UK Atomic
Energy
Authority

ExCALIBUR

Software Support Complementary Actions 2

M7.3 Version 1.00

Abstract
The report describes work for ExCALIBUR project NEPTUNE at Milestone M7.3. It describes aspects of coordination within the NEPTUNE project not covered in previous Year 3 reports, namely, the development of a GitHub repository for infrastructure code and project planning; the development of a project website for knowledge transfer within NEPTUNE; and a description of collaborations arising from NEPTUNE-related interactions, including the Fusion Modelling Use Case working group established create a connection between Project NEPTUNE and the wider ExCALIBUR programme.

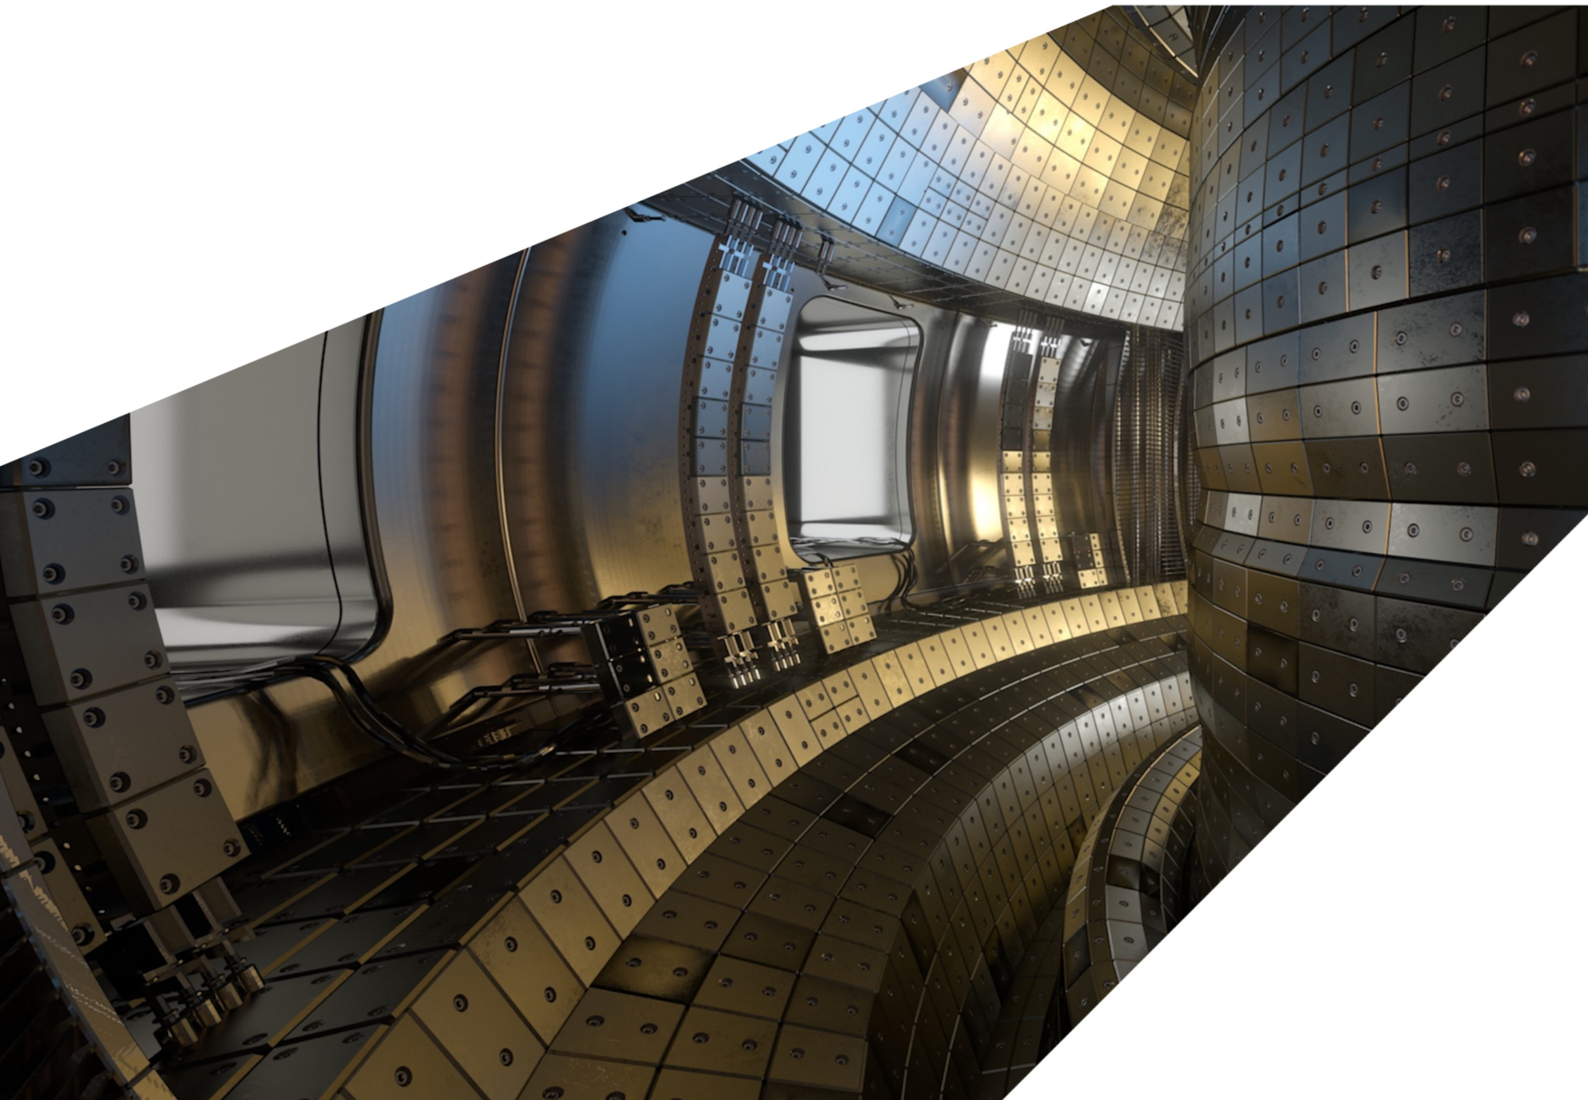

UKAEA REFERENCE AND APPROVAL SHEET

	Client Reference:	
	UKAEA Reference:	CD/EXCALIBUR-FMS/0065
	Issue:	1.00
	Date:	18 March 2022
Project Name: ExCALIBUR Fusion Modelling System		

	Name and Department	Signature	Date
Prepared By:	Joseph Parker Wayne Arter BD	N/A N/A	18 March 2022 18 March 2022
Reviewed By:	Wayne Arter Project Technical Lead	W. Arter	18 March 2022

1 Introduction

Activity 7 “Support and Coordination” is a programme providing support and coordinating work internally to complement Activity 3, and initiating Software Support Procurement to ensure continuity of external work. The main focus of the work is to facilitate the production of a DSL tailored to NEPTUNE requirements, to consider how best to couple NEPTUNE-developed software with other codes, and to determine how and when to benchmark NEPTUNE software against pre-existing tokamak edge modelling software.

In addition, this project provides management of external research, coordinating the additional expert input in support of the numerical analysis project around the topics of matrix preconditioning, asymptotic-preserving methods and multirate/Variable Stepsize-Variable Order timestepping algorithms to both spectral/hp element and particle-based methods.

As such, Support and Coordination contributes primarily to Work Package 4 “Code structure and coordination” (FM-WP4 from the Fusion Modelling System Science Plan [1]), with some work falling under the remits of FM-WP1 “Numerical representation”, and of the cross-cutting packages XC-WP1 “Common approaches & solutions” and XC-WP2 “Emerging technologies”.

In particular, this year’s work focussed on novel timestepping methods, performance benchmarking, and matrix preconditioning. Previous Year 3 Support and Coordination reports covered procurement of mathematical support for advanced timestepping and asymptotic-preserving methods [2], and benchmarking and code coupling [3]. Further reports from grantee work are due at the end of the Financial Year and so will be reported on at a later date.

