HORIZON-EUROHPC-JU-2021-COE-01

The European Centre of Excellence for Engineering Applications

Project Number: 101092621

D2.16

Report on the OpenFoam/Elmer Application Use Case

The EXCELLERAT P2 project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101092621. The JU receives support from the European Union's Horizon Europe research and innovation programme and Germany, Italy, Slovenia, Spain, Sweden and France.

Work Package:	2	Use-Case Execution
Author(s):	Matic Brank	Jernej Kovacic
	Janez Povh	Leon Kos
Approved by	Executive Centre Management	23.11.2023
Reviewer	Ansgar Niemöller	RWTH
Reviewer	Simone Bna	CINECA
Dissemination Level	Public	

Date	Author	Comments	Version	Status
26.10.2023	Matic Brank	Initial draft	V01.0	Draft
4.11.2023	Matic Brank	Reply to first review	V02.0	Draft
17.11.2023	Matic Brank	Reply to second review	V03.0	Draft

List of abbreviations

CEA	Centre of atomic energy, Cadarache
F4E	Fusion for Energy
FEM	Finite Element Method
IR	Infra-red
ITER	International Thermonuclear Experimental Reactor
UC	Use Case
UL	University of Ljubljana

Executive Summary

This document summarises the progress made in the OpenFoam/Elmer application use case (UC) **FA-1** by University of Ljubljana (UL). We present the objectives of the use case and the development of the use case workflow in detail. The introduction explains the purpose and extent of the use case with respect to research in magnetically confined nuclear fusion. Then, the objectives of the use case are explained and its aim with respect to the EXCELLERAT P2 work is discussed. In section 3 the workflow is presented in more detail, explaining individual simulation steps needed to achieve the desired output of the use case.

In sections 4 and 5 the current progress and next steps are presented. The progress achieved so far is presented with focus on a small-scale case and its role as a basis for large scale simulations. These are outlined in the conclusion section with the work layout for the future. The status of algorithmic optimisation is also presented with focus on parallelisation and storing data to binary files using the HDF5 file format.

Table of Contents

Table of Figures

1 Introduction

Fusion energy, as stated by Fusion for Energy (F4E), the European Union organisation managing Europe's contribution to the International Thermonuclear Experimental Reactor (ITER) [1], *"offers the promise for abundant, safe and sustainable energy for the future"*. To contribute to the fulfilment of this promise, further ground-breaking developments with respect to reactor design but also with respect to reactor operation are necessary. To this end, UL is expected to conduct the setup of a digital twin of the first wall of a tokamak fusion reactor. This use case will not only contribute to the design and optimisation of this crucial part of a tokamak type fusion reactor by better explanation of the physical processes around the inner wall during the fusion reactions but also pave the way towards a real-time control of fusion reaction, which is needed to achieve stable fusion.

In this use case, the setup of a digital twin of the first wall of a tokamak fusion reactor is targeted. Determining the real temperature at the surface of the first wall of the tokamak from infra-red (IR) cameras is a crucial task for appropriate engineering design of the first wall. During the operation of the fusion reactor, the goal is to detect overheating parts from temperature images in real time and thus being able to react before these parts of the reactor are damaged. To this end, the performance of the engineering parts of the tokamak, dependent on the reactor geometry and the temperature distribution in the system, are modelled with field line tracing for the magnetically conducted particles (using L2G software [2]), which contribute most to the power deposition to the first wall. Then, a thermal model using Finite Element or Finite Volume Methods is employed for the wall, i.e., via OpenFOAM. The corresponding camera signal is then calculated by performing an optical simulation with Raysect [3], taking into account the radiation from the first wall due to the temperature distribution obtained from the thermal modelling. The results will be validated using experimental data that is available from ITER first plasma experiments and WEST [4] (CEA tokamak), along with synthetic diagnostic that combines IR and visible light camera images. Experimental data from ITER will consist of two different components. In the beginning of the ITER operation a simplified component without active cooling will be installed, which is named temporary limiter. Then this simplified component will be replaced by more complex components with active cooling, i.e., first wall panels. Experimental data from WEST will consist of actively cooled components only. An example of a predicted camera output of the WEST tokamak is shown in Figure 1.

