UK Atomic Energy Authority

ExCALIBUR

NEPTUNE: Year 1 Summary

M1.2.1

Abstract

This report describes work for ExCALIBUR project NEPTUNE at Milestone 1.2.1. This End of Year One (Y1) summary report outlines the reasoning behind the Y2 plan, listing questions to be addressed (especially by 3rd parties), the proxyapps that will be developed, and how they will be connected to each other and exploited to steer thinking around the development of a next generation edge boundary code.

UKAEA REFERENCE AND APPROVAL SHEET

	Client Reference:	
	UKAEA Reference:	CD/EXCALIBUR-FMS/0011
	Issue:	1.00
	Date:	10 March 2020
Project Name: ExCALIBUR Fusion Modelling System. Version 1.00		

	Name and Department	Signature	Date
Prepared By:	Wayne Arter Lucian Anton Debasmita Samaddar	N/A N/A N/A	10/03/2020 10/03/2020 10/03/2020
Reviewed By:	Rob Akers	N/A	10/03/2020
Modified By:	Rob Akers MSSC		10/03/2020
Approved By:	Rob Akers MSSC		12/03/2020

Chapter 1

Introduction

This section begins with a subsection Section **Error! Reference source not found.** which places the NEPTUNE project in the global context. The following subsections Section **Error! Reference source not found.** and Section **Error! Reference source not found.** review more local UKAEA activities prior to commencement of NEPTUNE, and its concluding subsection brings history up-to-date with a summary of NEPTUNE Y1 activities. Much of the material in this Section 1 bears on the choices made for the Y2 programme discussed in Section 2 and summarised in Section 3.

1.1 Background

In the past few years, the focus of the UKAEA fusion programme has been increasingly directed towards the design of power-generating reactors, see e.g. papers presented at a Royal Sociey meeting in 2018 “Fusion energy using tokamaks: can development be accelerated?’, notably refs [1, 2]. At the start of November 2018, two documents of worldwide importance for the development of fusion appeared, namely the ITER Research Plan [3] and the updated European Research Roadmap [4]. They correctly emphasise the importance of modelling *in-silico* but are lacking in detail as to how accurate and reliable models are to be developed in a timely fashion for the designers and operators of the ITER and DEMO devices. For example there is lack of usage of the word “software", implying absence of the term “software engineering" (SWEng). However, the increased stress on model reliability for “digital twinning" might a priori be expected to require careful management of software throughout its entire life-cycle.

There are special features for the software engineer to account for in a leading-edge technology programme such as fusion reactor development, which is nonetheless long-term. Digitally twinned software must be flexible, to deal with issues which perhaps were not initially regarded as important and may require rapid resolution, for example in the light of new experimental data or during construction work. This amounts to a requirement that the software, based upon current fusion timelines needs to remain not only usable, but also modifiable, perhaps drastically, over at least the next 30 years needed to see DEMO into power production.

Concerning existing software, especially for plasma physics modelling, it can be argued that the finite difference software (fd) presently underlying much design work is obsolescent in the above special software engineering terms, being expensive and/or error-prone to maintain as the products of an era when SWEng was in its infancy. Specifically here is being considered the

Patankar scheme [5] that underlies the SOLPS fluid code known as B2 [6]. SOLPS is for many people the ‘gold standard’ edge modelling code because it incorporates more physical effects than more recent developments. From a practical point-of-view however, fd is unsatisfactory for design work because it fails to give robust error estimates, and inefficient because of 21st Century trends in computer architecture. Typically these schemes are not very accurate with errors of 1 % rising locally to 10 %. Recent developments of the *chebfun* approach by Trefethen and co-workers (eg. ref [**Error! Reference source not found.**, **Error! Reference source not found.**] are producing commodity software capable of solving PDEs to ‘spectral’ accuracy in double precision. Sustained progress here will render these plasma physics low-order-accurate fd schemes an embarrassment. There is perhaps an even more significant issue for the use of the particle-based methods described in say the review [7], since these are associated with relatively large amplitude sampling ‘noise’.

1.2 Spectral/hp Element Method

Two student projects involving the spectral/hp element method for modelling the tokamak edge were commissioned by UKAEA from 2015 onwards, see ref [8, 9], and a number of other interactions took place, notably presentations by Sherwin the joint author of the textbook [10], and in 2018 by both Sherwin and his former post-doc Moxey. Just prior to NEPTUNE commencement, Arter [11] gave a presentation at the Isaac Newton Institute explaining the desirability of using spectral elements for tokamak edge modelling.

Spectral element methods [12] are capable of providing usable error estimates and of exploiting the developing architecture to ensure maximum use of complex 2-D, 3-D and time-dependent physical models expected to be more robust under extrapolation to new situations and regimes. Their extreme accuracy much reduces the need for special physics-dependent coding, suggesting that a near-complete replacement of code essential for reactor design is feasible, and if carefully managed, such a new development can meet the demanding SWEng requirements of the fusion programme.

Amongst spectral element approaches, none has the maturity and the full range of advantages presented by spectral/hp elements. The spectral/hp element method (see 650-page textbook [10]) was developed jointly by Karniadakis and Sherwin around the turn of the Century. In the name, the “h" implies variable element size, and “p" implies use of a p^{th}-order accurate approximation, where p is an arbitrarily large order. “Spectral" reinforces the notion that as $p \rightarrow \infty$, spectral accuracy, ie. error smaller than any power of element size, is achievable. This is to be compared with commercial packages for comparable problems that use low order finite elements, often restricted to first or second order. The name spectral/hp does not exclude use

also of *r*-elements, ie. additional elements introduced dynamically to represent evolving features of say a fluid flow-field.

Figure 1.1 indicates the increasing application of the method in published papers to a wide range of problems. In addition, there has been commercially driven application of spectral/hp elements to the design of F1 cars. Other 21st Century developments have seen the the incorporation of special, geometrically accurate meshing techniques [13] into the NekMesh library, which overcome the often large errors associated with common triangulations of geometry [14], and production of the Nektar++ library [15]. Further investigation of geometry representation for NEPTUNE application is to be found in the Milestone Report [16]. Figure 1.1 indicates that the latter has facilitated widespread uptake of the method, giving evidence for a growing community with a large range of applications, many of direct interest to fusion.

