

Use of massively parallel computing to improve modelling accuracy within the nuclear sector

LI M Evans^{1,2*}, J D Arregui-Mena², P M Mummery², R J Akers¹, E Surrey¹, A Shterenlikht³, M Broggi⁴, L Margetts²

1. CCFE, Culham Science Centre, Abingdon, Oxon, UK

2. School of Mechanical, Aerospace and Civil Engineering (MACE), University of Manchester, UK

3. Mechanical Engineering Department, The University of Bristol, UK

4. Institut für Bauinformatik, Leibniz Universität Hannover, Germany

ABSTRACT

The extreme environments found within the nuclear sector impose large safety factors on modelling analyses to ensure components operate in their desired manner. Improving analysis accuracy has clear value of increasing the design space that could lead to greater efficiency and reliability.

Novel materials for new reactor designs often exhibit non-linear behaviour; additionally material properties evolve due to in-service damage a combination that is difficult to model accurately. To better describe these complex behaviours a range of modelling techniques previously under-pursued due to computational expense are being developed.

This work presents recent advancements in three techniques: Uncertainty quantification (UQ); Cellular automata finite element (CAFE); Image based finite element methods (IBFEM). Case studies are presented demonstrating their suitability for use in nuclear engineering made possible by advancements in parallel computing hardware that is projected to be available for industry within the next decade costing of the order of \$100k.

1. INTRODUCTION

Engineering design for the nuclear sector brings together a particular set of demands not observed elsewhere [1]. Nuclear power plants must be built in such a way that they must not suffer catastrophic failure under any given set of perceivable circumstances from earthquakes [2] to terrorist attacks [3]. The inevitable wear and tear of components over the plant's lifespan must be well understood and predictable. At the plant's end of life it must be possible to be decommissioned, safely depositing activated parts of the machine ensuring long-term safety and sustainability of the surrounding area [4]. During operation, the constituent materials used

*Corresponding Author: llion.evans@ukaea.uk

for manufacturing the components undergo some of the most extreme environments known to humankind. Future generation nuclear power plants (such as gen IV molten salt reactors or magnetic confinement fusion devices) aim to increase energy output creating more extreme conditions; high thermal loading, extreme pressures, interaction with corrosive fluids, mechanical loads induced by electromagnetic fields, plasma erosion [5]. In addition to this is irradiation damage where product of the nuclear reaction created by burning fuel damages the surrounding material. In certain regions of a fusion device it is expected that every single atom will be “knocked off its perch” (displaced) up to twenty times per year, potentially either displacing the atom permanently or transmuting it into another element altogether [5]. The culmination of this is a constantly evolving and complex set of material properties due to the changing atomistic and microstructural make-up caused by a varied range of damage mechanisms.

In order to withstand these demands a new family of ‘super-materials’ are proposed which will have a set of material properties designed by materials scientists for a specific set of conditions typically exhibiting complex anisotropic non-linear behaviour. Such classes of materials contain functionally graded materials [6], novel three-dimensional composite weaves [7], nanomaterials [8], high entropy alloys [9], self-passivating alloys [10] etc.

Regardless of the difficulties in modelling such extreme environments, accurately describing the behaviour of such materials presents a significant challenge compounded by their changing properties. This has historically been achieved using homogenisation methods [11], a process that assumes average material properties. Changes in properties, due to damage evolution, are implemented via a (often sparse) database of empirical data. Therefore, large and restrictive safety factors are imposed on modelling analyses. Improving analysis accuracy has a clear value of increasing the design space that could lead to greater efficiency and reliability. Indeed, it is questionable whether the design of the 1st generation fleet of nuclear reactors would be allowed using current design tools; however, we know with experience that they can be operated safely.

The last nuclear fission power plant to be built in the UK started construction in 1987 [12] when the modelling techniques presented in this work were in their infancy and computational hardware was inadequate to perform them at scales relevant for engineering design. Due to recent advancements in high performance computing (HPC), it is now feasible to utilise such methods on sufficiently large simulations so that they can be used effectively to inform design decisions. However, they cannot be of general use to the community at large if they are to be the preserve of national research facilities. Predictions for computing hardware suggest that current HPC systems capable of performing these techniques will cost of the scale affordable to the industrial sector by the end of this decade [13]. As such, these techniques are being developed to maturity in tandem ready for such a time.

This work firstly presents a brief overview of computational hardware advances that have enabled advanced materials modelling techniques. Then three general technique areas are presented as avenues for improving modelling accuracy within the nuclear sector, i.e. uncertainty quantification (UQ), cellular automata finite element (CAFE) and image based finite element methods (IBFEM). Case studies are presented for each demonstrating their suitability for the nuclear sector. Finally, conclusions are drawn along with comments on future developments in the field.

2. ADVANCES IN COMPUTING HARDWARE ENABLING IMPROVED MODELLING

The last nuclear power station to be built in the UK, Sizewell-B, started construction in 1987 and therefore was designed years earlier using technology of that era. To put things into perspective the world's fastest computer in 1988 was the Cray Y-MP system capable of 2.6 GFlop/s (floating point operations per second) which is roughly comparable to an iPhone 4 (released in 2010) or the Intel Atom N2600 (released in 2011) both used as low power consuming mobile processors. Suffice to say that computing hardware has developed drastically since then [14]; currently the world's most powerful computer is Tianhe-2, China, which is capable of 33.9 PFlop/s. That is to say, high performance computing power has increased by over 13 million fold in less than 30 years. The power of Tianhe-2 alone is equivalent to the entire global population solving 4.7 million calculations per second; this hardware enables us to approach problems that were previously impractical to solve but only if software makes efficient use of this technology.

As with all current HPC systems, Tianhe-2 achieves faster computing times by utilising a greater number of computing cores rather than increased speed on a single processor [15]. This has been the standard practice since the mid-90s when vector computing fell out of favour and frequency scaling was abandoned [16]. Advances since 2010 have mostly been achieved through the introduction of heterogeneous supercomputers that use a mixed processing approach [17]. Typically, this consists of standard CPU processors coupled with GPUs but may also include field programmable gate arrays (FPGA) or bespoke coprocessors. This allows offloading of certain tasks to a different processing architecture better suited to the task, e.g. GPUs are particularly well suited to high-throughput tasks.

Although this hardware offers additional computing capability, this can only be used if software is written such that it can make use of what is available [17]. As with parallel computing, where problems need to be sub-divided for distribution over processors, heterogeneous computing requires determining which parts can be offloaded to the coprocessor. Although efforts have been made to automate this process, little headway has been made. Therefore, just as parallel coding involves an additional layer of complexity so does coding for coprocessors, which increases development time.