This report describes other work undertaken to coordinate and foster collaboration within the NEPTUNE community. In Section 2.1 we describe work setting up a repository for NEPTUNE code, prototyping working practices, and writing a prototype software framework in Data Parallel C++ (DPC++). In Section 2.2 we describe the development of the Project NEPTUNE website, designed to act as a record of project planning and management, and to facilitate knowledge transfer among project participants. In Section 2.3 we describe the activities of a working group set up in support of the ExCALIBUR Fusion Modelling Use Case, bringing together UKAEA staff, NEPTUNE grantees and external partners from STFC and academia. In Section 2.4 we describe ongoing collaborations arising from Project NEPTUNE but not due specifically to other Activities. Finally, in Section 3, we summarize.

2 Task Work

2.1 PolyRepoPracticeCore

The NEPTUNE development philosophy is “separation of concerns”, where domain specialists write proxyapps that focus on some aspect of the overall code. Production of proxyapps is ongoing in accordance with the Science Plan [1]. These proxyapps need to be drawn together into a single coherent codebase, something that requires infrastructure both for code coupling and for code maintenance and testing.

As recommended in [4], each proxyapp will be hosted in a separate GitHub repository under the NEPTUNE organization. Then a main NEPTUNE repository hosts software to interface the proxyapps, which pulls in the proxyapps as git submodules. This enables separation of concerns as a proxyapp's developer can focus largely on their own work, making pull requests, raising issues etc. in their proxyapp's repository without extraneous pull requests, issues and project planning relating to other proxyapps. The main repository is then the place for discussions relating to interactions between proxyapps. For example, a bug fix in a proxyapp would lead to an issue in its own local repository, but if the fix would change the proxyapp's expected output or external API, it would lead to an issue in the main NEPTUNE repository.

To prototype ideas and refine working practices, this main repository has been created [5]. This repository is being used to host development of infrastructure code (I/O, build systems, testing frameworks), project planning (issues and kanban boards) and a prototype framework for the NEPTUNE physics software. The software presently targets the solution of the 1+1D Vlasov–Poisson system using a particle-in-cell (PIC) method, since this is a well-studied and benchmarked system with many existing test problems and implementations.

The implementation is modular, allowing other physical systems to be solved by changing the information passed to the particle pusher and elliptic solver. Similarly, the dimensionality and geometry of the problem may be changed by defining new implementations of the classes `Mesh`, `Space` and `Velocity`. In particular, the space and velocity space/particle implementations are separated so that it is possible to add the option of spectral/hp implementations in space alongside the grid-based approach already implemented. Thus it will be possible to rapidly prototype a spectral/hp in space + PIC code when a Nektar++ API becomes available through NEPTUNE project work.

2.1.1 SYCL parallelization

The prototype code is parallelized using oneAPI's DPC++/SYCL implementation [6]. The conclusions from this exercise are mixed. The process of parallelizing a small serial code with DPC++ is moderately straightforward. On a local machine, the code's integrated test that took 12 seconds on a single thread before parallelization, ran in just over 3 seconds on 4 threads, with no optimizations besides replacing normal `for` loops with SYCL's `parallel_for`.

The developer experience however was less positive. While much documentation is provided by Intel, it is not always clear which documents are current, and which approaches and coding styles are recommended. There is a library of sample code [7] but this often proposes multiple approaches to the same problem, without making clear which approach is best. For example, the sample reduction code provides six different methods for performing a reduction from a vector to a number. However, there is no library function for performing a reduction into a vector (such as a `sum all reduce`), and this had to be implemented manually in the prototype code. In addition the language is very verbose, and it seems unlikely that details of the SYCL implementation could be fully abstracted from the rest of the code without using macros.

2.2 NEPTUNE Website

A vital aspect of support and coordination is to ensure that accurate, up-to-date technical information concerning the entire ExCALIBUR project NEPTUNE is available to all the community. It has been argued [8] that this is best met by means of a website, see ref [9]. Work has continued to develop the website, primarily directed towards design of the NEPTUNE software, since the two reports appeared in October. There is now a GitHub action to verify that all input is in ASCII.

