for multiphysics driven design. Being such a long multi-decadal research programme, the community has borrowed codes and software from other domains, e.g. ANSYS [1] is largely used, mostly due to its ubiquity and ease of use, MCNP [2] similarly is widely used and has an extensive set of fission relevant usage and benchmarking. There has never been an effort (to the authors knowledge) that a proper requirements gathering has ever been executed. Thus in order to cover the whole range of analysis to support the design of fusion reactors from construction, through startup to full operations, dynamic behaviour during discharge, post shutdown effect and off-normal operations; we must be capable of performing the following physics coupled together: thermo-mechanical, radiation (photons, neutrons and electrons), fluids (conjugate heat transfer and turbulence), electro-magnetics (static, quasi-static and full wave), macroscopic material deformation (solid mechanics), tritium transport, magneto-hydrodynamics, nuclear transmutation, and chemical reaction rates. This is not to say that this is an exhaustive list, as some of these collections of physics rely upon other simulations to predict coefficients and inputs.

#### 3.1 Geometry

Not only is there a computational challenge, bringing together a number of traditionally disparate physics regimes, there are also geometric complexity requirements; fusion reactor geometries are complex *matryoshka*-like with large components adjacent to thin components, with narrow intricate 3D printed cooling channels and composite materials. Performing whole component or even whole reactor calculations using traditional FEA based techniques may take extraordinary numbers of mesh elements estimated at between  $10^{10}$  -  $10^{14}$ . Generating meshes in geometries with such disparate element sizes is a challenge, so is the critical act of applying boundary conditions to model, and there is significant challenge generating that number of elements in a reasonable time frame. Generating meshes with  $10^9$  elements could take more than 24 hours, implying that any future meshing schemes must be not only

be robust, but must also scale across nodes due to memory requirements, but also across the given CPUs on a node for performance reasons. There are a few distributed parallel mesh generators; BoxerMesh from Cambridge Flow Solutions, which is a hex-dominant mesh generator, but uses an octree representation of geometry to remove defects in the underlying CAD geometry and Bolt from Coreform, which operates under a similar procedure. Thus, there is a workflow challenge in terms of getting high quality simulation ready meshes of the appropriate number of elements. Additionally on the horizon are high order methods such as Iso Geometric Analysis (IGA) and HP refinement methods which confer computational benefits but strongly differ from the current methods of mesh generation.

### 3.2 Coupling & Computation

We must simulate the physics that is needed to accurately represent our problems, both in isolation for comparison with traditional Commercial Off The Shelf (COTS) software e.g. ANSYS and for validation purposes. In terms of comparison, the COTS tools tend to employ monolithic coupling of only some problems, and commercial tools are not usually High Performance Computing (HPC) relevant. This means that with traditional coupling techniques, one must perform the iterative solve for a given physics solution, pass that solution by mesh transfer or by nodal result, and then converge the next physics solution, and so on until around the last physics package is reached, returning to the beginning and solving the whole loop again until convergence is reached. One could imagine leaving the COTS software solution and migrating to a more easily deployed open source toolchain; using ERMES [3] for Electro-Magnetics, Calculix [4] for thermo-mechanical, OpenFOAM [5] for fluids, and so on, but each of these codes must then be coupled using some mechanism, for example using preCISE [6]. Whilst perhaps philosophically satisfying, this solution is not particularly performant, particularly if any one of the coupled codes does not support distributed parallelism or does not scale well, producing a serial bottleneck and creating a limitation through Amdahl's law. If instead, we embrace FEA tools and libraries that have been shown to scale e.g. MFEM [7] or MOOSE [8]; we will not be limited by the slowest and non-parallel part of the sequence, since a) the sequence will be entirely block coupled and b) the fundamental libraries are MPI-hybrid (MPI & Threads) and some have GPU co-compute.

### 3.3 Multi-Scale

If we consider either the failure of materials due to cracking, the origins of cavitation in fluids, or the origins of superconductivity; they are all effects whose origins are at the micro or meso-scale but their impacts are upon the macroscopic world. In particular the propagation of cracks due to microstructural defects and the nature of magnetic field propagation in superconductors is of particular concern, as these have measurable impacts upon the performance of large scale structural and magnetic components. Focusing upon materials, there are a number of macroscopic material behaviours that we ascribe, such as elasto-visco-plasticity, that have coefficients that can be determined through phase field (crystal grain) modelling and some coefficients that can be determined through Density Functional Theory (DFT).