The Intel Xeon Phi processors, used in Tianhe-2, improve efficiency by including more computing hardware (i.e. computational cores) on a single processing board. In doing so it can process vast amount of data very quickly, the current limitation to speed-up is how quickly it can access this data. The architecture is made up of several tiers of memory, i.e. cache, RAM, HDD, each with increasing amount of space but 'further' away from the processor. Very large simulations can often have datasets that are terabytes in size, causing the data input/output (I/O) section of code to be the bottleneck. As computational power increases so too will the desire to handle larger datasets. The SAGE project [18], led by Seagate, aims to address this issue of hierarchical memory by using 'percipient storage methods' to allow computations that could be performed on any tier of data via advanced object based storage. This will be achieved by embedding the computational capabilities directly onto the storage thus drastically reducing data movement between compute and storage clusters, shown schematically in Figure 1.

As computing systems move towards exascale capability (a thousand petaflops), if processor power usage continues on its current trend, the demand for electricity will be prohibitively large. A single exascale system would require over a gigawatt of power, equal

to the output of the UK nuclear power station Dungeness-B. This issue has now been given great precedence in development of future systems such that there is a move away from measuring a HPC system's Flop/s to measuring Flop/s/W, as shown in the emergence of the Green500 list to rival the Top500 [19]. Intel's Blue Gene HPC systems have been designed to forgo raw speed for low power as an early prototype for technology that could be used in the first exascale system [20]. The case study presented in section 5.2 of this paper used one of these low power systems.

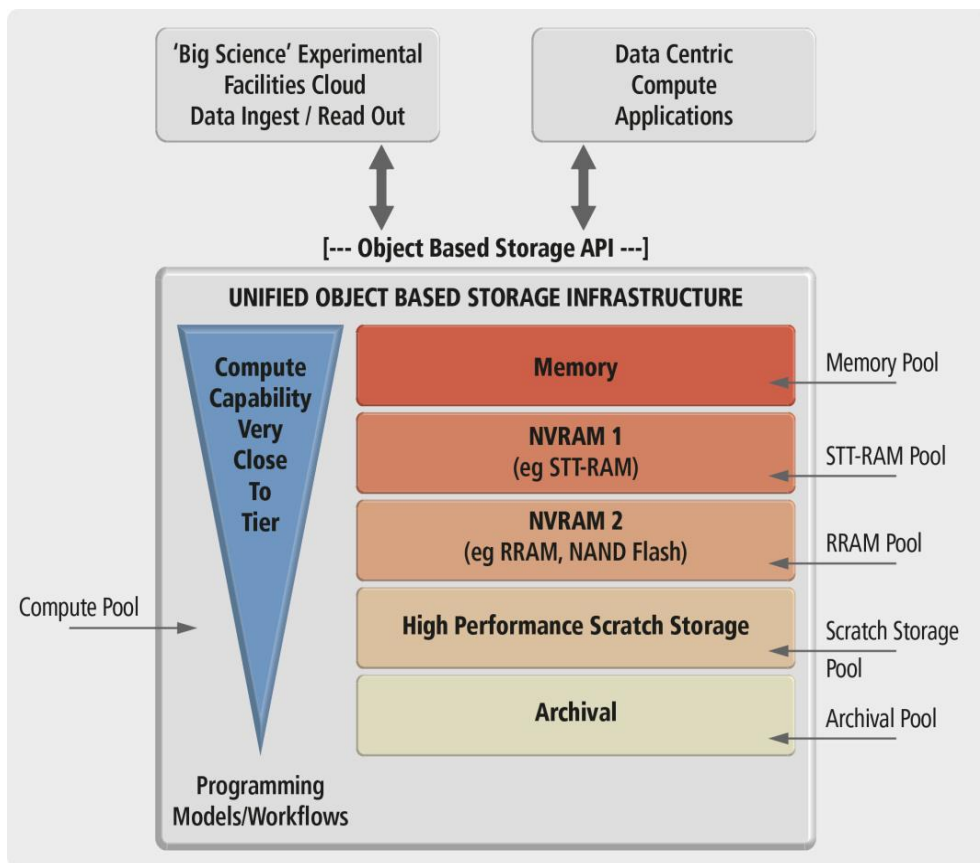


Figure 1. Data centric computing architecture as proposed by the SAGE project [18].

3. UNCERTAINTY QUANTIFICATION (UQ)

3.1. Concept background

Advancements in computational capabilities allow engineers to design structures which are more complex and with more precision than ever before. However, connecting the virtual digital world with the real one is still the major challenge of modern engineering. Engineering practice is to use deterministic modelling which often means one set of inputs such as material properties and boundary conditions that yield one set of results e.g. [21]. If the same computational analysis were to be performed multiple times, the output would be identical each time. Whereas repeat testing of similar samples would exhibit a range of outcomes.

Firstly, material property values are always an average of a large number of data points collected via experiments that aim to measure the response of a bulk volume, thus smoothing out localised variations [22]. Secondly, the environmental conditions or loading of a component (mechanically or otherwise) will rarely be identical; even under controlled laboratory conditions this scatter is taken into consideration.

A good example of this is a high heat flux component within a fusion reactor that undergoes thermal cycling. Finite element analysis (FEA) of the component will always show the peak stresses in exactly the same location because the inputs (material properties and boundary conditions) do not change, as shown in Figure 2a) [23]. In reality, not all parts fail identically; this is due to variations in these characteristics both locally within one sample and globally from sample to sample, as shown in Figure 2b) [24]. Additionally, it is unlikely that the manufacturing process will yield identical samples; therefore, variations in geometry, i.e. deviations from design, are also possible.




Figure 2. Fusion heat exchanger component having been tested a) computationally where results are identical each time and (isometric view) [23] b) experimentally where failure location changes due to range of variations (plan view) [24].

When data is reported from a part qualification testing campaign the scatter of the experimental results is an important component of the outcome [25]. It gives an understanding to the performance predictability and therefore a sense of what limits should be adhered to during operation. Furthermore, what the likelihood of failure is at the specified limit even though it is beneath the maximum allowable loading level [26]. This is also a more appropriate approach to simulation rather than the deterministic ‘one set in, one set out’ approach. By employing stochastic modelling this is what the uncertainty quantification (UQ) technique aims to achieve. Utilising a distribution of inputs results in a range of outputs, the challenge is in determining how the inputs and outputs should be distributed and interpreted, respectively. This is shown schematically in Figure 3.