The high-level website structure is summarised as follows:

PN	Program name - analogous to book title and preface, contains important administrative information
BD	Business design - analogous to book flyleaf or introductory chapter, designed to ‘sell’ the project to potential customers
RB	Requirements baseline - where use cases define functional requirements
TS	Technical specification - documents decisions taken to meet functional requirements, includes non-functional requirements such as software standards and conventions
DJF	Design justification file - place to store all ideas considered to meet spec.
DDF	Design definition file - definition of software using selected ideas
MGT	Management file - contribution from ref [4]
MF	Maintenance file - not yet needed
OP	Operational documentation - contribution from ref [4]
REF	Reference Material - to ensure consistency of usage of mathematical notation, acronyms and terminology (‘ontology’)
IND	Index - only needed in absence of search engines

It is envisaged that this structure will cascade down to proxyapps, which will be able to link upwards to common material such as in item REF above, in order to minimise the amount of new documentation to be produced.

All entries have been overhauled as part of the ongoing work. Item RB now contains order-of-magnitude estimates for physical quantities relevant to tokamak edge modelling and use cases for NEPTUNE devised by each member of the team. These cases are intended to be relevant to the needs of a physicist, a design engineer, a finite element expert and a particle specialist respectively. Item TS now includes order-of-magnitude estimates for calculation costs of highly detailed edge computations. DJF has a list of reports and a brief survey of each report’s contents is in the process of being added.

Item REF has been considerably expanded, so that there are now over 300 acronyms covering both the highly technical (*e.g.* SSA for Singular Spectrum Analysis) to the purely administrative

