ExCALIBUR

Performance of spectral-hp element methods for the referent plasma models

M2.2.1

Abstract
The report describes work for ExCALIBUR project NEPTUNE at Milestone 2.2.1. Minutes of meeting to form report on technical progress.

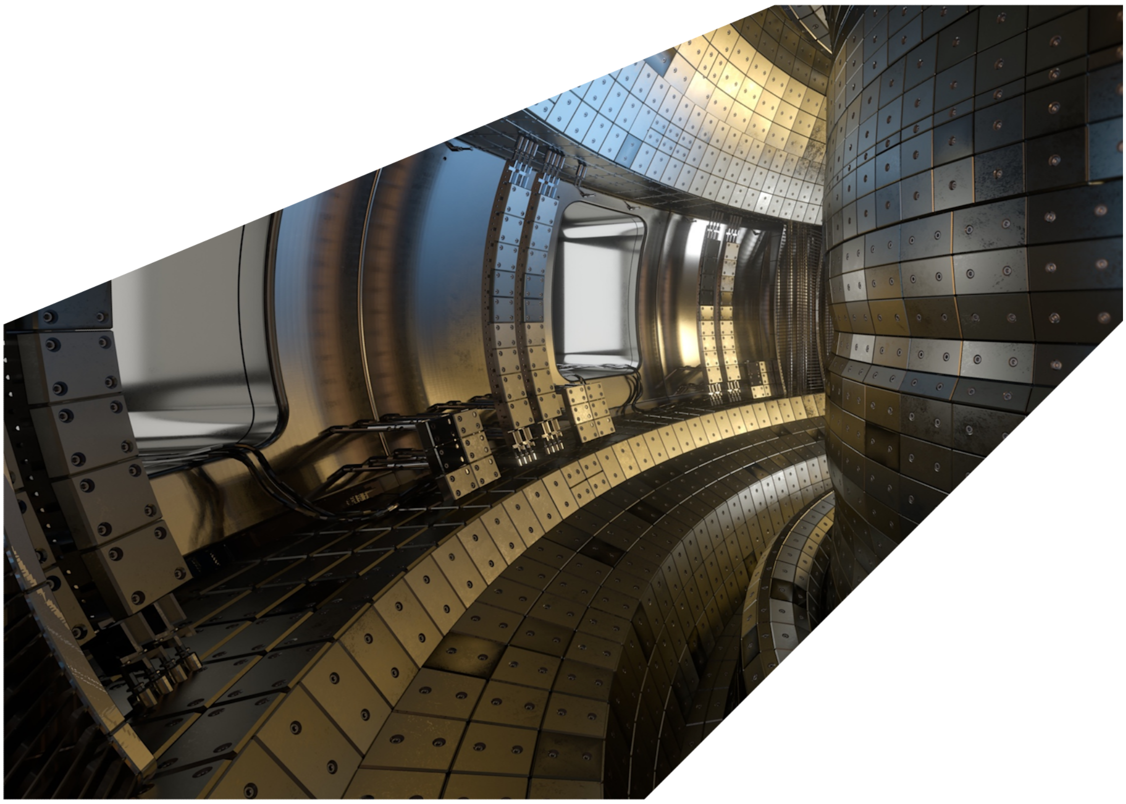

UKAEA REFERENCE AND APPROVAL SHEET

	Client Reference:	
	UKAEA Reference:	CD/EXCALIBUR-FMS/0034
	Issue:	1.00
	Date:	April 27, 2021
Project Name: ExCALIBUR Fusion Modelling System		

	Name and Department	Signature	Date
Prepared By:	Ed Threlfall Wayne Arter BD	N/A N/A	April 27, 2021 April 27, 2021
Reviewed By:	Rob Akers Advanced Computing Dept. Manager		April 27, 2021
Approved By:	Rob Akers Advanced Computing Dept. Manager		April 27, 2021

NEPTUNE Meeting: 23 April 2021 10.00-11.00am BST

Present

- Chair: Wayne Arter, UKAEA
- Ed Threlfall, UKAEA
- David Moxey, Exeter
- Chris Cantwell, Imperial College London

1 Minutes

WA explained that the purpose of this meeting is for the grantees to report on progress to-date and outline plans for the next few months' work.

DM had prepared slides outlining mesh generation, solvers and training work. D1.3 - report on surface mesh generation is currently outstanding due to DM's personal circumstances. WA asked for an estimated ETA for the report; DM responded that he would probably be able to submit a draft by the end of the current month, adding that he has undertaken significant technical work for the report during the last few weeks. WA advised that reports by the other grantees were available for viewing on the NEPTUNE documents repository, to which all grantees should have read access. DM had looked and noticed STFC's reports there.