Even within the FEA engineering community ‘uncertainty quantification’ is a broad term with a range of applications and implementations but with similar goals in aiming to quantify

safe bounds of operation. In practice UQ is normally implemented by repeating the same analysis whilst performing sweeps (i.e. changing input values) of all the variables of interest. The output of this is a large database of results that require interrogating to define a multidimensional safe operational zone. Depending on the number of variables to be swept and the degree of detail required, the total number of simulations required may increase exponentially which is one of the greatest challenges in implementing UQ. If the original deterministic model was computationally expensive, performing UQ may be prohibitively so. For example, to reliably compute failure probabilities to the order of 10^{-2} , thousands of simulations might be needed. When increasing to a target failure probability to the order of $10^{-6} - 10^{-8}$ the number of simulations would be too numerous even with current HPC capabilities. Thus, ‘smart’ algorithms such as advanced Monte Carlo simulations would be required to reduce the number of simulations to a level appropriate for HPC [27].

Until recently, performing stochastic FE for UQ purposes with standard commercial FEA packages was not straightforward [28]. This was because it required substantial manual interaction to set up variable sweeps and then collate the required data from the vast output before UQ analysis could be performed. COSSAN is a general-purpose software package for managing the approach to UQ by integrating the various codes required to perform all stages of the analysis [29]. By efficient management of the computing tasks within a HPC architecture and streamlining the workflow its aim is to make UQ more approachable and give added value with bespoke toolboxes for stochastic analysis. COSSAN’s use has previously been demonstrated successfully with large FE models [27].




Figure 3. Schematic representation of differences between deterministic and stochastic modelling input and output [30].

3.2. Case study – Stochastic material properties with random finite element methods (RFEM)

Graphite is used as a moderator in the core of most fission reactors in the UK. It is known that graphite has large localised variations in its properties [31] meaning that the response of

different parts made from the same material to the extreme environments of a nuclear reactor can vary significantly [32]. More recently, it has been reported that sections of the graphite core in Hunterston-B have failed before they were expected to do so resulting in costly maintenance periods [33]. Better understanding of the variability of the graphite would aid engineers in planning for such events. Here a case study is presented on work by Arregui-Mena et al [34] that used random finite element methods (RFEM) for the purpose of UQ with graphite bricks designed for the core of an advanced gas-cooled reactor (AGR).

As previously discussed, one of the more difficult tasks in UQ is choosing a method by which the model inputs are distributed. RFEM is a technique originating from soil mechanics where it is known that there are large localised variations in material properties [35]. Within FEA each element can contain its own unique set of material properties, this capability is not usually used as it is more convenient to assign the exact same properties to all elements belonging to the same material type. RFEM takes deterministic material properties, such as elastic modulus, and varies the material properties by randomly assigning each element a different value. So that the material still behaves similar to the 'real' material globally, the randomisation happens within distribution parameters set by the user with the deterministic property acting as the mean of the distribution. Depending on the intended use, the spatial spread could be completely random or linked to a characteristic length scale, using a 'local average subdivision' method.

The statistical distribution (i.e. mean, variance and spatial correlation length) of the coefficient of thermal expansion (CTE) of a graphite core brick was collected experimentally by dissecting a brick at various locations. Once the geometry of the brick had been created and meshed for FE the CTE was applied in-line with the RFEM procedure. Each random distribution of the properties is termed a 'realisation'; for this analysis one hundred realisations were created. As the number of realisations increase, the average of all their results should tend towards that of the deterministic analysis. By comparing displacement (strain) results for the deterministic and stochastic models, it was decided that one hundred realisations was sufficient for this study. One realisation of random CTE distribution as compared to the deterministic version can be seen in Figure 4. For the FE analysis, the bricks were subjected to a linear gradient temperature change from their centre to represent the heating from a fuel rod, as shown in Figure 5. To calculate the results of the hundred realisations simultaneously, ParaFEM [36–39], an open source parallel FE code was used.

By visual inspection, it is clear from Figure 6 that the results for the deterministic model are periodically symmetric, as would be expected for periodically symmetric geometries and boundary conditions. However, this is not the case for the stochastic model (see Figure 7) where it is clear that the variation in CTE impacts the distribution of Von Mises stress. Additionally, the distribution of stress is unique for the three realisations shown in Figure 7, which is as expected. The power of RFEM does not come from investigating the individual realisations but when considering the results holistically. Figure 8 shows that the spread of data is large meaning there would be a wide variation in performance of these bricks. Additionally, when considering the average, the maximum stresses are around 10 % higher than reported by the deterministic model meaning that the maximum allowable temperatures would be overestimated. If it was assumed that operational limits have been chosen such that no bricks were to fail, the stochastic realisation with the largest stresses are more than twice that of the deterministic model. These results demonstrate the potential value UQ has for the nuclear sector not only in design of new reactors but also in continued operation of the UK's current fleet.




Figure 4. Coefficient of thermal expansion for a) the deterministic model and b) one stochastic realisation [34].

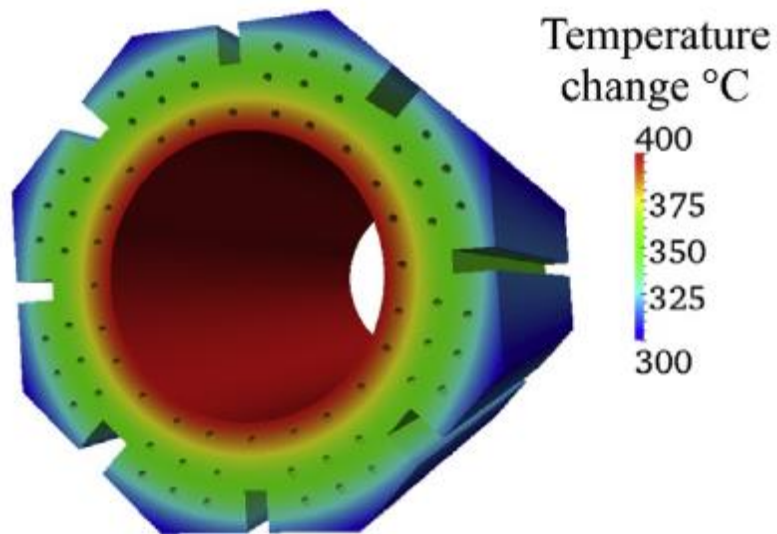


Figure 5. Temperature profile used to represent heating from fuel rod in the centre [34].




Figure 6. Von Mises stress for the deterministic model as calculated by FEA [34].




Figure 7. Von Mises stress for three realisations of the stochastic model as calculated by FEA [34].




Figure 8. Von Mises stress profile along inner bore wall travelling the length of the brick [34].

4. CELLULAR AUTOMATA / FINITE ELEMENT (CAFE)

4.1. Concept background

The cellular automaton (CA) approach is similar to that of FEA by which a domain is discretised into grid of ‘cells’ for analysis, as shown in Figure 9. Rather than to calculate a ‘degree of freedom’, CA is used to describe and calculate the cell’s discrete ‘state’. Although this grid can have an infinite number of dimensions, the number of states is finite. Investigations to the use of CA in materials science have been varied [40]. A frequent focus is on crystalline materials due to discrete states that can be considered. In this instance, CA is coupled with finite element (FE) to create cellular automata / finite element (CAFE).