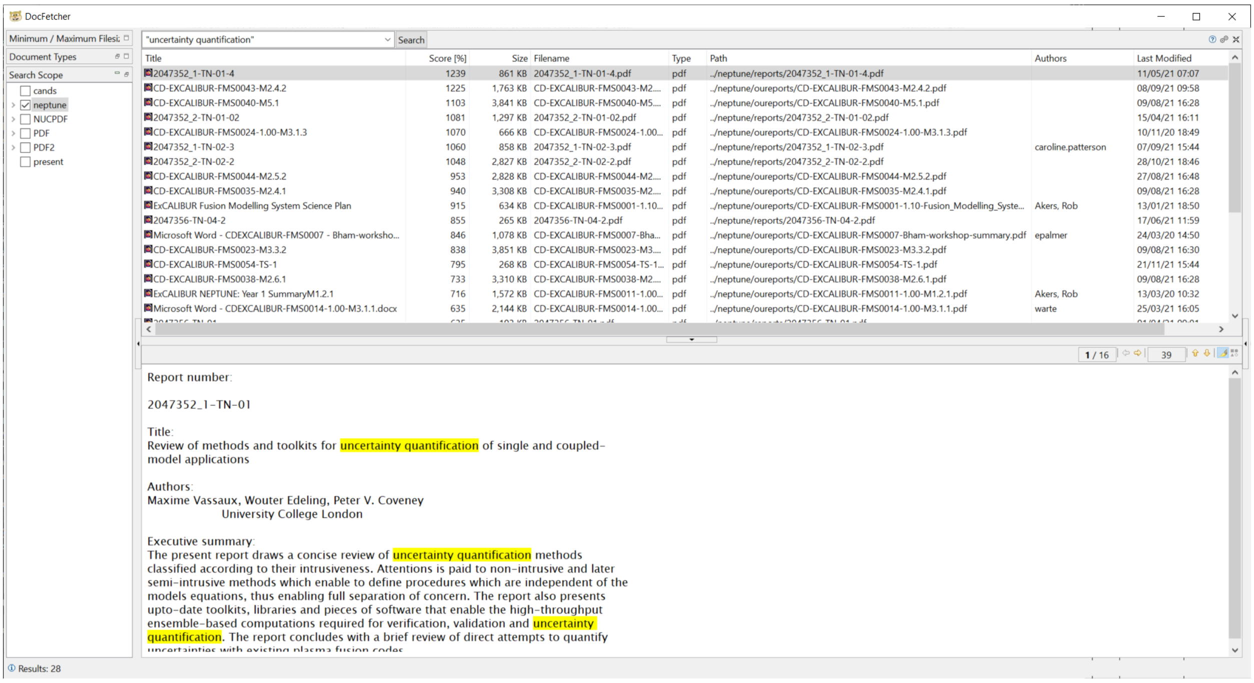

Title	Score [%]	Size	Filename	Type	Path	Authors	Last Modified
2047352_1-TN-01-4	1239	861 KB	2047352_1-TN-01-4.pdf	pdf	../neptune/reports/2047352_1-TN-01-4.pdf		11/05/21 07:07
CD-EXCALIBUR-FMS0043-M2.4.2	1225	1,763 KB	CD-EXCALIBUR-FMS0043-M2....	pdf	../neptune/oureports/CD-EXCALIBUR-FMS0043-M2.4.2.pdf		08/09/21 09:58
CD-EXCALIBUR-FMS0040-M5.1	1103	3,841 KB	CD-EXCALIBUR-FMS0040-M5....	pdf	../neptune/oureports/CD-EXCALIBUR-FMS0040-M5.1.pdf		09/08/21 16:28
2047352_2-TN-01-02	1081	1,297 KB	2047352_2-TN-01-02.pdf	pdf	../neptune/reports/2047352_2-TN-01-02.pdf		15/04/21 16:11
CD-EXCALIBUR-FMS0024-1.00-M3.1.3	1070	666 KB	CD-EXCALIBUR-FMS0024-1.00...	pdf	../neptune/oureports/CD-EXCALIBUR-FMS0024-1.00-M3.1.3.pdf		10/11/20 18:49
2047352_1-TN-02-3	1060	858 KB	2047352_1-TN-02-3.pdf	pdf	../neptune/reports/2047352_1-TN-02-3.pdf	caroline.patterson	07/09/21 15:44
2047352_2-TN-02-2	1048	2,827 KB	2047352_2-TN-02-2.pdf	pdf	../neptune/reports/2047352_2-TN-02-2.pdf		28/10/21 18:46
CD-EXCALIBUR-FMS0044-M2.5.2	953	2,828 KB	CD-EXCALIBUR-FMS0044-M2....	pdf	../neptune/oureports/CD-EXCALIBUR-FMS0044-M2.5.2.pdf		27/08/21 16:48
CD-EXCALIBUR-FMS0035-M2.4.1	940	3,308 KB	CD-EXCALIBUR-FMS0035-M2....	pdf	../neptune/oureports/CD-EXCALIBUR-FMS0035-M2.4.1.pdf		09/08/21 16:28
ExCALIBUR Fusion Modelling System Science Plan	915	634 KB	CD-EXCALIBUR-FMS0001-1.10...	pdf	../neptune/oureports/CD-EXCALIBUR-FMS0001-1.10-Fusion_Modelling_Syste...	Akers, Rob	13/01/21 18:50
2047356-TN-04-2	855	265 KB	2047356-TN-04-2.pdf	pdf	../neptune/reports/2047356-TN-04-2.pdf		17/06/21 11:59
Microsoft Word - CDEXCALIBUR-FMS0007 - Bham-worksho...	846	1,078 KB	CD-EXCALIBUR-FMS0007-Bha...	pdf	../neptune/oureports/CD-EXCALIBUR-FMS0007-Bham-workshop-summary.pdf	epalmer	24/03/20 14:50
CD-EXCALIBUR-FMS0023-M3.3.2	838	3,851 KB	CD-EXCALIBUR-FMS0023-M3....	pdf	../neptune/oureports/CD-EXCALIBUR-FMS0023-M3.3.2.pdf		09/08/21 16:30
CD-EXCALIBUR-FMS0054-TS-1	795	268 KB	CD-EXCALIBUR-FMS0054-TS-1...	pdf	../neptune/oureports/CD-EXCALIBUR-FMS0054-TS-1.pdf		21/11/21 15:44
CD-EXCALIBUR-FMS0038-M2.6.1	733	3,310 KB	CD-EXCALIBUR-FMS0038-M2....	pdf	../neptune/oureports/CD-EXCALIBUR-FMS0038-M2.6.1.pdf		09/08/21 16:28
ExCALIBUR NEPTUNE: Year 1 SummaryM1.2.1	716	1,572 KB	CD-EXCALIBUR-FMS0011-1.00...	pdf	../neptune/oureports/CD-EXCALIBUR-FMS0011-1.00-M1.2.1.pdf	Akers, Rob	13/03/20 10:32
Microsoft Word - CDEXCALIBUR-FMS0014-1.00-M3.1.1.docx	635	2,144 KB	CD-EXCALIBUR-FMS0014-1.00...	pdf	../neptune/oureports/CD-EXCALIBUR-FMS0014-1.00-M3.1.1.pdf	warte	25/03/21 16:05

