UK Atomic Energy Authority

ExCALIBUR

Identification of a preferred overall numerical scheme

M2.6.1

Abstract
The report describes work for ExCALIBUR project NEPTUNE at Milestone 2.6.1. Minutes of meeting to form report on technical progress.

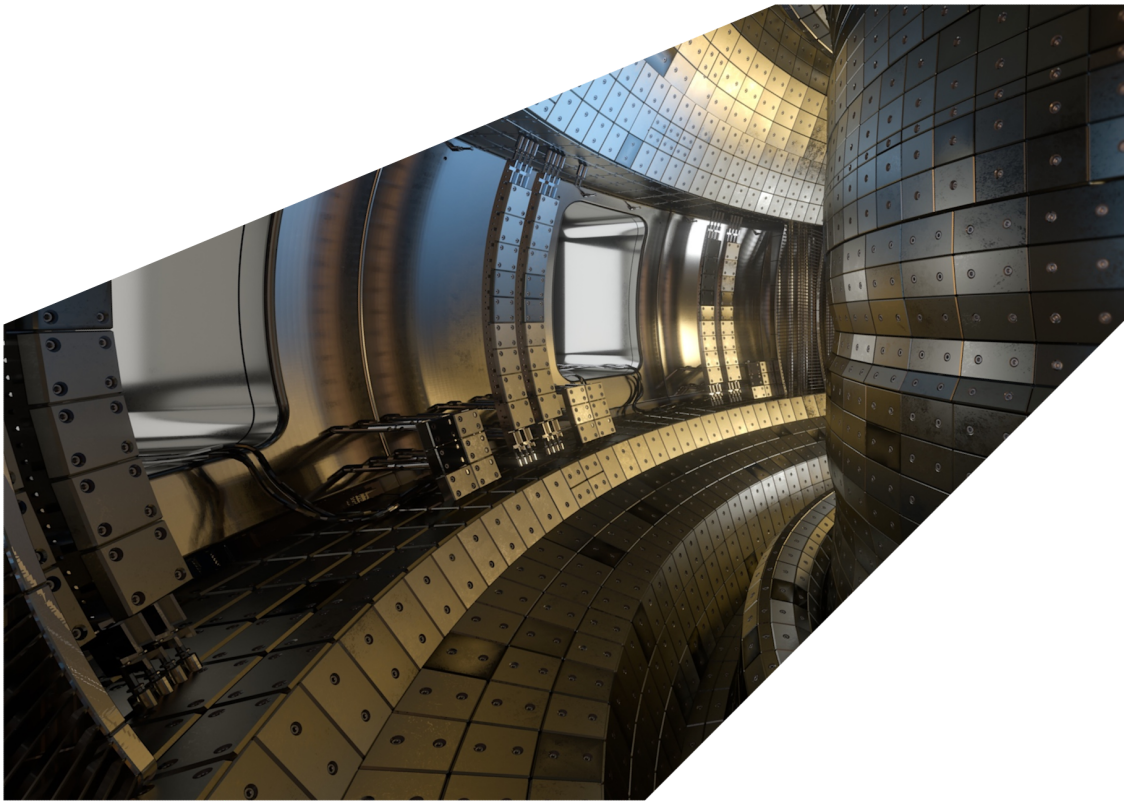

UKAEA REFERENCE AND APPROVAL SHEET

	Client Reference:	
	UKAEA Reference:	CD/EXCALIBUR-FMS/0038
	Issue:	1.00
	Date:	June 28, 2021
Project Name: ExCALIBUR Fusion Modelling System		

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

1 NEPTUNE Meeting: 17 June 2021 11.00-12.00 BST

Present

- Wayne Arter, UKAEA
- Ben Dudson, University of York
- Ed Threlfall, UKAEA

2 Minutes

This was a one-on-one meeting between UKAEA and Ben Dudson, holder of the grant T/NA083/20, to discuss progress.

WA explained that this meeting was the basis for a progress report to the Met Office, serving to identify the latest intelligence on promising research directions. Another reason for this particular session was to discuss future plans in light of staff movements at York, which was discussed first.

2.1 BD's relocation and consequences

BD provided an update on the structure of York's future role in NEPTUNE : David Dickinson is currently spending 10% of his time on the project, to rise to 20%; he is taking on tasks linked to Oxford (ie. 1-D kinetics). BD gave an update on Ed Higgins - BD is liaising with Emma Barnes, the effective 'head' of RSEs at York, re EH taking over from Peter Hill (n.b. EH now plans to stay at York). EH has background in material science but has worked on automated testing framework, community coordination etc (typical RSE activities) but lacks the fusion experience of PH. WA expressed some concern at only getting 20% of DD's time - BD said the latter had other commitments inc. doctoral training (CDT) and STEP. WA said he would have to recommend a deceleration of affected parts of the NEPTUNE schedule. BD agreed that some aspects should be prioritized eg. kinetics (DD's main interest) and conceded that Task 5 (complicated 2-D fluid model) would have to be assigned lower priority. WA said the latter was something of a pity as collaboration with TSVV would have been possible and a superior version of that code would have been an asset.