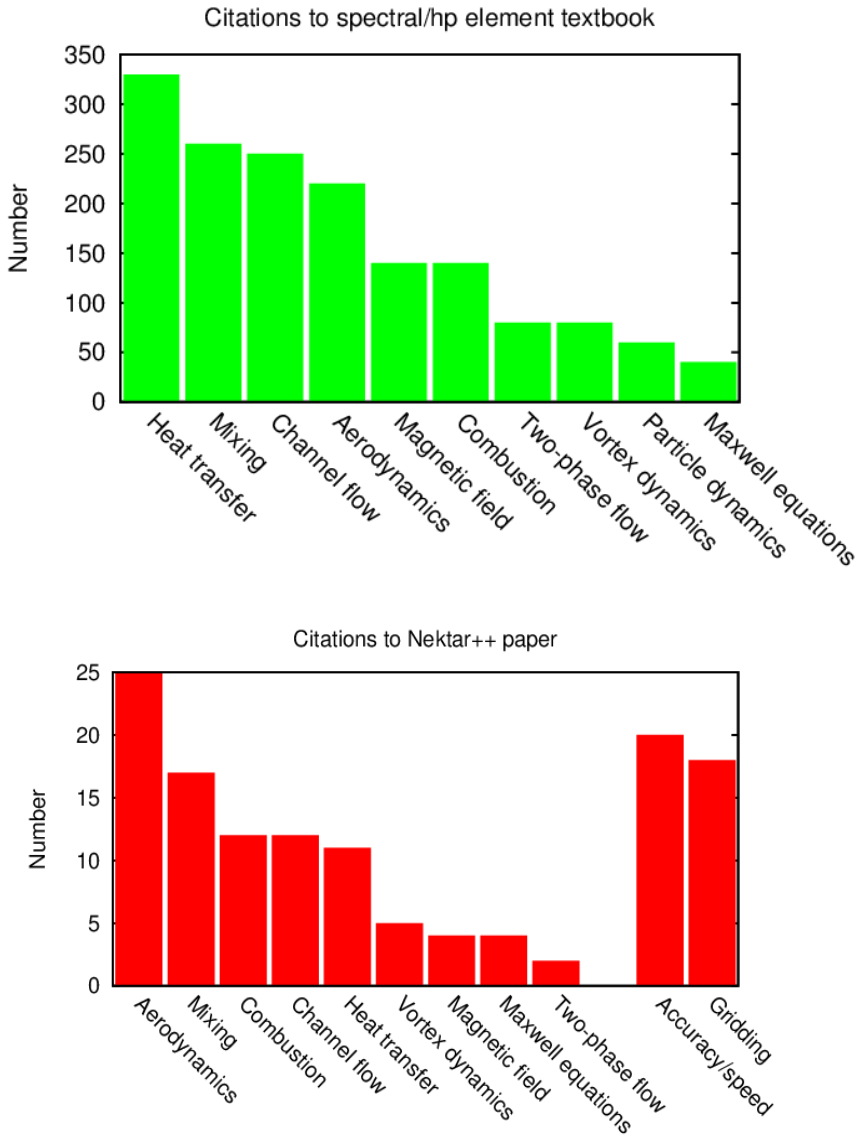

Figure 1.1: Citation counts for applications of the spectral/hp element textbook and Nektar++ library. Search performed on google scholar(1/5/18) for keywords indicated. Spectral/hp element textbook from

1997/2003 has 2300 citations, Nektar++ library description paper from July 2015 has 120 citations (151 as of 1/11/18).

The controllable accuracy of spectral element schemes is important despite the inevitable errors in any experimental confrontation with reality, because of great advances in techniques for fitting models to data, in fields such as optimisation, uncertainty quantification and data assimilation, now often bracketed as Machine Learning - see recent work by Arter and students [17, 18]. These advances in VVUQ techniques give more robust fitting with improved error estimates.

One advantage of extreme accuracy that is highlighted here, because of its relevance to the serious problem of heat deposition to the first wall, occurs when there is very anisotropic heat transport, found in some layered materials or as here, because of a large-scale magnetic field. If the numerical grid is not aligned with the field, then heat leaks in the normal direction at every timestep at a rate proportional to cell-size h, leading to an effective heat diffusivity $D_{\perp,num} \propto h^2/\Delta t$ (Δt is time-step size) which may greatly exceed the intended perpendicular transport.