Figure 9. Illustration of using CA to simulate a 3D solidification process, here a 2D slice of a single iteration is shown. ‘0’ is the liquid phase with other numbers representing unique grains, the arrows show the copying of cell states [41].

For standard FE the typical workflow uses boundary conditions to describe the environmental state (e.g. global variables, loading, sinks and sources) with the material’s response calculated from its inherent material properties. If any of the boundary conditions were to cause a change in the material properties, these would need to be updated from a pre-known look up table [42]. For example, material properties are well known to have strong thermal dependence, therefore if the temperature within the model deviates significantly from the initial conditions the properties (e.g. modulus, thermal conductivity etc.) are updated to an interpolated value within a known given range previously measured experimentally.

The drawback for standard FE is that, under certain circumstances, this can be problematic because this data is oversimplified by only taking into consideration one dimension. Whereas in actuality a whole range of mechanisms such as temperature, pressure and loading could be contributing to the changes in properties. Alternatively, the data may not exist because it is unfeasible to measure experimentally e.g., extreme environments can be too challenging or expensive to recreate under laboratory conditions. In these cases, large extrapolations or approximations are made for material properties that lead to lower confidence in modelling results and therefore larger safety factors must be imposed.

CAFE aims to provide a solution to this issue by coupling CA to the FE workflow in order to calculate changes to the material structure and how this impacts the material's response. This is done by taking the continuum field output such as temperature and strain fields from the FE layer at each time increment to be superimposed on the CA layer. The CA layer calculates the material response accordingly at a sub FE element level by subdividing each element into a CA 'neighbourhood'. Once complete this updates the state of the material microstructure and returns back to the FE layer a set of damage variables. This could include quasi brittle fracture or void evolution etc. [43].

Because each FE element requires subdividing into a CA grid, the computational expense of this process significantly increases with the number of elements. In the past, this has restricted investigations to relatively small volumes of material but advances in HPC systems now make this process useable for engineering applications. In the nuclear sector of particular interest would be use of CAFE for predicting ductile to brittle transition temperature, grain instability, solidification, recrystallization and dynamic strain induced transformation.

4.2. Case study - Cleavage propagation across crystal boundaries

Cleavage propagation in polycrystalline materials is a key issue in the design of nuclear vessels. Although alloying elements can improve material performance by creating sinks which mitigate crack propagation it is not possible to suppress completely. Thus, accurate modelling of their behaviour can lead to a better understanding of how they may impact material degradation and thus structural integrity. Here a brief overview of recent progress by Shterenlikht & Margetts [41] to develop the CAFE technique to model cleavage propagation in polycrystalline materials is presented as a case study relevant to nuclear materials.

A cube geometry with dimensions of 10 x 10 x 10 mm was specified. This was meshed for FE purposes using an edge seeding of 0.5 mm for hexahedral elements. The material was treated as linearly elastic and given an elastic modulus of 200 GPa. A polycrystalline volume was digitally engineered using randomized generation methods resulting in a distribution of both grain sizes and orientations. A mean grain diameter of 1 mm was specified resulting in a total of 1000 grains. By comparison to theoretical grain distributions [44] this method was shown to reliably create realistic size and orientation distributions in addition to boundary topology. When subdividing for CA, the resolution required to ensure mesh independence is 105 grid cells per crystal [45].

Other than standard fixing conditions to avoid free movement the only boundary condition prescribed within the FE layer was a distributed loading of 1 kN normal to the plane near one corner of the cube (see Figure 10). To initiate cracking a site was set at the coordinate location 0, 0, 5 mm. To represent pre-existing nano-cracks it would be possible to place any number of crack initiation sites whose locations were randomly distributed. In this case, it would be possible to investigate their interaction in addition to propagation.

Modelling was performed by coupling the CA code CGPACK [46] with the FE code ParaFEM [39]. The CE and FE codes were solely responsible for the cleavage and mechanical iterations, respectively and the calculated variable result values were passed from one to the other between time iterations.

To advance the cleavage iterations the CA code scans over all cells to see if any intact cells have a cleaved neighbour. When it finds such a cell, it checks whether cleavage conditions are satisfied, if so the model crack advances for a characteristic length that is related to the mean grain size. In this example, the cleavage criterion was an equation that linked the normal

stresses to the surface energy and relaxation distance. If the crack reached a grain boundary further criterion were considered before the crack could propagate, e.g. if a grain already included a crack. Using this CA method, a mesh independent cleavage is achieved based on the critical stresses and characteristic length scales.

For this cube model, three runs of the same simulation were performed each with unique results, as shown in Figure 11. The yellow and green cracks show clusters of cracks on the $\{100\}$ and $\{110\}$ planes respectively. $\{111\}$ planes have very high surface energies and are therefore unlikely to exhibit cleavage in practice. It can be seen from the results that the crack clusters combine to form a large crack normal to the direction of maximum stress. If this simulation was repeated a statistically significant number of times it would be possible to collate enough data to quantify the associated scatter, a form of UQ.




Figure 10. FE results of elastic 10 x 10 x 10 mm cube with modulus = 200 GPa loaded by 1 kN [41].




Figure 11. Three runs for CAFE simulation of cube with each showing unique cleavage propagation results [41].

5. IMAGE BASED FINITE ELEMENT METHODS (IBFEM)

5.1. Concept background

Image based finite element methods (IBFEM) is the process by which three-dimensional images of a part are converted directly into a finite element mesh for the purpose of simulation. The technique encompasses a range of methods that can be used to achieve the final model and can cover a range of length scales. These meshes are used in place of those created by computer-aided design (CAD). Some of the benefits associated with IBFEM are:

- Real components can be simulated at a microstructural level [47].
- It is easily possible to model complex architectures such as foams [48].
- Fabrication defects (e.g. micro-cracking or porosity) are inherently included [49].

Because of this, it is then possible to perform both an experimental test and a simulation of that test on the same exact sample for validation purposes. This has been performed in the past comparing thermal stresses at the granular level simulated with IBFEM with those measured experimentally using synchrotron X-ray diffraction [50]. An extension to this would be to simulate quality assurance tests on a scanned manufactured component, thus performing a virtual part qualification.

The workflow can be separated into three distinct components; three-dimensional imaging; mesh creation; finite element analysis.