WA asked about the position re BD working on NEPTUNE after taking up his new post at Lawrence Livermore National Laboratory (LLNL). BD made it clear that his terms preclude accepting any payment except through the US DoE, so would need to propose formally to the DoE in order to work on NEPTUNE ; he thought there might be an acceptable case for this. WA cited his own private sector experience and also mentioned that UKAEA have ongoing programmes with some US labs, especially for fusion work (Rob Akers' area). BD had in the past liaised with Rob Akers and Fulvio Militello re setting up a UKAEA-LLNL memorandum of understanding (MoU), with an agreement to develop software for fusion activities.

BD's future employer is keen for him to continue BOUT++ development - WA asked what the key aims were and whether this included extensions to higher-order methods. BD replied that the

main interest is in incorporating additional physics (he has in the past tried adding a fourth-order accurate finite-difference scheme but it was fragile on a non-uniform grid).

2.2 Main Discussion

WA asked whether BD was still happy to implement the multispecies plasma model (Task 4 for BOUT++, System 2-5 in ref [1]); BD answered that much of this is already underway under the auspices of STEP (BD shared a link to Hermes-3 [2] via the chat); the concept is a modular multi-species model. WA asked whether the design could handle adding different species, to which BD replied yes and offered the manual as a source of further information [3]. WA asked whether the design had a generic species with attributes such as mass and charge. BD gave an explanation of two designs he had tried. 1) In the existing (1-D) SD1D code he had introduced a species class and also reactivity functions which acted to provide sources/sinks - this worked well in 1-D but was cumbersome in 2-D/3-D. 2) A design with a state object was tried - reactions etc. modify this state, which passes down the task chain; this approach was more flexible but it lacked 'checks and balances', also any task could modify the state, leading to ordering and calling convention issues and the possibility of using unset values. In summary, the latter scheme was not fully robust. WA agreed that design 2) would involve a lot of work. BD promised to write up both designs. WA made the point that design 2) has similarity to the mechanics of computer games and BD admitted that the ENTT framework used by MINECRAFT had been an inspiration; WA mentioned in this context that MAYA is well-documented, see book by Gould [4], although of course MAYA now as an Autodesk product is commercial software. BD indicated that he has specifically considered ENTT and provided a github link [5]. Regarding actual equations, BD said he would look at the terms in the equation system 2-5 of the call and check that they can be added. WA said documentation would be very important for future work.

The discussion moved onto the area of progress. WA mentioned work done by Will Saunders on higher-order methods for 1-D fluids in FIREDRAKE: boundary conditions were being investigated but WA admitted the grid used was currently 'generous' in number of FE nodes- this work is progressing. BD agreed this was promising. WA said UKAEA is leveraging WS's existing FIREDRAKE experience, though he has a background in molecular dynamics and not plasmas. It was suggested that some of BD's workload could be taken up by WS. However, there is still a need for a 1-D multispecies code, for which UKAEA needs evidence that it works with higher-order methods. WA said the intention is for Sarah Newton to help guide WS toward solutions of more realistic problems to explore the limits of FIREDRAKE, although the recently received report [6] on test problems had now also been passed on to WS. WA said the current lesson from NEKTAR++ simulations is that if there is a problem, decrease the mesh size, possibly by whole orders of magnitude. BD asked whether the FIREDRAKE DSL was still a candidate for becoming a common language for the various frameworks involved by NEPTUNE. WA mentioned the opinion of Steven Wright that FIREDRAKE's generated C code is opaque and not readily maintainable in itself (BD called it a one-way translation). WA said NEPTUNE potentially could fund software that was better in this regard: UFL is not an acceptable option for non-FEM-conversant physicists (eg. because of the need to understand weak form). WA mentioned the issue for FEM of spatially-dependent parameters such as the thermal conductivity, and the many different terms appearing in eg. Zhdanov [7]. The standard solution is to assign nodal (point) values; BD confirmed that BOUT++ does something

similar eg. at cell boundaries, but made the point that proofs of convergence become equivocal in the presence of such complexities. ET mentioned his own experience with discontinuous Galerkin was for physical parameters to be elementwise constant. WA said this might be acceptable for slowly-varying quantities, but BD brought up the fact that, in plasma physics, some quantities vary as $T^{\frac{5}{2}}$ which implied rapid spatial variation. WA said the opinion of Spencer Sherwin should be sought.

