UK Atomic Energy Authority

ExCALIBUR

Assessment of which UQ methods are required to make NEPTUNE software actionable

M2.4.1

Abstract
The report describes work for ExCALIBUR project NEPTUNE at Milestone 2.4.1. Minutes of meeting to form report on technical progress.

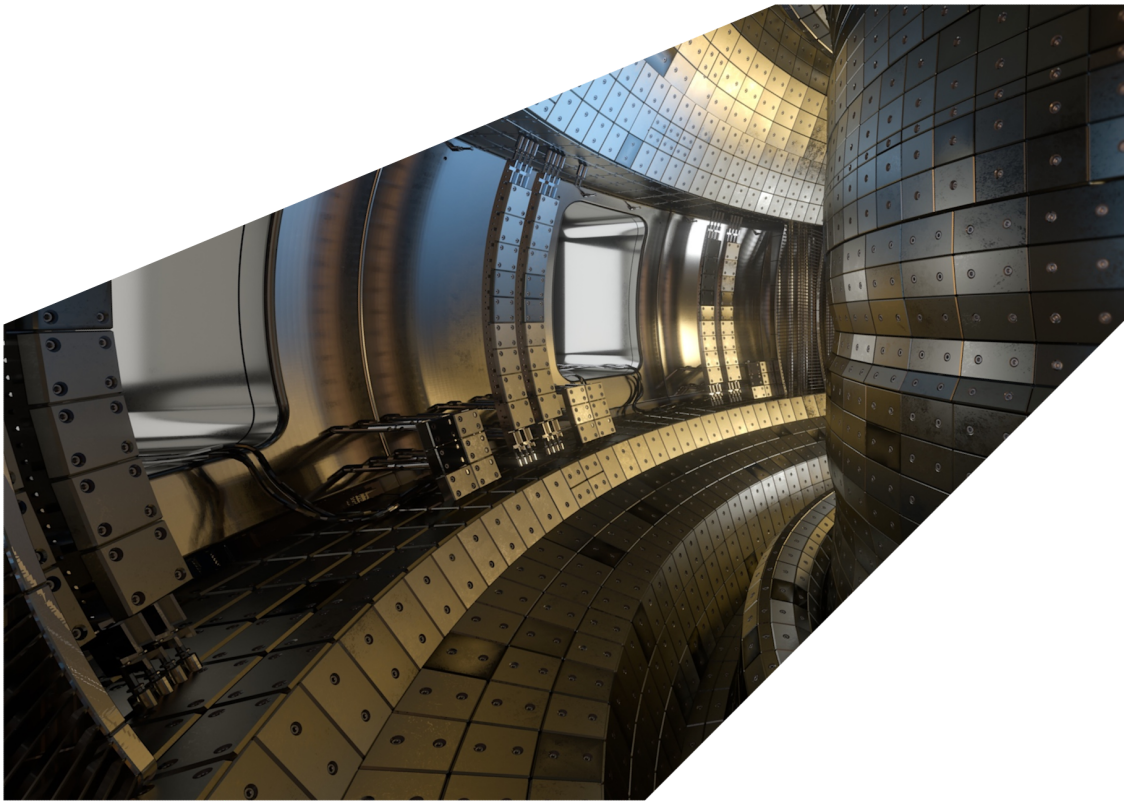

UKAEA REFERENCE AND APPROVAL SHEET

	Client Reference:	
	UKAEA Reference:	CD/EXCALIBUR-FMS/0035
	Issue:	1.00
	Date:	April 28, 2021
Project Name: ExCALIBUR Fusion Modelling System		

	Name and Department	Signature	Date
Prepared By:	Ed Threlfall Wayne Arter BD	N/A N/A	April 28, 2021 April 28, 2021
Reviewed By:	Rob Akers Advanced Computing Dept. Manager		April 28, 2021
Approved By:	Rob Akers Advanced Computing Dept. Manager		April 28, 2021

1 NEPTUNE Meeting: 20 April 2021 10-11am BST

Present

- Chair: Wayne Arter, UKAEA
- Maxime Vassaux, UCL
- Ed Threlfall, UKAEA
- Wouter Edeling, CWI Amsterdama (apologies) *Minutes augmented with information provided by WE who was unable to be present*

2 Minutes

Note that MV has obtained recently a postdoc position at Rennes University and wishes to stay in France; he remains involved in NEPTUNE for the coming six months but his future on the project afterwards is uncertain.