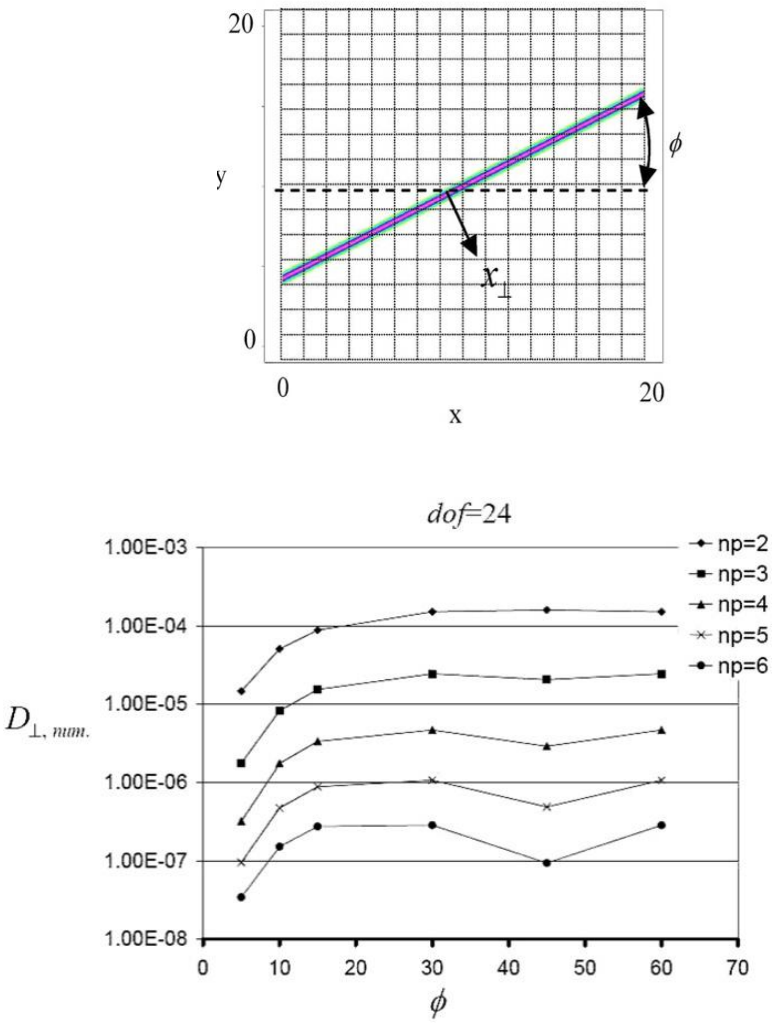

Figure 1.2: Spectral element methods can model the entire range of misalignment angles φ without any

special treatment, from ref [19]. Figure at top indicates the field direction relative to the fe mesh, in the graph below np corresponds to the order p and dof to the number of degrees of freedom per unit length.

Even if the grid is field-aligned, either by special local construction or by use of an unstructured, finite element representation, since a good first wall power distribution requires almost tangential incidence, there will be elements with very small angles leading to poor convergence properties for most solvers. However, as Figure 1.2 shows, for spectral elements, the effective relative heat diffusivity remains small over the whole range of φ, even having a local minimum at the angle of maximum misalignment.

1.3 High Level Integration

Work in [20] examined how by use of appropriate software technology, spectral element software could be integrated into an object-oriented approach permitting easy integration into the overall reactor design package. This is a two-way process, in that design can only proceed at an acceptable speed using surrogates for full 3-D to 6-D models even on multi-core desktops, whereas as in critical cases, 3-D to 6-D models may need to be invoked to reduce uncertainty. Preliminary indications were that this should be achievable using a graph-based approach, implemented in matrix form - see Figure 1.3.

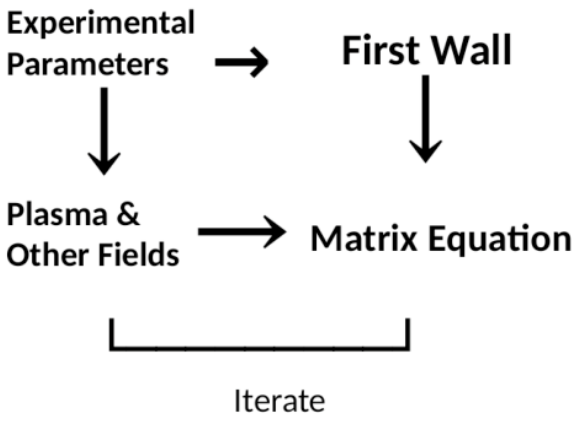

Figure 1.3: Schematic of top-level algorithm.

A bundled capability for producing surrogates needs to be included, and emphasis needs to be placed upon generating well-documented arms-length interfaces that can be used by external packages. The programming style described in [21] is particularly appropriate to this purpose, and should be applicable in any object-oriented language. There will also be a need for an easy

way to specify scalar and vector differentiation operators.

1.4 UKAEA Activities to date

Work to-date on the NEPTUNE project has been driven by the above considerations. Of all the tokamak modelling codes, SOLPS is the one of which development has been most weighed down by its history, making edge modelling the prime candidate area for producing new software. The EUROfusion Theory and Advanced Simulation Coordination (E-TASC) programme has also recognised this, in that two of the first TSVV (Theory, Simulation, Verification, and Validation) tasks are to pilot writing of the European Boundary Code (EBC), and to launch a major upgrade of the SOLPS particle code known as EIRENE. These pilot tasks are due to finish at the end of 2020, and particularly the EBC project has a challenging target to produce significant novel physics results, in order to secure funding for the proposed 5-year follow-up.

It has been decided that NEPTUNE will coordinate primarily with the TSVV EBC project (which necessarily has to coordinate with the particle code project). UKAEA staff working on NEPTUNE have attended the two major meetings of the EBC pilot phase, and relevant outcomes are discussed in the Milestone Report [22]. As pointed out at the 2019 internal meeting by UKAEA staff, some decisions taken to ensure the EBC targets are met are sub-optimal for a longer term project such as NEPTUNE , never mind a thirty-year horizon.

NEPTUNE also has to meet other UK SPF aims, based around the four pillars. However, it seems NEPTUNE can coordinate with EBC to the extent that it looks to examine a common 2-D fluid model in Y3, so that where NEPTUNE Y2 research tasks are directed towards this same model, NEPTUNE should be able to assist EBC with any problems encountered related to those tasks. Most other research, such as interface design, integration with particle methods etc. to be performed under NEPTUNE should be complementary to the EBC pilot phase development, and of immense value to a 5-year follow-on programme.

Other UKAEA activities are based around findings from the 2019 internal workshop for requirements capture and an external, UK wide workshop held in Birmingham on 05-02-2020; these are reported upon in ref [23]. The above activities have been accompanied by and to some extent driven by identification and preliminary investigation of the critical research needed into the physical model, numerical algorithms and software engineering through extensive community engagements and literature review.

Further investigation of particle methods for NEPTUNE application is to be found in the

Milestone Report [24].

Chapter 2

Task Work

Subsection Section 2.1 represents ‘outreach’ to the Computational Fluid Dynamics (CFD) community, to give them ‘some feel’ for the basic physical issues at stake, similarly Section **Error! Reference source not found.** tries to illustrate the sheer power of Exascale. The following subsection Section 2.3 discusses approaches that were raised during Y1 by a range of potential collaborators, and the last subsection Section 2.4 elaborates a potential addition to the NEPTUNE tasks.