WA then asked BD specifically if any 'showstopper' problems had come to light; BD replied that he was not aware of such; however, David Moxey had identified the need to make significant changes to NEKTAR++ in order to incorporate velocity-space effects (eg. a different basis set). WA said that he had hoped to be further advanced by now, notably it was still unclear whether kinetics problems should be solved via particle methods or 5-D/6-D version of NEKTAR++. WA asked BD to confirm that BOUT++ still does only finite-difference - BD confirmed. WA asked about conservative vs. non-conservative finite-difference - BD's opinion was that the choice sometimes makes rather little difference, being of greatest import in a closed system; in the scrape-off layer, it may not be such an issue due to the large throughflow of energy. WA mentioned concerns about mass conservation; BD said this was mostly important for high-recycling systems where absolute fluxes are large compared to sources. WA said charge conservation was also important eg. in PIC schemes. It was agreed that any invariants of a code should be explicitly checked; and also anything physically conserved but not constrained as constant should also be checked. (WA said that he understood that the ideal scheme conserving mass, energy, and momentum in the absence of dissipation was very hard to realise in practice.) WA asked about these sorts of checks in BOUT++. BD said there were routines for mass conservation checking but that energy conservation was more complicated as the operators for eg. $|\nabla u|^2$ need to be implemented identically in post-processing code. Post-processing in BOUT++ is via a Python back-end. WA mentioned that another problem of finite-differences is that refinement to ensure local energy conservation means global refinement due to the nature of the grid, but BD had little to say on specifically mesh refinement, though there is some non-dynamical boundary refinement in BOUT++. WA agreed meshing questions were David Moxey's area (BD's main role lies in the physics, though he is currently learning FEM).

WA asked whether any 'showstoppers' has arisen in Hermes-3 - BD answered that there were no new issues, just slowness due to equations that already were known to be extremely stiff. WA asked about coupling - BD replied that the system was tightly-coupled with all species solved simultaneously. WA said it would be a good problem for a preconditioner and BD agreed, stating also that the choice of time-integrator was key - he had been searching for quick and robust schemes eg. backward Euler. WA said there might be NEPTUNE resource to investigate time-stepping; BD mentioned SUNDIALS, which is the default choice for BOUT++. WA asked for BD's opinion on time-steppers, given that higher-order Runge-Kutta (RK) seems currently fashionable. BD replied that he has not tried implicit RK but that SUNDIALS is generally hard to beat. WA favoured investigating high-order implicit RK; BD said that he had heard good things in that area from Met Office. WA said that it was natural to use RK for discontinuous Galerkin, which is a kind of RK in spatial coordinate in that each completed time advance requires only values within the timestep.

WA summarized that the main emphasis was in the physics and that no new problems had so far been revealed (especially issues that scale badly), leaving only known problems such as the N^2 complexity with number of species N. BD agreed, reminding that the size of the state vector

also scaled with the grid size. WA said the problem of multispecies coupling was one area where the power of Exascale HPC appeared to be absolutely necessary (though obviously Exascale would also help in many other areas). BD said the problem should be decoupled if possible. WA responded that local matrix inverses can be stored but he agreed that they were unlikely to remain usefully constant.

The conversation turned to uncertainty quantification (UQ), with WA citing mentions of intrusive UQ earlier in the project (two years ago). The current position, following involvement from the VECMA community, was to favour non-intrusive UQ, which is easy to parallelize. BD mentioned auto-differentiation and its similarity to the kind of sensitivity analyses done in UQ. WA opined that intrusive UQ is not necessarily difficult to implement but that it involves significant amounts of labour. WA asked BD to postpone if not cancel work on intrusive UQ, adding that non-intrusive UQ was more in line with the NEPTUNE philosophy of separation of concerns. BD agreed, although indicating that he felt it should ultimately be pursued. WA emphasized the need to keep the UQ grant holders busy ('fed' with work) as those contracts end soon and the grant holders ideally should focus on fusion as a use case.

WA asked whether BD had any further matters to raise. BD answered largely in the negative, but again mentioned the time-stepping issue; global communications and small features (stability limit) constrain the time-stepping; the possibility of non-local or asynchronous time-stepping was aired. WA made the point that time is a 'additional coordinate' in most of NEPTUNE numerics; and he mentioned a recent UKAEA Advanced Computing talk on parallel-in-time evolution (by Debasmita Samaddar of UKAEA) which applied parareal time-stepping to SOLPS - showing significant improvements in wall-clock time but accompanied by a similarly increased demand for data storage.

Acknowledgement

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

References

[1] W. Arter. Equations for EXCALIBUR/NEPTUNE Proxyapps. Technical Report CD/EXCALIBUR-FMS/0021-1.00-M1.2.1, UKAEA, 2020.

[2] Hermes plasma edge simulation model: Hermes-3, a hot ion multifluid drift-reduced model. `https://github.com/bendudson/hermes-3`, 2021. Accessed: June 2021.

[3] Hermes-3 Code structure. `https://hermes3.readthedocs.io/en/latest/code_structure.html`, 2021. Accessed: June 2021.

[4] D.A.D. Gould. *Complete Maya Programming: An extensive guide to MEL and C++ API*. Elsevier, 2003.

[5] Entity component system (ECS) in C++. `https://github.com/skypjack/entt`, 2021. Accessed: June 2021.

[6] B. Dudson, P. Hill, E. Higgins, D. Dickinson, S. Wright, and D. Moxey. Task 2.1 1D fluid model tests. Technical Report 2047356-TN-04-2, UKAEA Project Neptune, 2021.

[7] V.M. Zhdanov. *Transport processes in multicomponent plasma*. CRC Press, 2002.