Report number:

2047352_1-TN-01

Title:
Review of methods and toolkits for uncertainty quantification of single and coupled-model applications

Authors:
Maxime Vassaux, Wouter Edeling, Peter V. Coveney
University College London

Executive summary:
The present report draws a concise review of uncertainty quantification methods classified according to their intrusiveness. Attentions is paid to non-intrusive and later semi-intrusive methods which enable to define procedures which are independent of the models equations, thus enabling full separation of concern. The report also presents upto-date toolkits, libraries and pieces of software that enable the high-throughput ensemble-based computations required for verification, validation and uncertainty quantification. The report concludes with a brief review of direct attempts to quantify uncertainties with existing plasma fusion codes

Figure 1: Result of a search for “Uncertainty Quantification” of NEPTUNE reports, using DOCFETCHER software. Search made on 15 March 2022 of current GitHub repository.

(*e.g.* RO for Responsible Officer), and over 350 mathematical symbols listed, each now with physical units where appropriate. A convention for converting the LaTeX representations to ASCII for use in C++ has been devised as a stopgap pending an agreed, detailed design of a class hierarchy. Considerable effort, involving all the UKAEA NEPTUNE team and external Nektar++ experts, has been put into designing a class or object-oriented hierarchy, preliminary results from which have been added to the website under item DDF.

Information has been extracted from internal reports where appropriate, in particular, the items MGT and OP have been filled as indicated above, by information from ref [4]. Elsewhere links have been provided to pre-existing repository material in LaTeX using the LaTeX `\input` command, and a larger number of citations now has links to the web. All approved NEPTUNE reports by both UKAEA and the grantees have been placed in the GitHub repository [10], although access to it is controlled and users who are not logged into GitHub (or who do not have a GitHub account) will get an ‘Error 404’. Ultimately it is hoped that all reports for which LaTeX source has been provided, will be more tightly linked to the website. However, currently it is recommended that the GitHub repository contents be ‘pulled’ to a local machine and indexed by a desktop search engine such as DOCFETCHER [11] or RECOLL [12]. Figure 1 shows the result of a test search.

2.3 Fusion Modelling Working Group

The Fusion Modelling Working Group is a new working group of people involved in the UKAEA Fusion Modelling Science Use Case (part of the wider ExCALIBUR programme). Currently 15 people are involved, 9 from Project NEPTUNE (6 UKAEA staff and 3 from Nektar++), 3 from STFC (including the chair Stephen Longshaw), 2 from the University of Cambridge and 1 from the University of York.

The STFC members are all involved in a number of EPSRC ExCALIBUR projects – most notably "Integrated Simulation at the Exascale: coupling, synthesis and performance" – but all of which include an element of fusion modelling. The purpose of this group is to highlight areas of ongoing work, fostering new collaboration and prevent duplication of efforts.

An inaugural (virtual) meeting was held on 11th February 2022. The meeting began with a round table of individuals' interests and relevant project involvement. Interests naturally covered a broad range of plasma physics topics (*e.g.* magnetohydrodynamics, gyrokinetics, PIC codes, plasma-wall interactions), as well as related physics (*e.g.* complex gas dynamics, smoothed particle hydrodynamics, strongly anisotropic systems) and code development topics (*e.g.* fluid-to-particle code coupling, grid generation for spectral/hp, uncertainty quantification, software development practices).

The discussion then moved on to what the working group could achieve, and what it could do to support NEPTUNE and wider fusion projects. The following suggestions were proposed:

- creation of a living document on coupling methods to inform the Project NEPTUNE
- definition of data structures used to pass information internally between the components of NEPTUNE
- definition of a working schema for output from fusion modelling, similar to that which exists for climate modelling (and investigation into whether the IMAS framework (summarized in a previous NEPTUNE report [3]) is appropriate for this)
- work to determine appropriate input/output file formats for NEPTUNE
- work on the delivery of software/package management, in support of the existing work package for delivering reproducible coupled software and workflows as part of ExCALIBUR cross-cutting project on code coupling

It was a agreed that the working group should not introduce too many new meetings, and where possible collaborative working should be favoured over meetings. Future meetings will occur roughly monthly, but will be themed so only those interested need attend. More general, larger meetings will be held on a six-monthly basis.