The simulations in which we determine the engineering response of structures will sometimes need to consider these meso and micro-scopic effects, as we have many fundamental forces at play which can impact the performance of materials. Critically, it is the response of materials to the various loads present in a tokamak that determines the durability, performance and the overall economics of fusion in terms of time to replacement of components, quantity of high grade heat extracted and several others critical measures. Further, if one

considers the diffusion of tritium through materials, there are a number of concurrent processes that are driven at different scales; diffusion, porous flow through surface of materials, chemical exchange and reactions, and traditional fluid flow. In order to perform simulations at multiple scale lengths, there must be a way of solving the physics of the smaller subsystem multiple times under different conditions, for example the way that MOOSE solves with problem is through the “MultiApp” [11] scheme, which through parent-child links define a directed acyclic graph through which the whole macro-scale problem can run, with the micro (and subsequently meso) scale functions being called several times. This type of scheme can be used to address shortcomings in macroscopic modelling, for example where two materials are bonded in traditional FEA there will be discontinuities in material response, but a multi-fidelity scheme can handle the bonded region more accurately.

### 3.4 Digital Twins

The term *Digital Twin* (DT) is a new term and has a multiplicity of definitions, The Centre for Digitally Built Britain (CDDDB) [12] defines a DT as “A realistic digital representation of something physical. What distinguishes a digital twin from any other digital model is its connection to the physical twin”. In the case of fusion, the current tranche of experiments are largely on machines that were built in the 1980’s and the record keeping that goes with older devices is enough to limit the utility of the DT. There are, however, a number of newer devices world wide that ultimately should be digitally twinned; those being ITER, MAST-U (UK), NSTX (USA), DTT (Italy) and SPARC (USA private). The value of twinning these devices demonstrates that we have a) a sufficient understanding of the physics at play and can reproduce real results, b) predict the dynamic response of these devices and therefore start to use DT’s to perform predictive modelling of components likely to fail and c) act as a source of validation information for multiphysics driven systems. However, the simulation tools and inputs as indicated in the previous sections are far from being capable of simulating all the physics in a fully coupled manner, therefore the prospect of having a DT of a full tokamak in the near term is slight. There are several other facilities in and beyond the fusion domain that one could digitally twin in the nearer term that proves out the concept. At the United Kingdom Atomic Energy (UKAEA) Culham site there are the HIVE and H3AT facilities that provide small scale tokamak relevant experiments for high heat flux and temperature. Both of these facilities could be twinned and the appropriate simulation toolchain used to predict performance, thus providing a set of validation data and some qualification of the software stack.

## 4 Digital Qualification Strategy

In terms of qualifying materials and whole systems according to regulatory requirements is a country specific endeavour and beyond the scope of this script. However, we layout arguments for the strategy for the cross validation of a suite of software tools to be used for fusion reactor design. It is clear from the timeline of fusion devices that are currently planned, along with other supporting facilities, there is no facility currently being planned that covers the appropriate radiation, thermal, magnetic, and plasma loads to be used to qualify materials and entire components. There are plans to qualify materials in IFMIF [13] which will irradiate small coupons of materials to power plant relevant Displacements Per Atom (DPA) which can be used to validate findings of atomistic Density Functional Theory (DFT) predictions of performance and DPA prediction, but those materials will not be exposed to appropriate plasma, temperature or magnetic conditions.

**Table 1.** (non-exhaustive) List of existing (or under construction) facilities useful for cross validation  
 note \* LHC here means detailed Quadrapole simulation

Facility	Type	EM	Fluids	Thermo-Mech	Radiation	MHD	Material
DRESDYN	MHD	X	X			X	
Diamond	Photon	X	X	X	X		
MAX IV	Photon	X	X	X	X		
H3AT	Fusion Test		X	X			
HIVE	Fusion Test	X	X	X			X
Chimera	Fusion Test	X	X	X			
IFMIF	Fusion Test		X	X	X		X
SNS	Spallation		X	X	X		X
ISIS	Spallation		X	X	X		X
nTOF	HEP		X	X	X		
LHC*	HEP	X	X	X	X		
MAST-U	Fusion	X	X	X	X		
ITER	Fusion	X	X	X	X		

Thus, in absence of the appropriate qualification facilities, we must find the next best set of facilities and perform overlapping validation. There are a number of operating facilities that can be used to validate single or dual physics, for example thermal hydraulic or liquid metal test facilities, but there are few sites that can handle multiple physics and as indicated earlier there are none that can perform all the physics (and chemistry) that are needed. A list of potential facilities that could be used are shown in Table 1.