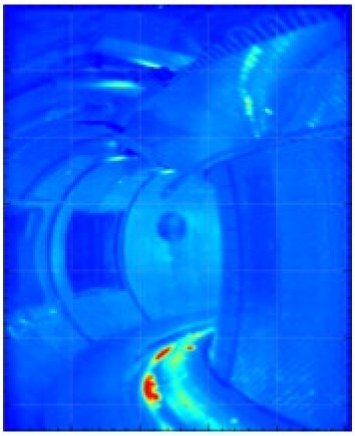

Figure 1: Synthetic camera image of WEST tokamak during operation.

2 Objectives of the Use Case

In a nuclear fusion reactor, particles that are conducted along the magnetic field lines deposit large amounts of power to the wall of the reactor. The resulting heat flux heats up the wall of the reactor, while critically high temperatures can result in melting and deformation of the wall. IR cameras are thus installed in the reactor to detect rising wall temperatures. However, the camera sensors can be affected by particles arriving from other parts of the reactor as a consequence of surface reflection, while some photons might also arrive directly from the plasma itself. This results in false positive detected hot spots, i.e., the camera detects a temperature rise that is not present in the reactor. The goal of the use case is to establish a simulation workflow to obtain synthetic camera signals, resulting in a digital twin of the fusion reactor, which can be directly compared to the experimental data to distinguish between real and false hot spots. The simulations to produce a digital twin are thus used to correctly describe the visible light in cameras that see not only surface temperature but also reflected light from other surfaces. When reflections are subtracted, correct surface temperatures are determined. Therefore, only complete simulations (digital twins) can give correct results to be seen by cameras. The ray-tracing and thermal modelling of the cooled surfaces in a complete fusion reactor are exascale relevant simulations that are expected to achieve good scalability. The inverse problem when having a real image from a camera for which surface temperatures are to be determined with the knowledge that the surfaces are reflective, can be only solved having previously performed forward simulations of the problem and taking some AI correlation algorithms into account to determine the correct heating scenario. The latter case is an essential requirement for real-time control of the engineering system. The algorithms used for the calculation of a digital twin are well suited to benefit from parallel processing and hardware acceleration. For this problem FPGA and accelerators are thus a good testbed for real use cases.

It is the ambition of UL as the owner of the further application FA-1 use case to advance beyond the state of the art with the setup of a digital twin and consequently the engineering design of the first wall of a tokamak type fusion reactor. This will be achieved by improving the coupled application of the community codes OpenFOAM and Raysect towards highly detailed simulation setups and more efficient workflow execution to acquire simulation results that will better explain physical processes around the inner wall during fusion reactions.

UL is to execute a smaller scale use case targeting the Slovenian peta-scale system VEGA and then move to large scale simulations based on the presented codes. In terms of fulfilling the objectives of EXCELLERAT P2, algorithmic optimisations are planned as well, such as the porting of the field line and ray tracing codes to GPUs. For thermal modelling, the focus is also on saving the results using the HDF5 file format. The code will then be deployed to different petascale systems. UL will prepare and run the simulations by porting the workflow and moving it to large scale. Thereafter, significant workflow improvements will be pursued to reduce time to solution and work towards the solution of the real time problem of the digital twin. To this end, UL, who was not part of the consortium during EXCELLERAT P1, but was informed about the work done via its participation in the interest groups, will join the consortium via the EXCELLERAT on-boarding, use case development and use case monitoring processes. These processes ensure proper integration of UL or any other new consortium member into the established EXCELLERAT P2 consortium and include assessment of the use case requirements, use case development in the respective tasks of the technical WPs as well as knowledge and competence exchange and progress monitoring.