To facilitate collaborative working, a group Slack channel was made (as a channel with the NEPTUNE Slack workspace). A git repository [13] was also created under the ExCALIBUR-NEPTUNE GitHub organization.

2.4 Grantee collaborations

There are two pieces of grantee work under the remit of Support and Coordination which are due by the end of the financial year 2021/22, and which will appear in the NEPTUNE documents repository [10] after review. The first of these is updates to the reports on approaches to performance portability [14, 15, 16] giving recent developments in the software and hardware for Exascale. The second is a report describing the use of ReFrame as a framework for performance benchmarking (ReFrame was discussed briefly in NEPTUNE report [3]).

In addition, there are a number of new collaborations arising from Project NEPTUNE, either resulting from project definitions or from new connections made via the NEPTUNE progress meetings. Of these, the collaboration between STFC and the University of Oxford on matrix preconditioning falls within the scope of the Support and Coordination activity, and is described in the following subsection.

2.4.1 STFC and Oxford: matrix preconditioning

As a result of the Matrix Preconditioning project (NEPTUNE ITT T/NA084/20), STFC (Sue Thorne, Tyrone Rees, Hussam Al Daas, Andrew Sunderland, Emre Sahin, Vassil Alexandrov) have an ongoing collaboration with the University of Oxford (Patrick Farrell). The collaboration focusses on the application of new preconditioning techniques for FEM in NEPTUNE-relevant problems. Initial work presented at a NEPTUNE progress meeting in March 2022, considered solving a 2D Poisson problem using the continuous Galerkin method with different preconditioners: (1) an algebraic two-level Schwarz preconditioner recently derived by the STFC team [17]; and (2) the long-established algebraic multigrid preconditioner [18] as implemented in parallel with BoomerAMG in HYPRE [19].

The 2D Poisson test problem was implemented in Firedrake [20] and used 13th order continuous Galerkin on an 80 by 80 Cartesian grid, giving a total problem size of $1,083,681^2$. This was parallelized on 256 cores. The algebraic multigrid implementation took 744 iterations and 18.7s seconds to converge to a tolerance of 10^{-7}. In contrast, the algebraic two-level Schwarz preconditioner required only 23 iterations and 2.1 seconds to reach the same tolerance, a reduction in the number of iteration of $\sim \times 30$ and a speed up of $\sim \times 9$.

3 Summary

This report summarizes work undertaken as part of the Year 3 “Support and Coordination” Activity that have not been described in the previous reports [2, 3]. The report discusses activities which enable collaboration and knowledge transfer, namely the creation of a Project NEPTUNE software repository (Section 2.1) and website (Section 2.2). Further, it discusses collaborations which either support the NEPTUNE project or have arisen from NEPTUNE project work, namely the Fusion Modelling Use Case working group (Section 2.3) and other collaborations between NEPTUNE grantees (Section 2.4). Grantee work due to be submitted by the end of this Financial Year Y3 will be reported on in due course.

Acknowledgement

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

References

[1] W. Arter, L. Anton, D. Samaddar, and R. Akers. ExCALIBUR Fusion Modelling System Science Plan. Technical Report CD/EXCALIBUR-FMS/0001, UKAEA, 2019. `https://www.metoffice.gov.uk/binaries/content/assets/metofficegovuk/pdf/research/spf/ukaea-excalibur-fms-scienceplan.pdf`.

[2] J.T. Parker and W. Arter. Literature review for Call T/AW086/21: Mathematical Support for Software Implementation. Technical Report CD/EXCALIBUR-FMS/0052-M7.1, UKAEA, 2021. `https://github.com/ExCALIBUR-NEPTUNE/Documents/blob/main/reports/ukaea_reports/CD-EXCALIBUR-FMS0052-M7.1.pdf`.

[3] J.T. Parker, E. Threlfall, W. Arter, and W. Saunders. Code coupling and benchmarking. Technical Report CD/EXCALIBUR-FMS/0053-M7.2, UKAEA, 2021. `https://github.com/ExCALIBUR-NEPTUNE/Documents/blob/main/reports/ukaea_reports/CD-EXCALIBUR-FMS0053-M7.2.pdf`.