Numerous three-dimensional imaging techniques exist e.g. X-ray or neutron tomography, magnetic resonance imaging, LIDAR, multibeam sonar, SEM sectioning etc. The method of choice will depend on the compatibility of the signal with medium being imaged, e.g. LIDAR requires translucence. Other considerations include the scale of the features of interest or if the signal may detrimentally affect the sample. For example, neutron tomography is unlikely to be appropriate for in vivo scanning although possible. The final length scales at which the modelling will be performed will depend on the initial image, owing to the fact it is not possible to improve resolution downstream in the workflow path. In recent literature, X-ray tomography is the technique most commonly used for IBFEM, likely due to its roots stemming from biomechanics [51] where three-dimensional imaging is standard practice to diagnose patients.

Converting three-dimensional images requires specialised software that first distinguishes and separates various regions of the image into segmented volume sections representative of constituent materials. Labels and material properties are assigned to each of these sections before an algorithm is used to subdivide the volume into elements connected by a global mesh of nodes and vertices. Because one of the main reasons for using IBFEM is to model a very complicated geometry, meshes are usually highly unstructured.

FEA of an image-based model is very similar to CAD based FEA, requiring setting of boundary conditions and solution parameters. The main difference is that meshes tend to be comparatively large to describe the complex geometries accurately. CAD based meshes normally have an element count of the order of 50,000 whereas image-based meshes are of the order of 100,000,000. It is for this reason why IBFEM is still an emerging technique as the current range of commercial FEA packages are known to scale poorly on parallel computing systems and are thus ill suited for IBFEM [52]. They also include a prohibitive licence per compute node pricing strategy that significantly increases the cost of performing such simulations. Several open source, therefore free, codes exist that claim to be ultra- scalable

and developed specifically for utilising HPC architecture. For the IBFEM case study presented in this paper, the open source code ParaFEM [39] was used.

5.2. Case study – Ceramic-Metal joining in fusion reactor component

The divertor section in a tokamak is used to exhaust heat and helium ash produced by the fusion reaction and is the region that will experience the highest thermal loads under normal operation [53]. ITER is currently being built in Cadarache, France and when operational it will be the world's biggest fusion device. Thermal fluxes in its divertor will be around 10 MW/m²; to withstand such thermal loads the divertor will contain an armour region made up of hundreds of thousands of actively cooled tiles. An early design concept, presented here, consisted of carbon fibre composite (CFC) tiles (approx. 30 x 30 x 4 mm) with a hole in the middle through which a copper alloy coolant pipe (approx. 10 mm diameter) is passed whereupon both parts are joined together (coined a 'monoblock') [54]. The difficulty in achieving this is that the copper alloy has a significantly greater coefficient of thermal expansion than that of the CFC. Therefore, on thermal loading the pipe would expand at a greater rate than the surrounding CFC armour inducing stresses. To mitigate this, a compliant copper interlayer was introduced to the design, located at the interface between the armour and pipe. However, increasing the number of joining interfaces also increased the potential for creating thermal barriers, which would reduce the component's efficiency at extracting heat and would increase likelihood of failure (or reduce lifespan). Investigations were made into various joining techniques for suitability of monoblock manufacturing considering cost vs quality of join. In work by Evans et al [55] IBFEM was used to investigate one candidate technique which managed to reduce manufacturing costs by using brazing for joining but also increased join quality by pre-coating 'off the shelf' commercially available brazing foils with chromium which serves to enhance bonding at the interface. A brief overview of that work is presented here as a case study to demonstrate how IBFEM can be used for modelling in nuclear engineering.

X-ray tomography imaging of the part was performed at the Manchester X-ray Imaging Facility [56], University of Manchester, UK, using a Nikon Metrology 225/320 kV system (with the 225 kV source). The resultant image had a voxel (three-dimensional pixel) width of 21.8 x 10⁻⁶ m with a total of 208 million voxels used to describe the part; this would have resulted in a FEA mesh of around a billion tetrahedral elements. To reduce computational expense, the image was downsampled to 50% of its original resolution, achieving a compromise between computational expense whilst retaining microstructural features of interest. The final voxel width was 43.6 x 10⁻⁶ m resulting in 137 million tetrahedral elements. Conversion of the image to FEA mesh was achieved using the Simpleware [57] suite of programmes, version 6 (Simpleware Ltd., Exeter, Devon, UK).

FEA analysis was performed using ParaFEM, (revision 1796) [58], an open source parallel finite element platform developed by the authors using a IBM BlueGene/Q system hosted at HPC The Hartree Centre, STFC, UK. Analysis and visualisation of results used ParaView version 3.14.1 64 bit (Kitware Inc., Clifton Park, New York, USA) [59].

X-ray imaging of the part showed a significant amount of porosity within the CFC, which, importantly, had a preferential alignment with the fibre direction. Thus, it was possible to input isotropic material properties for the carbon phase as the inclusion of porosity within the mesh would lead to anisotropic behaviour on the macroscopic scale. In practice, it is known that CFCs have higher thermal conductivity in plane with fibre layers compared to across fibre

layers. This is a twofold heat transfer mechanism with heat preferentially travelling along fibres and porosity alignment (which runs parallel to fibres) creating orthogonal thermal barriers.

The second feature of note within the sample was a region of debonding between the coolant pipe and armour. This occurred during the manufacturing process whilst cooling from its joining temperature and presented itself as a very thin but expansive region covering almost half of the area at the pipe-armour interface. The region was so small that it was not visible by eye and because it was an unexpected result of the manufacturing process, it would not have been included in analysis at the design stage.

To test the impact of the microstructure on the performance IBFEM modelling of this component was performed under fusion reactor conditions. To replicate these, the boundary conditions applied were a coolant temperature of 150 °C in the pipe and a thermal flux of 10 MW/m² on the plasma-facing surface. To further investigate the effect of the debonding region, the simulation was performed twice, firstly with the void between the heat sink (coolant) and heat source (plasma) then rotating the part through 180° such that the debonded region was 'behind' the pipe, as shown in Figure 12.

Results showed that both the porosity and debonded region acted as thermal barriers within the component but that the effect of the latter was significantly greater. An example visualisation of the results when the component had reached steady state can be seen in Figure 13. When comparing the two simulations it was observed that when the debonding was between the heat source and sink that the maximum temperature in the part was over 100 °C higher than when rotated through 180°, as shown in Figure 14. Additionally, thermal gradients were greater which would lead to higher thermally induced stresses thus increasing likelihood of failure or reducing lifespan. These are effects which are non-negligible and therefore must be considered for the part's proposed use but would not have been accounted for had only CAD based been used. This advanced form of modelling was only made feasible due to use of parallel computing.