DM reported on the mesh generation work, which he is leading. Re. the surface meshing work he has taken the approach of simplifying the problem to a spherical geometry (unique radius of curvature) and assessing the surface normal accuracy (ie. deviation from CAD normals) for various levels of mesh refinement. Overall results for his higher-order mesh appear promising, with normals accurate to better than 0.1 degree in a L_2-norm assessment; however, some individual normals showed deviations of 2 degrees from analytic. WA asked if DM was using CADFIX, to which DM answered that he was using OPEN CASCADE with a B-spline sphere (not a NURBS one - which would have rendered the normal test trivial as NURBS can exactly represent a spherical surface). DM said the CAD engine handles the representation of the unstructured sphere, and that he has been working with Joachim Peiro at Imperial on a functional to optimize the surface mesh quality (currently using a 'spring'-type metric, though there are other ways eg. the deviation of the normals could be used as a basis for optimization). DM said the surface distance quality was accurate within $10^{-5} - 10^{-6}$ for a sphere of radius 0.5. WA asked whether the deviations came from the CAD package itself and DM agreed this was a valid question; he proposed doing the analysis for a 'perfect' sphere. Other possible improvements are to use higher-order curvature approximations (currently linear, try quadratic or cubic). (One point that comes to mind here is that it may be better to interpolate the normal vector over a surface face, rather than interpolate a scalar quantity derived from the normal - cf. Phong vs. Gouraud shading, though this was not mentioned during the meeting.) DM says this meshing challenge represents an interesting, solvable problem; his report will outline metrics and meshes. He will also assess GMSH and POINTWISE in terms of the performance of their normal representations. WA added that the main thing was to keep variations smooth - that there can be deviations but they must vary gradually. DM said one question was whether the surface meshing was localizing the error, eg. *NekMesh* cannot currently do a periodic surface mesh so there might be error concentrated at the 'join' in the sphere surface. DM reiterated that his draft report should be ready by the end of next week.

DM moved on to discuss D 1.1 - mesh generation for 2-D cross-sections, in which a key issue is r-adaptivity of meshes at the magnetic field X-point. DM has a code that uses variational optimization (from work on ExCALIBUR project ELEMENT with a PhD student) - currently DM working on tidying this up and improving performance / memory use. WA asked whether this work might benefit from additional expertise in code optimization; DM acknowledged that this might be the case as there are a couple of versions of this code (including CPU and GPU versions) and that expertise in optimization, benchmarking, and profiling might be useful. WA noted that we have future calls to issue that could be used to provide help in this area (eg. Sue Thorne at Rutherford). DM said that one current problem with the 2-D meshing (and also with D1.2 - quadrilateral meshing) was that *NekMesh* ignores curves within the computational domain so cannot currently mesh the X-point geometry (coincidentally, the quad meshing uses something called 'cross-points') and that this deficiency needed to be rectified. He also stressed the importance of keeping his own work aligned with UKAEA requirements (eg. field-aligned meshes in 3-D, not yet attempted). WA added that Ben Dudson was best-placed to provide this alignment. DM was trying to understand the nature of the 3-D meshing challenge; field lines that do not close. WA mentioned that many people use a local approximation to fieldline alignment when computing parallel derivatives, however as far as he knew, all employed a relatively low order finite difference approx., except near the X-point, and that different meshing was sometimes used in the region of that point. He stated the problem was the obvious discontinuity in the toroidal angle where element edges (or faces) will fail to be congruent; a field-aligned global coordinate system is impossible (Ben McMillan apparently has some ideas to treat this). With full FEM, local coordinates would be avoided at the expense of moving all the problems to one plane - where they could be handled by allowing hanging nodes. An alternative is to open a say a 15^o toroidal gap and mesh using another method say with triangular prisms and some relaxation of field-following (possibly propagated through 360^o). DM said NEKTAR++ can handle hanging nodes in discontinuous Galerkin (same problem as a periodic boundary). However, the question of numerical performance with either hanging nodes or a meshed 15 degree gap is open. DM stated that the 2-D problem was tractable within the available 6 month window but that this 3-D problem probably was not. He mentioned that NEKTAR++ can do axially-symmetric geometries using inhomogeneous methods ie. Fourier decomposition in one angle. WA clarified that what he was seeking was DM's expert opinion on how to proceed, but he acknowledged that it might not be possible to provide such opinion at this juncture. DM emphasized that he has yet to understand the physics impact of various meshing options and said he might seek advice from Spencer Sherwin and also Ben Dudson. WA made the point that BOUT++ generally does 2-D cross section plus axisymmetric extrusion runs but that this would not work for real device geometries including details such as tiles, ports and antennas. DM replied that he might understand more once task 2 was complete; this task is CC's work trying to quantify the effect of the mesh on the solution quality, in which task WA commented that the problems of field-alignment and surface alignment are combined.