3 Workflow Description

The wall of a fusion reactor receives the most power from particles conducted along magnetic field lines. There are other contributors (neutrons, photons, …), however their effect to thermal loading is negligible compared to magnetically conducted particles. The particles heat up the wall due to the deposited power, which can be predicted by magnetic field line tracing. The heat load can be converted into a temperature distribution through an appropriate thermal model, which considers material properties such as heat capacity, conductivity and density of the wall components. The effect of actively cooled components with appropriate heat transfer coefficients for the coolant must also be included. Once temperatures are computed the optical simulations can take place. Photonic particles are generated on the sensor of the IR camera and traced backwards. Reflections are calculated at wall intersections and radiation caused by the heated walls needs to be incorporated.
The accurate assessment of synthetic images thus requires three principal numerical simulation steps whose execution needs to be performed in sequence. That is, the output of the field line tracing solver is an input to the thermal model, that in turn provides the input to the optical simulation yielding results in the form of synthetic measurements, i.e., images. This simple scheme is depicted in Figure 2.

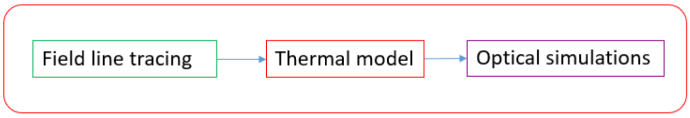

Figure 2: Sequential execution of codes for digital twin.

For the field line tracing the algorithms checks for an intersection of the field line with the wall, which is given as a triangular mesh. The wall consists of multiple panels which are further divided into smaller components attached to the beam (see panel in middle figure in Figure 3). There is no direct thermal contact between these smallest components. They are only connected by the beam, which experiences only small amounts of heat and heat conduction through the beam is neglected, such that each small component can be analysed independently from each other. For rather simple first wall geometries, i.e, components with simple shapes without much detail or components which are not actively cooled, the assessment of heat fluxes and temperatures is expected to be delivered faster than for more complex shapes such as actively cooled components or different castellations. However, for optical simulations a full temperature distribution is needed in order to give accurate results. Along with large input and output data volumes, there is a significant challenge for an appropriate coordination of the computational resources, i.e., the utilisation of different hardware and queuing systems.

For preliminary simulations and post-processing, it is planned to use the VEGA system through development access mode. As mentioned above, the described modelling work has been in development at the University of Ljubljana and work has progressed according to the schedule defined in deliverable D2.1, even though the executed simulations are still limited in size and the produced results are of lab type character. Therefore, UL will follow the classical task path:

- getting access to a peta-scale system (VEGA),
- deploying code and data to the system,
- preparing and running the simulations by porting the workflow and moving it to larger scales,
- evaluating and validating the results.

The goal of UL is to achieve significant workflow improvements to reduce time to solution and work towards the solution of the real time problem. The workflow is depicted in Figure 3:

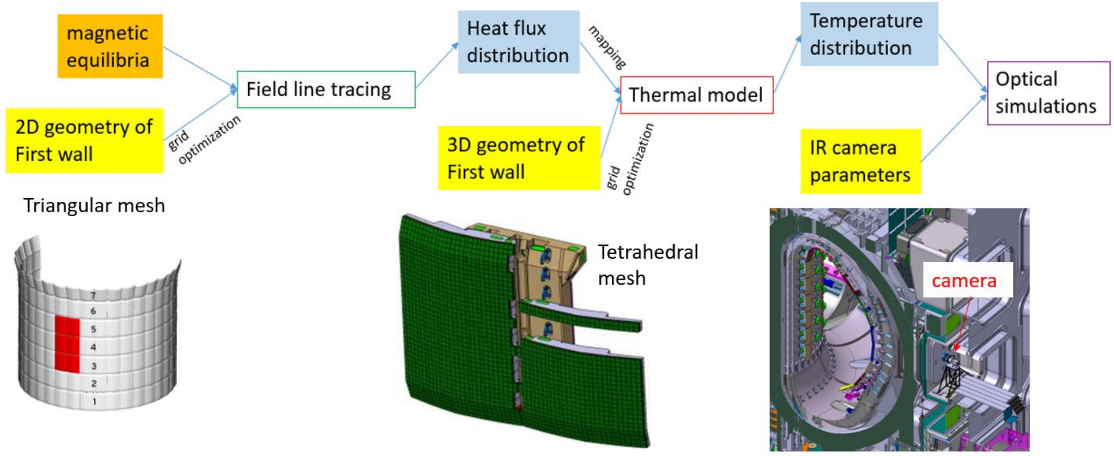

Figure 3: Workflow components for the simulation of IR camera.