2.1 SOL and PFC as convection problem

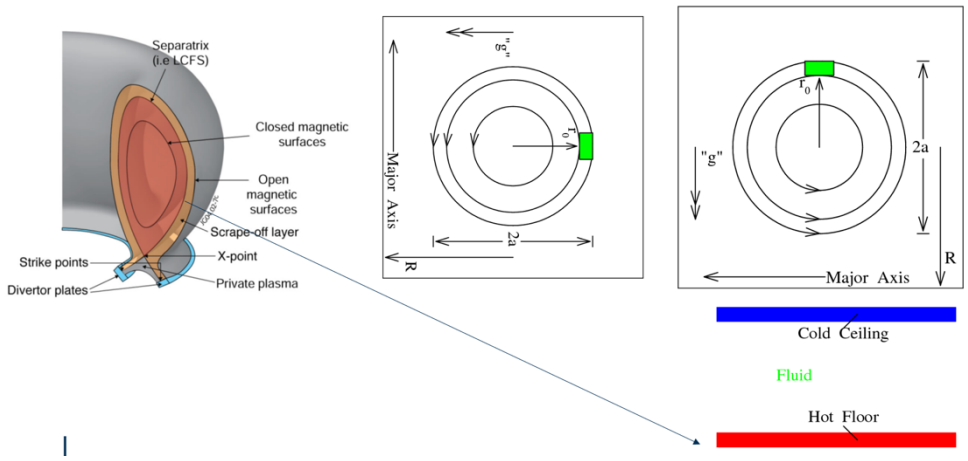

Figure 2.1: Relationship between the tokamak edge and classical fluid dynamics. The plot at far left is courtesy of EUROfusion (eurofusion.org).

To give the specialists in Computational Fluid Dynamics (CFD) a feel for the problem, Figure 2.1 shows how the outer part of the midplane (green box) can be approximated as a classical Rayleigh-Benard convection (RBC) problem - see the recent work [25]. Plasma wants to flow along fieldlines, but there is no escape from the effects of momentum conservation, and the net result is that plasma experiences an inward acceleration *g*. Rotating the picture through 90^{o} so this acceleration looks vertical makes clearer the relation. The hot central plasma represents a heated floor capable of driving convection in the SOL. For RBC, Reynolds numbers and boundary layer thicknesses have been fitted by Grossmann & Lohse [26]. Estimates of the

dimensionless pressure/temperature difference or Rayleigh number *Ra* range in the tokamak edge [25] plausibly from 10^5–10^{12}. Assuming a Prandtl number *Pr*≈1, where *Pr* is the ratio of viscous to thermal diffusion, the Reynolds *Re* number varies from *Re*=10−4×10^4. Such values imply relatively narrow boundary layers, thickness∝1/ $\sqrt{Re}$ and turbulence for most of the range.

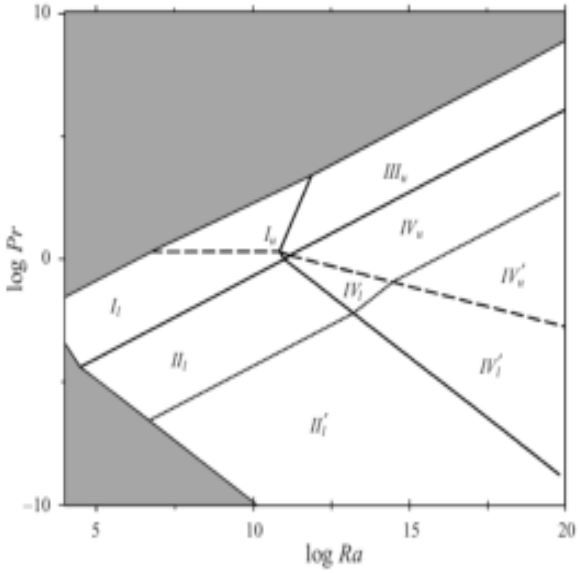

Figure 2.2: Ten regimes of RBC depending on the parameters Ra *and* Pr*. Reproduction of Fig. 2 from ref [26].*

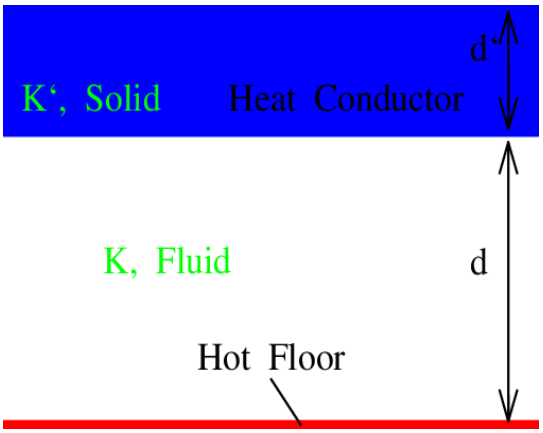

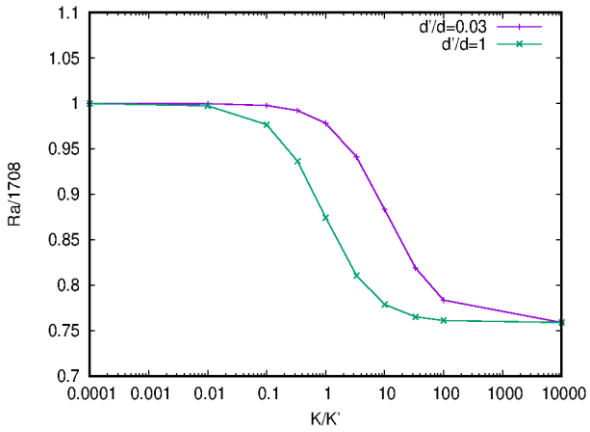

Figure 2.3: Sketch at top of the so-called Rayleigh-Jeffreys problem, where a solid of thermal conductivity K' *overlies a fluid with thermal conductivity* K. *Below is a plot of normalised critical* Ra *as the ratio of conductivities changes. Data from ref [27].*