Fusion will need to seek collaborations and access to some of the facilities in order to seek validation of physics packages in regimes of interest to fusion. Non-fusion candidates are useful when it allows validation of some but not all the physics, for example nTOF which would offer access to fluids, radiation, and thermo-mechanical data, LHC Quadrapole magnets which offers data pertinent to quench and beam divergence data. In most cases, quantities of interest are performance monitoring data such as flow rates, strain gauges, temperature monitor data, radiation count rates, currents, voltages, etc; all the data that could be used to validate real world engineering performance. These data will allow validation to occur across a range of temperatures, strains, fields and conditions and facilitate regulatory acceptance in fusion regime.

## 5 Conclusions

In order to achieve controlled magnetically confined fusion power in a short amount of time (before 2050), we must use *insilico* qualification of components in absence of experimental and operational data. We must leverage validated simulation, which itself must be cross validated with the available data and experiments from multiple domains. A key suite of experimental data across a wide range of temperatures, irradiations, magnetic fields, fluid types and flows and operational scenarios will result in the widest and most appropriate set of validation data. These simulations must embrace the exascale due to the sheer quantity of compute required. The software stack must be capable of utilising all the compute available and therefore must be able to use MPI, OpenMP (or other threading technologies) and will need to be able to leverage GPU co-compute in order to be able to scale to the whole exascale node. Fusion specifically is missing many validation experiments and thus we must look to other domains, especially spallation neutron sources and HEP experiments. These non-fusion

facilities are critical to provide validation data, and to gain regulatory acceptance. We must also start to integrate mesoscopic and microscopic effects into our macroscopic domains via the use of multiscale modelling, the inclusion of such effects will allow a much more physically accurate analysis. The simulation and the design workflow must be performed under the presence of uncertainty, and thus uncertainty propagation techniques will be important. Once we have access to a suite of validated codes, capable of performing multiphysics at large scale, wrapped with uncertainty quantification techniques; we will be able to revolutionise the way fusion engineering design is done currently, with the potential for impact on other engineering domains.

## References

- [1] Gabriel J. DeSalvo and John A. Swanson, ANSYS Engineering Analysis System User's Manual, Houston, Pa., Swanson Analysis Systems, 1985.
- [2] C.J. Werner, J.S. Bull, C.J. Solomon, et al., "MCNP6.2 Release Notes", LA-UR-18-20808, 2018
- [3] R. Otin, "ERMES: A nodal-based finite element code for electromagnetic simulations in frequency domain", *Computer Physics Communications*, **184**, 11, 2588-2595pp, 2013
- [4] P. A. Gustafson, F. A. Yapor Genao, et al, "Integration of MAC/GMC into CalculiX, an open source finite element code", *AIAA SciTech 2019 Forum*, 7-11th January 2019
- [5] H. G. Weller, G. Tabor, et al, "A tensorial approach to computational continuum mechanics using object-oriented techniques", *Computers in Physics*, **12**, 6, 1998
- [6] H.-J. Bungartz, B. Gatzhammer, et al, "preCICE - A Fully Parallel Library for Multi-Physics Surface Coupling", *Computers and Fluids*, **141**, 250-258pp, 2016
- [7] R. Anderson, A. Barker, et al, "MFEM: A Modular Finite Element Methods Library", DOI: 10.13140/RG.2.2.10719.92322
- [8] C. J. Permann, D. R. Gaston et al, "MOOSE: Enabling Massively Parallel Multiphysics Simulation", arXiv, 1911.04488, 2019
- [9] RA Pitts, S Bardin et al, "Physics conclusions in support of ITER W monoblock shaping", *Nuclear Materials and Energy*, 12, 6074pp, 2017
- [10] Songtao Wu, Yican Wu et al, "Concept of Multi-function Fusion Reactor", 2nd IAEA Technical Meeting on First Generation of Fusion Power plants, NE-O-12, 2014
- [11] D. R. Gaston, C. J. Permann, John W. Peterson, et al, "Physics-based multiscale coupling for full core nuclear reactor simulation", *Annals of Nuclear Energy*, **84**, 45-54pp, 2015
- [12] The Centre for Digitally Built Britain, <https://www.cdbb.cam.ac.uk/system/files/documents/TheGeminiPrinciples.pdf>, 2019
- [13] J. Knaster, P. Garin et al, "Overview of the IFMIF/EVEDA project", *Nuclear Fusion*, **57**, 10, 2017