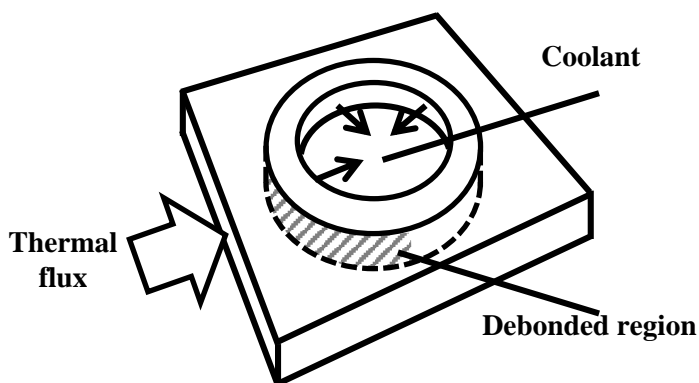


Figure 12. Schematic of part modelled by IBFEM, applied boundary conditions and debonding region are shown [55].




Figure 13. Visualised results from IBFEM analysis of part with dashed line denoting temperature profile used for Figure 14 [55].

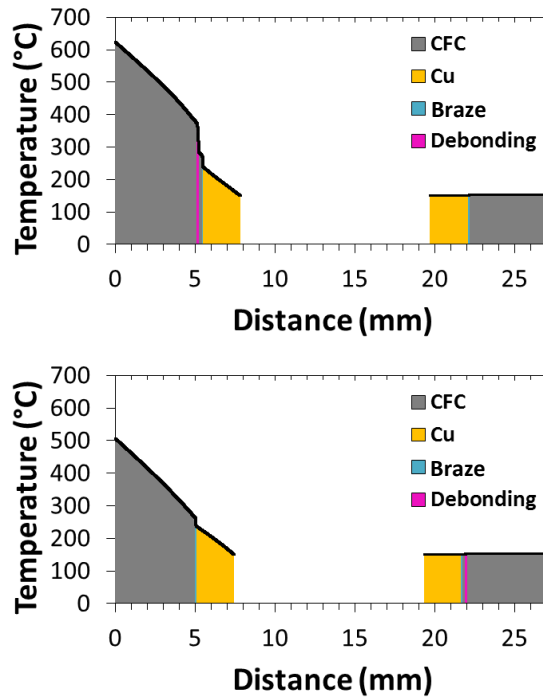


Figure 14. Temperature profile between the front and rear surfaces of the part with debonding region a) between the heat source and sink and b) 'behind' the heat sink [55].

6. CONCLUSIONS AND PREDICTIONS FOR FUTURE DEVELOPMENT

In this paper, recent advancements in computing architecture were discussed and how additional computational power, via heterogeneous systems, enables solution of complex problems where it would previously have been unfeasible to do so. Of particular interest to the nuclear sector is advanced modelling of novel materials, which typically exhibit anisotropic non-linear behaviour. Additionally, the future challenges to further computing advancements were noted (i.e. I/O and energy efficiency) and which avenues are being investigated to resolve this. The potential financial and efficiency gains of fewer or shorter enforced maintenance windows due to improved modelling are significant because availability is a key driver in determining the cost of electricity from a nuclear plant.

Then three modelling techniques were introduced that aim to simulate more realistic behaviours and have direct relevance to the nuclear sector. The practical use of these methods use has only been made possible by the aforementioned computational advances. The techniques introduced were:

- Uncertainty quantification (UQ) for predicting the degree of scatter expected in ‘real world’ scenarios.
- Cellular automata / finite elements (CAFE) which is used for analysing changes in discrete states at the sub element level and how this impacts material response.
- Image based finite element methods (IBFEM) for modelling components ‘as manufactured’ rather than ‘as designed’.

Case studies relevant to nuclear engineering were presented for each technique. In order to perform these studies suitable software able to effectively utilise HPC systems was needed. As commercial FEA packages do not scale well an open source alternative, ParaFEM, was used instead. This issue poses a significant barrier in the uptake of these techniques in the nuclear sector. Open source codes are not widely adopted by industry because they can be less ‘user friendly’, documentation can be sparse and when support is needed, it is often provided by a community of volunteers. However, the benefits of open source code are its low cost, efficiency, ability to view and customise all operations [60]. It is often assumed that commercial code is more accurate, but it has been proven that this is not the case [61]. However, commercial packages have a proven track record and undergo extensive validation testing to become certified for use within the nuclear sector. Because the HPC systems previously mentioned will be affordable to industry in the near future, it is pertinent to consider now which software will be used on such systems at that point in time. If the techniques presented here are to be used by the nuclear sector with open source code, it will require their initiation to drive certification effort because the open source community do not have the mechanism to recuperate the cost for this effort intensive certification procedure.

7. ACKNOWLEDGEMENTS

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053 and from the RCUK Energy Programme [grant number EP/I501045]. To obtain further information on the data and models underlying this paper please contact PublicationsManager@ccfe.ac.uk. The views and opinions expressed herein do not

necessarily reflect those of the European Commission. This work made use the HPC resources of The Hartree Centre (project fusionFEM) made available within the Distributed European Computing Initiative (DECI-12) by the PRACE-2IP, receiving funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement RI-283493.

Copyright © 2016 United Kingdom Atomic Energy Authority, all rights reserved.

REFERENCES

- [1] K.L. Murty, I. Charit, Structural materials for Gen-IV nuclear reactors: Challenges and opportunities, *J. Nucl. Mater.* 2008. 383 (1-2) pp. 189–195.
- [2] N.M. Newmark, W.J. Hall, Seismic design criteria for nuclear reactor facilities, in: *Proc. 4th World Conf. Earthq. Eng. Santiago, Chile, 1969*: pp. 37–50.
- [3] C. Behrens, M. Holt, Nuclear Power Plants: Vulnerability to Terrorists Attack [ADA444804], Washington, DC, US, 2005. <http://oai.dtic.mil/oai/oai?verb=getRecord&metadataPrefix=html&identifier=ADA444804> (accessed January 11, 2016).
- [4] L. Stamford, A. Azapagic, Sustainability indicators for the assessment of nuclear power, *Energy*. 2011. 36 (10) pp. 6037–6057.
- [5] S.J. Zinkle, J.T. Busby, Structural materials for fission & fusion energy, *Mater. Today*. 2009. 12 (11) pp. 12–19.
- [6] G.N. Praveen, J.N. Reddy, Nonlinear transient thermoelastic analysis of functionally graded ceramic-metal plates, *Int. J. Solids Struct.* 1998. 35 (33) pp. 4457–4476.
- [7] L.L. Snead, T.D. Burchell, A.L. Qualls, Strength of neutron-irradiated high-quality 3D carbon fiber composite, *J. Nucl. Mater.* 2003. 321 (2-3) pp. 165–169.
- [8] I.J. Beyerlein, A. Caro, M.J. Demkowicz, N.A. Mara, A. Misra, B.P. Uberuaga, Radiation damage tolerant nanomaterials, *Mater. Today*. 2013. 16 (11) pp. 443–449.
- [9] J.-W. Yeh, Recent progress in high-entropy alloys, *Ann. Chim. - Sci. Des Matériaux*. 2006. 31 (6) pp. 633–648.
- [10] F. Koch, S. Köppl, H. Bolt, Self passivating W-based alloys as plasma-facing material, *J. Nucl. Mater.* 2009. 386-388 pp. 572–574.
- [11] J.F. Bourgat, Numerical experiments of the homogenization method, in: R. Glowinski, J.L. Lions, I. Laboria (Eds.), *Comput. Methods Appl. Sci. Eng.* 1977, I, Springer Berlin Heidelberg, 1979: pp. 330–356.
- [12] B.V. George, J.A. Board, The Sizewell B design, *Nucl. Energy*. 1987. 26 (3) pp. 133–148.
- [13] L. Popa-Simil, Using High Performance Scientific Computing to Accelerate the Discovery and Design of Nuclear Power Applications, in: R. Segall, J. Cook, Q. Zhang (Eds.), *Res. Appl. Glob. Supercomput.*, IGI Global, 2015: pp. 119–148.
- [14] E. Strohmaier, H.W. Meuer, J. Dongarra, H.D. Simon, The TOP500 List and Progress in High-Performance Computing, *Computer (Long. Beach. Calif)*. 2015. 48 (11) pp. 42–49.
- [15] X. Liao, L. Xiao, C. Yang, Y. Lu, MilkyWay-2 supercomputer: system and application, *Front. Comput. Sci.* 2014. 8 (3) pp. 345–356.