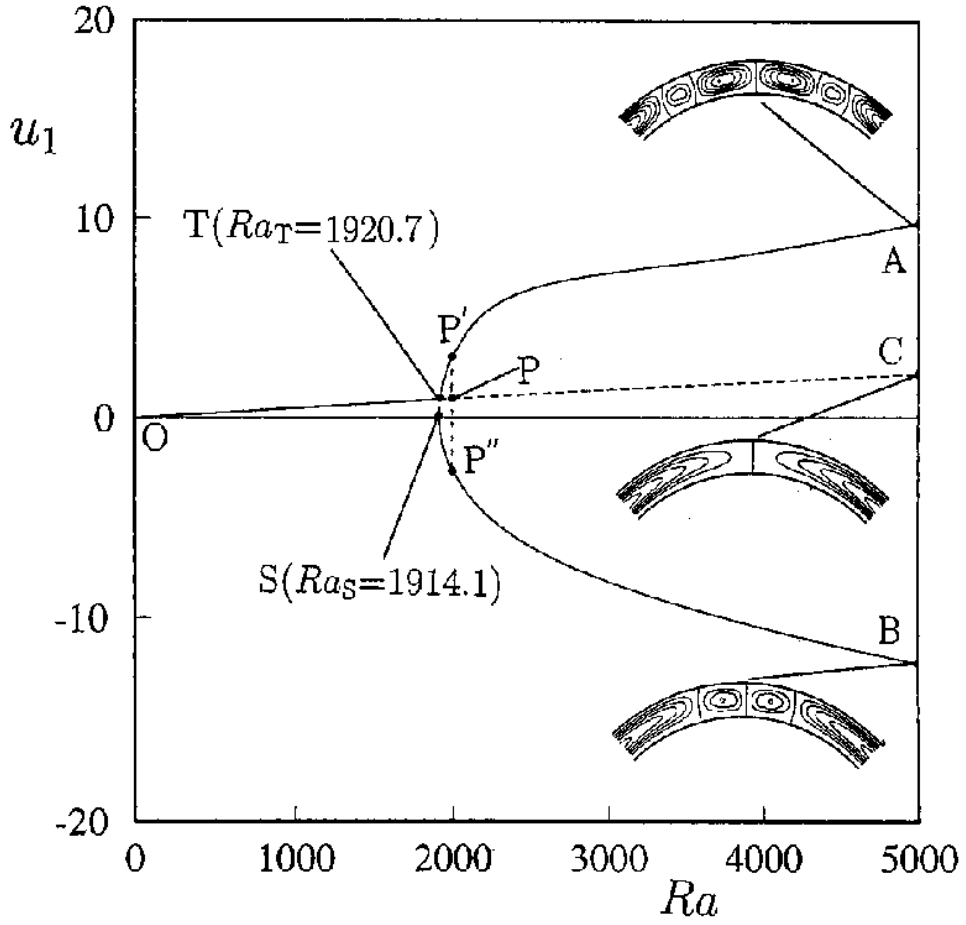

Figure 2.4: Amplitude plotted against Ra *at fixed* Pr *showing the interaction between the 'double-glazing' solution branch which becomes unstable at small* Ra *and the classical RBC branch which appears beyond* Ra*≈1900. Reproduction of Fig. 2 from ref [28].*

It is worth remembering that the Rayleigh-Benard problem, while simple to pose, is hard to understand fully, and sensitive to details of both boundary conditions and geometry. For example:

1. For the scaling of dimensionless heat flux Nu with convection, there are ten regimes of RBC depending on the temperature difference measured by Rayleigh number Ra and Prandtl number $Pr=Ku/K$, see Figure 2.2. The scaling law exponents γ_1, γ_2 which are defined so that $Nu \propto Ra^{\gamma_1} Pr^{\gamma_2}$ range over $1/4<\gamma_1<4/7$ and $-5/6<\gamma_2<1/8$.

2. There is an approx. 70 % reduction in heat flux if the boundary condition on the flow is changed from free-slip to stagnation. This follows because the critical Ra, $Ra_C \approx 657$ increases to $Ra_C \approx 1708$ [29], and from assuming that heat flux $Nu \propto (Ra - Ra_C)^{1/2}$.

3. There is an approx. 10 % increase in Nu possible due to a change outside the fluid region, namely the replacement of perfectly conducting boundaries by poor insulators, see Figure 2.3, understood from the fact that a poorly conducting boundary constrains the convection less.

4. In a heated annulus, which Figure 2.1 suggests is a more accurate model of the edge, there is a complex bifurcation structure possible, see Figure 2.4.

5. There are indications that spectral (and probably other numerical) schemes converge non-monotonically as spatial resolution increases, e.g. the reduction of RBC to the Lorenz equations [30] does not permit smaller-scale instability of thermal boundary layers because it represents them too crudely. However, more detailed, but unresolved models of these layers may exhibit time-dependence when the full solution with thinner layers does not, because layers of thickness d are destabilised by a term $\propto d^3$.

2.2 Capabilities of Exascale

In the tokamak edge, following [25], suppose plasma number density $n \approx 10^{18}\,\mathrm{m}^{-3}$. Order of magnitude dimensions are $L \approx 0.1\,\mathrm{m}$ for SOL thickness, reactor minor and major radii say $a = 3\,\mathrm{m}$ and $R_0 = 10\,\mathrm{m}$, so volume of SOL $\approx 4\pi^2 a R_0 L \approx 100\,\mathrm{m}^3$. Hence total number of electrons $\approx 10^{20}$.