Moving on to D2.1 - Baseline solver for anisotropic transport, DM said that work at Imperial was leveraging an undergraduate project but Imperial was recruiting a PDRA also for the task; CC cautioned that their candidate requires a work visa, but should be in place by July (the alternate choice had declined the job due to family relocation difficulties). WA asked as to the status of DM's interviewing and DM replied that they had assembled funding for an 18-month PDRA contract (6 months funded via UKAEA, but might be able to spend more than 6 months on NEPTUNE due to a certain flexibility in the remainder of the funding). WA added that future NEPTUNE grants might help, but that they are probably a year away. CC discussed work so far - now that the undergraduate is familiar with the code, good progress is being made. There was a brief mention of code issues to date including a non-physical value for the diffusion constant, but at least work had begun - other problems might be too hard for an undergraduate project and CC or Spencer Sherwin might step in, before July. WA advised that there is an updated, clearer version of the equations document [1]. CC opined that it would be useful to communicate with someone who knows the expected value of the diffusion coefficient to provide sanity checks; WA advised that there exists downloadable software (Fortran) that can produce derived values of the diffusion coefficient, see `https://github.com/wayne-arter/smardda-misc`.

DM moved onto D2.2 - matrix-free kernels, mostly in the hands of Spencer Sherwin and a former postdoc of his. DM has also merged old benchmarks of his into NEKTAR++ including AVX2 / AVX512 for all element types (eg. tetrahedra, pyramids). DM mentioned also the LIKWID environment which is capable of doing FLOP counts, estimates of arithmetic intensity and roofline analysis. His aim is for a performant, matrix-free proxyapp - development to commence in August. He noted that the spatially-varying anisotropy of the Laplacian will affect performance as operator data needs to be loaded from memory. CC added that he and Spencer Sherwin have another student working on variable coefficient Laplacian for hexahedra / quadrilateral meshes (done, triangles and tetrahedra in progress). WA mentioned that Sue Thorne at STFC was looking for relevant matrices to precondition and asked if DM or CC could provide examples. DM replied that he could easily extract matrices from the proxyapp solver and noted that Sue Thorne had emailed him during his recent leave period; he said he would respond to this, acknowledging that one of the challenges of the matrix-free implementation was the preconditioning. WA said that Sue Thorne has matrix free methods eg. Monte-Carlo. DM added that he has code for a p-multigrid preconditioner that works well for purely elliptic problems, but is not so effective for hyperbolic problems or eg. the Euler equations where preconditioning the elliptic part does not suffice. DM mentioned a talk he had heard comparing p-multigrid and LU - but no firm conclusion. He also considers algebraic multigrid and geometric multigrid and is interested to talk with Sue Thorne as he has no experience with Monte Carlo preconditioning. WA noted that Sue Thorne's reports are available on the document repository, possibly on a yet as unmerged branch, see `github.com:ExCALIBUR-NEPTUNE/Documents.git`.

DM discussed then D3.2 - training. He has had 5 or 6 responses to questionnaire re. areas of NEKTAR++ that NEPTUNE partners want to know more about; there was no particular training direction indicated in the responses. He proposed a demonstration of coding eg. the Hasegawa-Wakatani equations in NEKTAR++, a lecture on higher-order methods, and an interactive session on how to write a solver in the framework. DM cited EPSRC calls as the reason why the training schedule had not advanced further; WA reassured him that the training was not proposed to start until August. DM is keen to know problems people have when using and extending NEKTAR++. WA mentioned that Vassil Alexandrov had requested a copy of Zhdanov's derivation of the 13-moment transport equations [2] but had appeared not to get much out of it thus far. WA asked if there was anything CC wanted to add but CC replied that DM had covered everything.

WA mentioned Ben Dudson's proposal planned to run work on every set of proxyapp equations taken from ref [1] in parallel from 1 April and pointed out that DM was listed against these tasks. DM said his work plan was that Task 1 started in June on the NEKTAR++ side, and that he planned to recruit a PDRA to start in July and needed to advertise for this; 6 months PRISM funding for a 12-month PDRA contract and he needs to talk to Rob Akers.

WA said a discussion of repercussions of Ben Dudson leaving the project would be needed at some point; Rob Akers is negotiating with York (WA not privy to details).

WA said a meeting with CADFIX developers would be a good idea, mentioning that UKAEA has two full commercial licences for the software. It turned out that DM has monthly meetings with them already and has done so for a few years. CADFIX has a quad-based adaptive meshing that could be useful for comparison purposes - DM has student involved with this. DM stated that the meshing part of CADFIX was not great. WA added that the CADFIX interface is also poor; DM concurred, adding that it recalls something from the 1990s. However, he said that the API was better than that of OPEN CASCADE, which is very poor (eg. c.40,000 classes and much repeated functionality). WA said that CADFIX is otherwise a good tool, and that most nuclear fusion CAD is done via CATIA^{TM}. DM said CADFIX is good for cleaning up geometries, eg. STEP files. WA commented that most of this CAD / meshing 'needs a human in the loop'. He recommended reconvening for an additional meeting, of having UKAEA attend DM's next meeting with the CADFIX developers; DM will check the date of the next of these meetings. WA commented that those developers might be eligible for a future NEPTUNE grant if the company have sufficient U.K. research personnel to qualify, then WA closed the meeting.

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. 10/3/20.

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