- [16] P. Gepner, M.F. Kowalik, Multi-Core Processors: New Way to Achieve High System Performance, in: *Int. Symp. Parallel Comput. Electr. Eng.*, IEEE, 2006: pp. 9–13.
- [17] P. Bientinesi, J.R. Herrero, E.S. Quintana-Ortí, R. Strzodka, Parallel computing on graphics processing units and heterogeneous platforms, *Concurr. Comput. Pract. Exp.* 2015. 27 (6) pp. 1525–1527.
- [18] Sage Storage [Online]. Available: <http://www.sagestorage.eu/> (accessed January 11, 2016).
- [19] S. Hemmert, Green HPC: From Nice to Necessity, *Comput. Sci. Eng.* 2010. 12 (6) pp. 8–10.
- [20] R. Haring, M. Ohmacht, T. Fox, M. Gschwind, D. Satterfield, K. Sugavanam, P. Coteus, P. Heidelberger, M. Blumrich, R. Wisniewski, alan gara, G. Chiu, P. Boyle, N. Chist, C. Kim, The IBM Blue Gene/Q Compute Chip, *IEEE Micro.* 2012. 32 (2) pp. 48–60.
- [21] E. Rincón, J. Alonso, G. Barrera, J. Botija, P. Fernández, M. Medrano, G. Pérez, F. Ramos, A. Soletto, P. Barabaschi, E. Di Pietro, L. Meunier, A. Sakasai, K. Masaki, Y. Shibama, Structural analysis of the JT-60SA cryostat base, *Fusion Eng. Des.* 2011. 86 (6-8) pp. 623–626.
- [22] P. Bowen, J.F. Knott, Size effects on the microscopic cleavage fracture stress, σF^* , in martensitic microstructures, *Metall. Trans. A.* 1986. 17 (2) pp. 231–241.
- [23] T. Barrett, D. Hancock, M. Kalsey, W. Timmis, M. Porton, Design Study of a Water-Cooled Divertor: Alternative Concepts, Report for TA WP12-DAS-02-T02 [EFDA D 2MA647], 2012.
- [24] P. Gavila, B. Riccardi, G. Pintsuk, G. Ritz, V. Kuznetsov, A. Durocher, High heat flux testing of EU tungsten monoblock mock-ups for the ITER divertor, *Fusion Eng. Des.* 2015. 98-99 pp. 1305–1309.
- [25] M.H. Ramsey, S.L.R. Ellison, H. Czichos, W. Hässelbarth, H. Ischi, W. Wegscheider, B. Brookman, A. Zschunke, H. Frenz, M. Golze, M. Hedrich, A. Schmidt, T. Steiger, Quality in Measurement and Testing, in: H. Czichos, T. Saito, L. Smith (Eds.), Springer Handb. Metrol. Test., Springer-Verlag, Berlin Heidelberg, 2011: pp. 39–141.
- [26] R.F. Stapelberg, Reliability and Performance in Engineering Design, in: *Handb. Reliab. Availability, Maintainab. Saf. Eng. Des.*, Springer London, London, 2009: pp. 43–294.
- [27] E. Patelli, H. Murat Panayirci, M. Broggi, B. Goller, P. Beaurepaire, H.J. Pradlwarter, G.I. Schuëller, General purpose software for efficient uncertainty management of large finite element models, *Finite Elem. Anal. Des.* 2012. 51 pp. 31–48.
- [28] J.D. Arregui-Mena, L. Margetts, P.M. Mummery, Practical Application of the Stochastic Finite Element Method, *Arch. Comput. Methods Eng.* 2014. 23 (1) pp. 171–190.
- [29] COSSAN-X [Online]. Available: <http://www.cossan.co.uk/> (accessed January 11, 2016).
- [30] E. Patelli, M. Broggi, M. de Angelis, M. Beer, OpenCossan: An Efficient Open Tool for Dealing with Epistemic and Aleatory Uncertainties, in: *Second Int. Conf. Vulnerability Risk Anal. Manag.*, Liverpool, UK, 2014.
- [31] G.M. Laudone, C.M. Gribble, G.P. Matthews, Characterisation of the porous structure of Gilsocarbon graphite using pycnometry, cyclic porosimetry and void-network modeling, *Carbon N. Y.* 2014. 73 pp. 61–70.