Shortest timescale is inverse eB/m_e, the electron cyclotron frequency, where $e/m_e = 1.76 \times 10^{11}\,\mathrm{C\,kg}^{-1}$, and $B \approx 10\,\mathrm{T}$, so $\tau_{Be} \approx 10^{-12}\,\mathrm{s}$. Hence number of particle-steps to evolve 1 s of physical time is $10^{20+12+1}$, and assuming 1000 flop per update, we would need 10^{36} Flops. So to complete in 1 s on an Exascale machine, one could only sample 1 particle-step in 10^{18}. (Unlikely to be adequate because of electrostatic and other effects).

Suppose then that we adopt a fluid approach instead, ie. representing the electron distribution by the first 3 moments. Electron temperature $T_e \approx 10\,\mathrm{eV}$, thermal speed $V_{Te} \approx 10^6\,\mathrm{m\,s}^{-1}$. Sample SOL at uniform 1 mm interval, number of sample-points $\approx 10^{11}$, timestep for explicit scheme $\approx 0.1 \times (10^{-3}/10^6)$ so number of sample-point updates is 10^{11+10}; assuming 1000 flop each update, we would need 10^{24} Flops, 6 orders of magnitude greater than exascale.

Another way, suppose the numerical problem is D-dimensional; we would need say 1000 Flops

each sample update and allowed N_D samples per spatial dimension and N_D^2 in time. Then to update in 1 s, we have $N_D^{D+2} \approx 10^{15}$. Thus if $D=3$, $N_3 \approx 1\,000$, and $N_5 \approx 100$. ***In conclusion, accuracy controlled, unstructured, implicit fluid models should be possible.***

Similarly, estimates of the maximum frequency at which neutral atom particle-steps can be sampled, assuming $n_0=n$ and $T_0=T_e$, indicates that useful results may be obtained from the particle approach. Moreover, repeating the above calculation for electrons but considering that only plasma ions have to be simulated on a drift type timescale$\approx L/V_{Ti}$, $V_{Ti} \approx 3\times10^4$ m s^{-1} gives much more encouraging results.

2.3 Alternative Approaches

During the requirements capture exercise, several other possible approaches were mooted. Discounting special pleading, those felt worth of examining were

1. Discontinuous Galerkin (DG) with usual (polynomial) finite element bases [**Error! Reference source not found.**]
2. Isogeometric finite element bases
3. Lattice Boltzmann (LB)
4. Smoothed Particle Hydrodynamics (SPH)
5. Sparse grids [**Error! Reference source not found.**, **Error! Reference source not found.**]
6. Asymptotic-Preserving Methods
7. Dedalus software [**Error! Reference source not found.**]

The work of Nikiforakis and coworkers, eg. [31] may be regarded as special pleading, but since it also has implications for geometry representation, it is discussed in the Milestone Report [16].

2.4 Potential Additions to NEPTUNE

2.4.1 Adjoint or ‘dual’ approach

The adjoint approach to sensitivity analysis exploits the fact that the sensitivity of only one or perhaps a small number of quantities such as Q, each a linear combination of the N_i is needed. The technique is best illustrated [32] by a simple, linear example.

Imagine that the problem being solved is the (‘primal’) set of linear equations $A\mathbf{x}=\mathbf{b}$ to find one quantity (the ’objective function’) $\mathbf{g}^T\mathbf{x}$, where the vectors **b** and **g** are given. It happens that the

same result may be achieved by solving instead the ('dual') equation set $A^T\mathbf{v}=\mathbf{g}$ and forming $\mathbf{v}^T\mathbf{b}$. This may be demonstrated by one line of matrix-vector algebra:

$$\mathbf{v}^T\mathbf{b}=\mathbf{v}^TA\mathbf{x}=(A^T\mathbf{v})^T\mathbf{x}=\mathbf{g}^T\mathbf{x} \qquad (2.1)$$

It should be apparent that if the objective function is to be evaluated for K different vectors $\mathbf{b}$, but only one $\mathbf{g}$, only one dual solution is needed, not K primal solutions. Provided one dual solution is not significantly more costly than one primal one, a significant saving may be realised.

The main tension arises when the model is not self-adjoint, in that the dual problem is then different from the one that the physicist would like to understand. This can be an additional reason for preferring a variational approach.

Chapter 3

Summary

Due to a significant amount of preparatory work, the requirements capture exercise from late 2019 to end of Y1 did not raise any significant issues that had not been anticipated in the Science Plan [33]. Concerning modelling and software, there was remarkably little dissension. The question of the physics to be included has been settled in the short term by the need to align with the E-TASC projects on edge modelling. Longer term there are questions still to be resolved concerning the details of the gyrokinetic (or kinetic) model to be employed, the inclusion of special relativistic effects, plasma chemistry, and for example whether the PFC boundaries should be allowed to 'melt', and whether particle dynamics within the top layers of PFCs should be included.

There is still a debate about how best to deal with situations when the software even at Exascale is incapable of resolving boundary or internal layers. A feature of the lower order (Patankar) fd scheme is that it can be formulated so that in an implicit time advance of a field advection-diffusion equation, it acts a contraction mapping upon the field, ie. it converges to a single, finite solution, regardless of lack of layer resolution or of size of timestep. Under such circumstances, a spectral scheme may fail completely cf. the dispersion analyses of say Ainsworth et al [34] and more recently for spectral/hp element schemes in the Galerkin [35] and discontinuous Galerkin [36] contexts - users generally prefer an answer to none at all. Of course, it may be argued that a manifestly erroneous result without an accuracy estimate will not be used to action large procurements, but in any event it would be better for the spectral scheme to produce a result at increasingly extreme parameters.

There are a number of ways to achieve this, the most obvious being the use of explicit artifical viscosity, to broaden thinner layers or small features to the point where full spectral accuracy is achievable. Different options have been explored, of which the most robust appears to be 'DG-mimicking spectral vanishing viscosity' (SVV), see ref [37]. Fernandez et al [38] specifically addresses robustness when mesh resolution becomes poor.

A better approach from the point-of-view of accuracy is to insert more resolution (more elements or increased order polynomials) where aliasing error has been detected. This has its limitations in terms of cost, but there are also practical issues concerning refinement that still need investigation. Some of these latter points might be more easily addressed by reformulating the problem in terms of a variational approach and/or using the Lie derivative formulation [39]. It was anticipated that these could be addressed in the cross-cutting programme.