WA began with a discussion of the recent report [1] by MV; he commended the report, particularly as MV is not a plasma physics expert (his main interest is materials and molecular dynamics, though he made the point that the UQ is somewhat orthogonal to the actual physics of the problem in question). WE has a background in UQ, mainly applied to (RANS) turbulence models, and the sub-gridscale parameterizations of simplified ocean models.

WA asked for an example of an intrusive UQ method; MV quoted polynomial chaos PC and the polynomial chaos expansion (PCE) and added that they did not favour these methods for NEPTUNE since they require the code in question to be modified (and presumably recompiled) for each change to the UQ campaign (eg. number of parameters), necessitating a lot of human work (Edeling had helped to convince MV of this). WE indicated that PC need not be intrusive (as discussed elsewhere see the report by Arter et al [2]) but that intrusive PC did have the advantage that the equations might only involve a relatively small number of coefficients of the PCE. Non-intrusive UQ uses existing libraries eg. *EasySurrogate* of VECMATK, or *MoGP* for as used by Peter Challenor's Exeter group and the Alan Turing Institute (Serge Guillas). Both the aforementioned use Gaussian Process surrogates and are more versatile and re-useable than intrusive methods, and still relevant for coupled models via semi-intrusive methods. WA mentioned that methods that were semi-intrusive as defined in the VECMA project allow for different scales and that the main problem with plasma physics was the presence of many coupled scales (eg. turbulence has much shorter timescale than the equilibrium solver and transport code). MV mentioned that data flows both ways between different scales in coupled models and that semi-intrusive UQ works with this. WA raised the issue of scheduling in this type of coupled simulation - difficult to optimize throughput. MV stated that this problem was inherent to multi-physics, not introduced by UQ but obviously exacerbated since UQ demands multiple simulations. VECMATK *QCG_PilotJob* can help as it does efficient job ordering / submission for HPC workflows. The development of the QCG software takes place in Poland, as part of VECMA.

WA started a discussion of enhanced sampling methods, asking if they were adaptive. MV answered in the affirmative; he is mostly familiar with stochastic collocation. Edeling is involved with active subspaces. WA asked if MV could describe the latter, but not possible as MV not the expert there, though he mentioned that this method was of great interest to Serge Guillas.

Active Subspaces

A description of the method provided by WE is as follows: "The original active subspace method was developed by Paul Constantine, and this method performs dimension reduction by finding directions in the input space along which the solution varies the most. This is done by computing a positive semi definite matrix of gradient samples (averaged over the input distribution), followed by an eigenvalue decomposition. The eigenvectors are orthogonal, and denote a rotated coordinate system in the input space. Dimension reduction is performed by keeping only the eigenvectors associated to large eigenvalues W, and using these to project the high-dimensional input vectors $\mathbf{y}$ to a lower dimensional one: $\mathbf{z} = W^T\mathbf{y}$. So instead of trying to figure out which input is important, this method tries to find a linear combination of input variables that is responsible for most variation in the output. The downside of this method is that you need to be able to compute gradients of your code output w.r.t. the inputs. Serge Guillas has one adaptation of the active subspaces method that does not require this, using Gaussian Processes. Another option that does not require code gradients is to use Deep Active Subspaces (developed by Tripathy), which uses neural networks. WE has spent two weeks looking into this method, and has small report and preliminary EasySurrogate implementation available including application to *CovidSim*. It was worth noting that Deep Active Subspaces was not inherently an iterative method."

WA asked about methods to treat the curse of dimensionality; MV mentioned adaptive methods used in *CovidSim* and molecular dynamics to find what were the main parameters driving variation in model outputs. WA asked whether the methods were local gradient-type analyses but it became clear that the analysis involved simulating over a distribution of inputs and then eg. finding the Sobol indices determining from which input the output variation stems. Stochastic collocation methods here use a regular grid capable of selective refinement according to input parameter (so anisotropically-refined in input parameter space); without refinement is like Latin hypercube sampling. WA mentioned that these methods might be challenged by a strange parameter distribution eg. a bimodal one; MV said this would just demand a large variance in the input sampling distribution. WA asked about resampling methods, which it turned out were one of WE's contributions to the report [1].