[4] E. Threlfall and W. Arter. Specification and Integration of Scientific Software. Technical Report CD/EXCALIBUR-FMS/0042-M3.1.4, UKAEA, 2021. `https://github.com/ExCALIBUR-NEPTUNE/Documents/blob/main/reports/ukaea_reports/CD-EXCALIBUR-FMS0042-M3.1.4.pdf`.

[5] ExCALIBUR-NEPTUNE PolyrepoPracticeCore GitHub Repository. `https://github.com/ExCALIBUR-NEPTUNE/PolyrepoPracticeCore`, 2022. Accessed: March 2022.

[6] J. Reinders, B. Ashbaugh, J. Brodman, M. Kinsner, J. Pennycook, and X. Tian. *Data Parallel C++: Mastering DPC++ for Programming of Heterogeneous Systems using C++ and SYCL*. Springer Nature, 2021.

[7] oneAPI Samples Repository. `https://github.com/oneapi-src/oneAPI-samples`, 2022. Accessed: March 2022.

[8] W. Arter, J. Parker, and E. Threlfall. Development Plan. Technical Report CD/EXCALIBUR-FMS/0033-D3.4, UKAEA, 2021. `https://github.com/ExCALIBUR-NEPTUNE/Documents/blob/main/reports/ukaea_reports/CD-EXCALIBUR-FMS0033-1.10-D3.4.pdf`.

[9] W. Arter, E. Threlfall, J. Parker, and W. Saunders. Software Specification Website. Technical Report CD/EXCALIBUR-FMS/0054-D3.1, UKAEA, 2021. `https://github.com/ExCALIBUR-NEPTUNE/Documents/blob/main/reports/ukaea_reports/CD-EXCALIBUR-FMS0054-D3.1.pdf`.

[10] ExCALIBUR-NEPTUNE Bid Documents, Calls, Reports and Notes. `https://github.com/ExCALIBUR-NEPTUNE/Documents`, 2021. Accessed: March 2021.

[11] DocFetcher is an Open Source desktop search application. `http://docfetcher.sourceforge.net/en/index.html`, 2022. Accessed: January 2022.

[12] Recoll finds documents based on their contents as well as their file names. `https://www.lesbonscomptes.com/recoll`, 2022. Accessed: January 2022.

[13] Fusion Modelling Use Case GitHub Repository. `https://github.com/ExCALIBUR-NEPTUNE/Cross-Cutting-Organisation`, 2022. Accessed: March 2022.

[14] S. Wright, B. Dudson, P. Hill, D. Dickinson, and G. Mudalige. Approaches to Performance Portable Applications for Fusion. Technical Report 2047358-TN-01-02, UKAEA Project Neptune, 2021. `https://github.com/ExCALIBUR-NEPTUNE/Documents/blob/main/reports/2047358/TN-01.pdf`.

[15] S. Wright, B. Dudson, P. Hill, D. Dickinson, and G. Mudalige. Identification of Testbed Platforms and Applications. Technical Report 2047358-TN-02, UKAEA Project Neptune, 2021. `https://github.com/ExCALIBUR-NEPTUNE/Documents/blob/main/reports/2047358/TN-02.pdf`.

[16] S. Wright, B. Dudson, P. Hill, D. Dickinson, and G. Mudalige. Evaluation of Approaches to Performance Portability. Technical Report 2047358-TN-03, UKAEA Project Neptune, 2021. `https://github.com/ExCALIBUR-NEPTUNE/Documents/blob/main/reports/2047358/TN-03.pdf`.

[17] H. Al Daas, P. Jolivet, and T. Rees. Efficient Algebraic Two-Level Schwarz Preconditioner For Sparse Matrices. *arXiv preprint arXiv:2201.02250*, 2022.

[18] S.F. McCormick. *Multigrid methods*. SIAM, 1987.

[19] BoomerAMG parallel implementation of the algebraic multigrid method. `https://hypre.readthedocs.io/en/latest/solvers-boomeramg.html`, 2022. Accessed: March 2022.

[20] F. Rathgeber, D.A. Ham, L. Mitchell, M. Lange, F. Luporini, Andrew T.T. Mcrae, G.-T. Bercea, G.R. Markall, and P.H.J. Kelly. Firedrake: Automating the finite element method by composing abstractions. *ACM Trans. Math. Softw.*, 43(3), December 2016.