Interaction with a reactor design framework awaits a better definition of data structures for such tools. There seems no reason however why this should not be addressed inside-out, as regardless, techniques have to be produced for moving data at Exascale.

Subject to the above qualifications, the production of proxyapps should proceed in Y2 as indicated in the Activities Plan [40], viz. a larger task :

- to define a referent physics model for the tokamak edge region, accounting for magnetised plasma behaviour in the presence of significant numbers of neutral atoms and molecules, allowing for radiation and chemical reactions, and identifying important wall interactions such as sheath formation.

This task will be expected to interact with other tasks to ensure a feasible implementation. It is desirable that theroretical support be provided for the EBC pilot code development, which is a 2-D fluid model, as well as for the 1-D fluid cases with kinetic effects explicitly listed below.

The candidate algorithms are expected to employ spectral finite element and particle representations. There are 5 tasks for which advanced mathematical skills will be important:

1. to assess performance of spectral elements for NEPTUNE,
2. to examine the optimal replacement of plasma species properties represented on high-order spatially accurate meshes by a particle representation and vice versa,
3. to study uncertainty quantification (UQ) techniques for NEPTUNE,
4. to study model order reduction techniques for NEPTUNE,
5. to investigate matrix-preconditioning techniques for NEPTUNE. There are 4 tasks to develop proxy-apps for NEPTUNE of demonstrable, high accuracy in challenging test-cases, namely:

a) a 2-D model of anisotropic heat transport,

b) a 2-D elliptic solver in complex geometry,

c) a 1-D fluid solver with simplified physics but with UQ and realistic boundary conditions,

d) a 1-D plasma model incorporating velocity space effects. There are 2 tasks concerning software engineering for NEPTUNE:

1. to investigate DSL and code generation techniques for NEPTUNE and
2. to investigate, in collaboration with UKAEA staff, data structures and design patterns for NEPTUNE.

The Science Plan [33] shows how the proxyapps should feed into the 5-year development. All the tasks may use any identified software packages including commercial software as part of the demonstration process, provided a feasible route to producing code freely usable by NEPTUNE is clearly indicated. It will obviously be better if a task is linked to the delivery of a proxy-app.

Acknowledgement

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

Bibliography

[1] H. Wilson. The impact of plasma physics on the timescale to a tokamak fusion power plant. *Philosophical Transactions of the Royal Society A*, 377(2141):20170435 (14 pages), 2019.

[2] E. Surrey. Engineering challenges for accelerated fusion demonstrators. *Philosophical Transactions of the Royal Society A*, 377(2141):20170442 (14 pages), 2019.

[3] Iter Organisation. ITER Research Plan within the Staged Approach. Technical Report ITR-18-003, Iter Organisation, 2018. https://www.iter.org/doc/www/content/com/Lists/ITER\%20Technical\%20Reports/Attachments/9/ITER-Research-Plan_final_ITR_FINAL-Cover_High-Res.pdf.

[4] Eurofusion. European Research Roadmap to the Realisation of Fusion Energy. Technical Report SOFT 2018 Version, Eurofusion, 2018. https://www.euro-fusion.org/fileadmin/user_upload/EUROfusion/Documents/2018_Research_roadmap_long_version.pdf.

[5] S.V. Patankar. *Numerical heat transfer and fluid flow*. Hemisphere Publishing, Washington, 1980.

[6] B.J. Braams. *Computational studies in tokamak equilibrium and transport*. PhD thesis, Rijksuniversiteit Utrecht, Netherlands, 1986.

[7] W. Arter. Numerical simulation of magnetic fusion plasmas. *Reports on Progress in Physics*, 58:1–59, 1995. http://dx.doi.org/10.1088/0034-4885/58/1/001.

[8] Z. Carter. An Application of the Spectral/hp Element Method to Plasma Fluid Modelling in BOUT++. Technical Report MSc Thesis Report, University of York, 2015.

[9] M. Bhojwani. Simulating edge localised modes within Tokamak plasmas using hp/spectral finite elements in Nektar++. Technical report, Imperial College London, Department of Aeronautics, 2015.

[10] G. Karniadakis and S. Sherwin. *Spectral/hp element methods for computational fluid dynamics 2nd Ed.* Oxford University Press, 2005.

[11] W. Arter. Challenges for modelling fusion plasmas. Technical Report Presentation at Isaac Newton Institute Plasma Physics Day, UKAEA, 2019. https://www.newton.ac.uk/seminar/20191028140014452.

[12] C. Pozrikidis. *Introduction to finite and spectral element methods using MATLAB*. CRC Press, 2005.

[13] M. Turner, D. Moxey, J. Peiró, M. Gammon, C.R. Pollard, and H. Bucklow. A framework for the generation of high-order curvilinear hybrid meshes for CFD simulations. *Procedia Engineering*, 203:206–218. http://dx.doi.org/10.1016/j.proeng.2017.09.808.

[14] L. Botti. Influence of reference-to-physical frame mappings on approximation properties of discontinuous piecewise polynomial spaces. *Journal of Scientific Computing*, 52(3):675–703, 2012. http://dx.doi.org/10.1007/s10915-011-9566-3.

[15] C.D. Cantwell, D. Moxey, A. Comerford, A. Bolis, G. Rocco, G. Mengaldo, D. De Grazia, S. Yakovlev, J.-E. Lombard, D. Ekelschot, B. Jordi, H. Xu, Y. Mohamied, C. Eskilsson, B. Nelson, P. Vos, C. Biotto, R.M. Kirby, and S.J. Sherwin. Nektar++: An open-source spectral/hp element framework. *Computer Physics Communications*, 192:205–219, 2015. https://doi.org/10.1016/j.cpc.2015.02.008.

[16] W. Arter, L. Anton, and D. Samaddar. Options for Geometry Representation. Technical Report CD/EXCALIBUR-FMS/0012-1.00-M2.1.1, UKAEA, 2020.