- [32] C.N. Morrison, A.P. Jivkov, Y. Vertyagina, T.J. Marrow, Multi-scale modelling of nuclear graphite tensile strength using the site-bond lattice model, *Carbon* N. Y. 2016. 100 pp. 273–282.
- [33] New cracks in Hunterston reactor - BBC News [Online]. Available: <http://www.bbc.co.uk/news/science-environment-29481481> (accessed January 11, 2016).
- [34] J.D. Arregui-Mena, L. Margetts, D.V. Griffiths, L. Lever, G. Hall, P.M. Mummery, Spatial variability in the coefficient of thermal expansion induces pre-service stresses in computer models of virgin Gilsocarbon bricks, *J. Nucl. Mater.* 2015. 465 pp. 793–804.
- [35] G.A. Fenton, D. V. Griffiths, *Risk Assessment in Geotechnical Engineering*, John Wiley & Sons, Inc., Hoboken, NJ, USA, 2008.
- [36] I.M. Smith, D. V. Griffiths, L. Margetts, *Programming the Finite Element Method*, 5th ed., Wiley, Chichester, 2013.
- [37] I.M. Smith, L. Margetts, The convergence variability of parallel iterative solvers, *Eng. Comput.* 2006. 23 (2) pp. 154–165.
- [38] L. Margetts, *Parallel Finite Element Analysis*, PhD Thesis, University of Manchester, 2002.
- [39] ParaFEM [Online]. Available: <http://parafem.org.uk/> (accessed January 11, 2016).
- [40] D. Raabe, Cellular automata in materials science with particular reference to recrystallization simulation, *Annu. Rev. Mater. Res.* 2002. 32 (1) pp. 53–76.
- [41] A. Shterenlikht, L. Margetts, Three-dimensional cellular automata modelling of cleavage propagation across crystal boundaries in polycrystalline microstructures, *Proc. R. Soc. A Math. Phys. Eng. Sci.* 2015. 471 (2177) pp. 20150039–20150039.
- [42] M. Xiaying, H. Hurang, A finite element elastic-plastic-creep analysis of materials with temperature dependent properties, *Comput. Struct.* 1988. 30 (4) pp. 953–956.
- [43] L. Saucedo Mora, M. Mostafavi, T.J. Marrow, D. Khoshkhou, B. Connolly, C. Reinhard, R. Atwood, S. Zhao, 3D cellular automata finite element (CAFE) modelling and experimental observation of damage in quasi-brittle nuclear materials: Indentation of a SiC-SiC-fibre ceramic matrix composite, in: *Struct. Mater. Innov. Nucl. Syst.*, Idaho National Laboratory, Idaho Falls, United States, 2015: pp. 243–250.
- [44] J.K. Mackenzie, Second paper on statistics associated with the random disorientation of cubes, *Biometrika.* 1958. 45 (1-2) pp. 229–240.
- [45] J. Phillips, A. Shterenlikht, M.J. Pavier, Cellular automata modelling of nano-crystalline instability, in: *Proc. 20th UK ACME Conf.*, Manchester, UK, 2012.
- [46] CGPACK on SourceForge.net [Online]. Available: <http://sourceforge.net/projects/cgpack/> (accessed January 11, 2016).
- [47] E. Coleri, J.T. Harvey, K. Yang, J.M. Boone, Development of a micromechanical finite element model from computed tomography images for shear modulus simulation of asphalt mixtures, *Constr. Build. Mater.* 2012. 30 pp. 783–793.
- [48] S.A. McDonald, G. Dedreuil-Monet, Y.T. Yao, A. Alderson, P.J. Withers, In situ 3D X-ray microtomography study comparing auxetic and non-auxetic polymeric foams under tension, *Phys. Status Solidi.* 2011. 248 (1) pp. 45–51.

- [49] Ll.M. Evans, L. Margetts, V. Casalegno, F. Leonard, T. Lowe, P.D. Lee, M. Schmidt, P.M. Mummery, Thermal characterisation of ceramic/metal joining techniques for fusion applications using X-ray tomography, *Fusion Eng. Des.* 2014. 89 pp. 826–836.
- [50] D. Gonzalez, A. King, M. Mostafavi, P. Reischig, S. Rolland du Roscoat, W. Ludwig, J. Quinta da Fonseca, P.J. Withers, T.J. Marrow, Three-dimensional observation and image-based modelling of thermal strains in polycrystalline alumina, *Acta Mater.* 2013. 61 (20) pp. 7521–7533.
- [51] P.G. Young, T.B.H. Beresford-West, S.R.L. Coward, B. Notarberardino, B. Walker, A. Abdul-Aziz, An efficient approach to converting three-dimensional image data into highly accurate computational models., *Philos. Trans. A. Math. Phys. Eng. Sci.* 2008. 366 (1878) pp. 3155–73.
- [52] A.M. Yamileva, A.V. Yuldashev, I.S. Nasibullayev, Comparison of the Parallelization Efficiency of a Thermo-Structural Problem Simulated in SIMULIA Abaqus and ANSYS Mechanical, *J. Eng. Sci. Technol. Rev.* 2012. 5 (3) pp. 39–43.
- [53] R. Tivey, T. Ando, A. Antipenkov, V. Barabash, S. Chiocchio, G. Federici, C. Ibbott, R. Jakeman, G. Janeschitz, R. Raffray, M. Akiba, I. Mazul, H. Pacher, M. Ulrickson, G. Vieider, ITER divertor, design issues and research and development, *Fusion Eng. Des.* 1999. 46 (2-4) pp. 207–220.
- [54] R.D. Watson, F.M. Hosking, M.F. Smith, C.D. Croessmann, Development and Testing of the ITER Divertor Monoblock Braze Design, *Fusion Sci. Technol.* 1991. 19 (3P2B) pp. 1794–1798.
- [55] Ll.M. Evans, L. Margetts, V. Casalegno, L. Lever, J. Bushell, T. Lowe, A. Wallwork, P. Young, A. Lindemann, M. Schmidt, P. Mummery, Transient thermal finite element analysis of CFC–Cu ITER monoblock using X-ray tomography data, *Fusion Eng. Des.* 2015. 100 pp. 100–111. doi:10.1016/j.fusengdes.2015.04.048.
- [56] MXIF [Online]. Available: <http://www.mxif.manchester.ac.uk/> (accessed January 11, 2016).
- [57] 3D image data visualisation, analysis and model generation with Simpleware software [Online]. Available: <http://www.simpleware.com/> (accessed January 11, 2016).
- [58] Ll.M. Evans, L. Margetts, J. Bushell, T. Lowe, A. Wallwork, W.E. Windes, P. Young, P.M. Mummery, Parallel processing for time-dependent heat flow problems, in: *NAFEMS World Congr.*, Salzburg, Austria, 2013.
- [59] ParaView - Open Source Scientific Visualization [Online]. Available: <http://www.paraview.org/> (accessed January 11, 2016).
- [60] J. West, S. Gallagher, Challenges of open innovation: the paradox of firm investment in open-source software, *R D Manag.* 2006. 36 (3) pp. 319–331.
- [61] Ll.M. Evans, Thermal Modelling of Composite Materials for Fusion Applications; Finite Element Analysis from X-ray Tomographic Images, PhD Thesis, University of Manchester, 2013.
- [62] R. Kurihara, S. Ueda, S. Nishio, Y. Seki, Fracture mechanics evaluation of a crack generated in SiC/SiC composite first wall, *Fusion Eng. Des.* 2001. 54 (3-4) pp. 465–471.