WE subsequently explained that he had developed stochastic resampling methods for turbulence-like problems. They used the stochastic surrogate in a two-way coupled simulation for a very simple atmospheric model (Lorenz 96), keeping the macroscopic part of the original model intact, and coupled it with a stochastic, neural-network based surrogate that resamples the reference data, conditioned on the macroscopic input features. Basically, conditional bootstrapping using neural networks. Although this approach works very well for Lorenz 96, for instance the work of Stephan Rasp indicates that the fidelity (or even the stablity) of the coupled physics-ML model is not guaranteed for larger models. Rasp suggests performing online learning, in which the ML model is trained to behave in a coupled system, rather than training the ML model just to represent (smallscale) data well, which WE is investigating.

WA asked whether MV recalled ref. 4, a (sizeable) fusion paper by C. Holland of the Californian

DIII-D group: *Validation metrics for turbulent plasma transport*, as UKAEA's main interest lies in whether our numerics capture correctly the features of turbulence; the paper distinguishes identifying experiment / numerical simulation.

A discussion of QoIs followed; MV agreed these must be defined at the start of a UQ campaign. Recent work by BOUT++ community members identified QoIs as well as input parameters and their uncertainty. MV says this sort of this is very specific to the physics problem in question. (One thought that ET has, not mentioned in the meeting, is that there may be two particular classes of QoI: averaged quantities, and more challenging extreme events.) WA added that there are two customers for the end-product software: engineers designing the first wall, who want to know the power deposition everywhere on the surface, and plasma physicists who may be more concerned with local details of the physics (eg. to construct physics-based surrogates for turbulence, as small-scale turbulence challenges computing power). Both classes of user need accurate results.

WA asked if anything else really stood out for MV in the context of plasma applications, mentioning that he (WA) has a student using the COSSAN toolkit (based on MATLAB). MV said that the advantage of VECMATK would be its integration with HPC workflows. MV mentioned that VECMATK had not in itself been verified (eg. vs. COSSAN) as it is based on standard techniques. Other toolkits are mentioned in the report eg. DAKOTA.

WA asked what NEPTUNE work was currently in progress (ie. post the report). Edeling is focussing on active subspaces and MV on semi-intrusive methods - currently in the context of his materials work - and also setting up the imminent (21-23 April 2021) VECMA hackathon. WA mentioned that the lack of a particle code was a problem for the hackathon; particles have been added to BOUT++ but are currently not working; as to the reason for this failure WA speculated that the problem is simply too hard for a postdoc to solve in a realistic timescale (MV sympathizes). WA mentioned some of our interest in stochastic methods arises from the fact that the particles code will give an effective stochastic source of mass and momentum. MV has done similar work in the molecular dynamics context where the question of the degree of stochasticity between coupled models is of interest. WA agreed the link with molecular dynamics is interesting and mentioned that recent UKAEA recruit Will Saunders' PhD was in this area. WA asked for opinion on how to mesh in the regions containing particles / interfaces with the fluid code; MV replied that the answer depends on the degree of coupling: in *weak coupling*, each model can be run separately. WA mentioned that an heuristic in *strong coupling* is the need to overlap computational domains of particle and fluid code - one cannot simply couple across an interface - and asks why this is. MV answered that in molecular dynamics, the energy transfer proves incorrect unless regions overlap, hence the last decade of developing this type of coupling technology, though the size of the overlap region can be small. WA added that the size of overlap regions was an interesting question. MV mentioned that he personally does mostly weak coupling as strong coupling is too computationally-expensive. He asked what subset of the space is particles vs. fluids in a tokamak simulation; WA answered that the problem starts with the sheath region near the machine walls where the cheaper fluid approximation is not valid, which spreads into a significant volume fraction. The issue with trying to use weak coupling is that the sheath region sets the boundary condition for the plasma, and if collisionality is low then the boundary effects can affect the bulk plasma. Complicated question hence our great need for UQ techniques.

WA concluded by asking if MV had further questions; MV reminded that UKAEA may wish to contribute a 15-min NEPTUNE use-case presentation to the all-hands VECMA meeting in mid-

May. WA already aware. MV noted significant overlap with interests of the UQ and NEPTUNE communities and recommended future correspondence re. coupling.

Acknowledgement

The support of the UK Meteorological Office and Strategic Priorities Fund is acknowledged.

References

[1] M. Vassaux, W. Edeling, and P.V. Coveney. Review of methods and toolkits for uncertainty quantification of single and coupled model applications. Technical Report 2047352_1-TN-01, UKAEA Project Neptune, 2021.

[2] W. Arter, E. Threlfall, and J. Parker. Report on user layer design for Uncertainty Quantification. Technical Report CD/EXCALIBUR-FMS/0024-M3.1.3, UKAEA, 2020.