[17] D.C. Calleja, W. Arter, M. De-Angelis, and E. Patelli. Strategy for Sensitivity Analysis of DEMO first wall. In *3rd International Conference on Vulnerability and Risk Analysis and Management (ICVRAM)*, April 8-11, 2018. https://www.researchgate.net/publication/325793196_Strategy_for_Sensitivity_Analysis_of_DEMO_first_wall.

[18] W. Arter, A. Osojnik, C. Cartis, G. Madho, C. Jones, and S. Tobias. Data assimilation approach to analysing systems of ordinary differential equations. In *2018 IEEE International Symposium on Circuits and Systems*

(ISCAS), Florence, Italy, number 8351751, 27-30 May 2018. http://dx.doi.org/10.1109/ISCAS.2018.8351751.

[19] E.T. Meier, V.S. Lukin, and U. Shumlak. Spectral element spatial discretization error in solving highly anisotropic heat conduction equation. *Computer Physics Communications*, 181(5):837–841, 2010. https://doi.org/10.1016/j.cpc.2009.12.018.

[20] W. Arter. Software Engineering for Fusion Reactor Design. Technical Report Presentation at Frontiers in Scientific Software - 2018 SEA conference, CCFE, 2018. http://dx.doi.org/10.13140/RG.2.2.18856.67843.

[21] W. Arter, N. Brealey, J.W. Eastwood, and J.G Morgan. Fortran 95 Programming Style. Technical Report CCFE-R(15)34, CCFE, 2015. http://dx.doi.org/10.13140/RG.2.2.27018.41922, https://scientific-publications.ukaea.uk/wp-content/uploads/CCFE-R-1534.pdf.

[22] L. Anton, W. Arter, and D. Samaddar. NEPTUNE: Background information and user requirements for design patterns. Technical Report CD/EXCALIBUR-FMS/0015-1.00-M3.3.1, UKAEA, 2020.

[23] D. Samaddar. Report on Y1 2019 Internal Workshop. Technical Report CD/EXCALIBUR-FMS/0018-1.00-M1.1.1a, UKAEA, 2020. Report on Y1 2020 External Workshop. Technical Report CD/EXCALIBUR-FMS/0010-1.00-M1.1.1b, UKAEA, 2020.

[24] W. Arter, L. Anton, and D. Samaddar. Options for Particle Algorithms. Technical Report CD/EXCALIBUR-FMS/0013-1.00-M2.3.1, UKAEA, 2020.

[25] F. Wilczynski, D.W. Hughes, S. Van Loo, W. Arter, and F. Militello. Stability of scrape-off layer plasma: A modified Rayleigh–Bénard problem. *Physics of Plasmas*, 26(2):022510, 2019. https://doi.org/10.1063/1.5064765.

[26] S. Grossmann and D. Lohse. Scaling in thermal convection: a unifying theory. *Journal of Fluid Mechanics*, 407:27–56, 2000.

[27] D.A. Nield. The Rayleigh-Jeffreys problem with boundary slab of finite conductivity. *Journal of Fluid Mechanics*, 32(2):393–398, 1968.

[28] J. Mizushima, S. Hayashi, and T. Adachi. Transitions of natural convection in a horizontal annulus. *International Journal of Heat and Mass Transfer*, 44(6):1249–1257, 2001. https://doi.org/10.1016/S0017-9310(00)00188-5.

[29] S. Chandrasekhar. *Hydrodynamic and hydromagnetic stability*. Clarendon Press, 1961.

[30] E.N. Lorenz. Deterministic Nonperiodic Flow. *Journal of the Atmospheric Sciences*, 20(2):130–141, 1963.

[31] Nandan Gokhale, Nikos Nikiforakis, and Rupert Klein. A dimensionally split cartesian cut cell method for hyperbolic conservation laws. *Journal of Computational Physics*, 364:186–208, 2018.

[32] M. Giles and N. Pierce. An introduction to the adjoint approach to design. *Flow, Turbulence and Combustion*, 65:393–415, 2000.

[33] W. Arter, L. Anton, D. Samaddar, and R. Akers. ExCALIBUR Fusion Modelling System Science Plan. Technical Report CD/EXCALIBUR-FMS/0001, UKAEA, 2019.

[34] M. Ainsworth and H.A. Wajid. Dispersive and dissipative behavior of the spectral element method. *SIAM Journal on Numerical Analysis*, 47(5):3910–3937, 2009. http://dx.doi.org/10.1137/080724976.

[35] R.C. Moura, S.J. Sherwin, and J. Peiró. Eigensolution analysis of spectral/hp continuous galerkin approximations to advection–diffusion problems: Insights into spectral vanishing viscosity. *Journal of Computational Physics*, 307:401–422, 2016.

[36] R.C. Moura, S.J. Sherwin, and J. Peiró. Linear dispersion–diffusion analysis and its application to under-resolved turbulence simulations using discontinuous galerkin spectral/hp methods. *Journal of Computational Physics*, 298:695–710, 2015.

[37] R.C. Moura, M. Aman, S.J. Sherwin, and J. Peiró. Spatial eigenanalysis of spectral/hp continuous galerkin schemes and their stabilisation via dg-mimicking spectral vanishing viscosity: Application to high reynolds number flows. *Journal of Computational Physics*, page 109112, 2019.

[38] P.Fernandez, R.C. Moura, G. Mengaldo, and J. Peraire. Non-modal analysis of spectral element methods: Towards accurate and robust large-eddy simulations. *Computer Methods in Applied Mechanics and Engineering*, 346:43–62, 2019.

[39] J. Larsson. A practical form of Lagrange–Hamilton theory for ideal fluids and plasmas. *Journal of Plasma Physics*, 69(3):211–252, 2003.

[40] W. Arter, L. Anton, D. Samaddar, and R. Akers. ExCALIBUR Fusion Modelling System Activities Y1-2. Technical Report CD/EXCALIBUR-FMS/0004, UKAEA, 2019.