For the field line tracing simulation, the magnetic equilibria inputs (information about the magnetic field, magnetic flux, plasma power etc.) will be supplied to UL from partners at ITER and CEA. UL has a collaboration agreement with both parties for the length of the EXCELLERAT P2 project and has access to the databases containing these inputs. For the thermal modelling, CAD models have already been supplied to UL. For the optical simulation, the current available camera design will also be supplied by CEA, which is responsible for the design of the camera in ITER. The camera parameters will include camera position and geometry, with a focus on the CCD sensor and aperture, including corresponding field of view, and spectrum range. Additional parameters will be reflectivity coefficients to properly simulate reflections. To optimise the workflow the inputs and output data between different simulation steps will be stored in appropriate parallel data formats (HDF5, netCFD, …).

4 Progress achieved

For the field line tracing codes, the code L2G [2] developed at UL and ITER has been optimised for better computational efficiency. The algorithmic improvements of the code are:

- Octree space partitioning
- Bounding box partitioning of large meshes to decrease time of line triangle intersection checks

The parallelisation of the field line tracing is also achieved by dividing the trace into multiple smaller traces for each field line as shown in Figure 5. Each field line trace is performed by one of the parallel processes, which checks for intersections with all triangles. The mesh is divided into a “target” mesh, i.e., the mesh to analyse, and a “shadow” mesh, which shadows the target mesh (see Figure 4a). A field line trace is launched in reverse, i.e., it starts from a triangle in the target mesh and is traced back towards the plasma. At each step the intersection is checked with all the triangles in the shadow mesh (see Figure 4b).

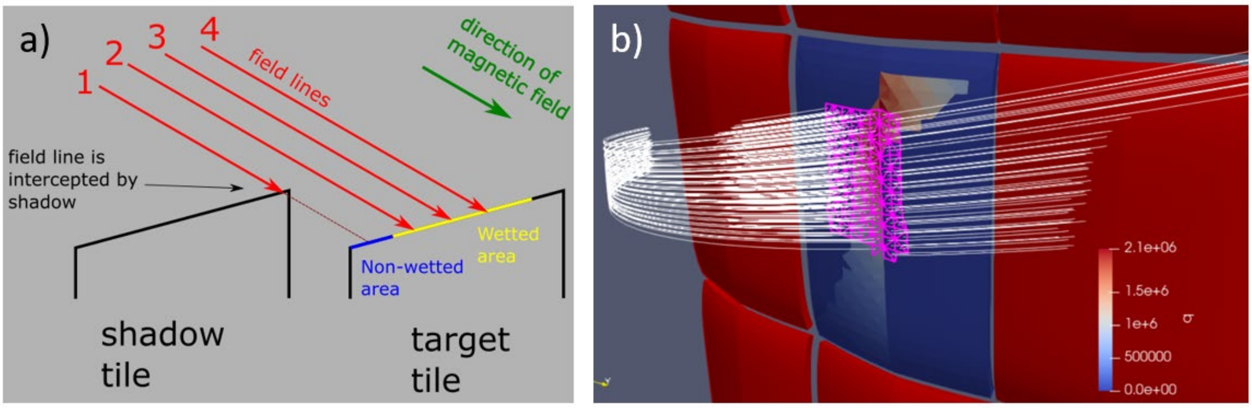

Figure 4: (a) Principle of shadow and target mesh. Field line 1 is intersected by a shadowing tile and does thus not affect the target tile (non-wetted - i.e., this part of target is shadowed by the shadow tile). Field lines 2-4 reach the target and deposit power (b) shadow mesh (red panels) surrounding target mesh, i.e., the part of the tokamak that we want to analyse. Field lines (white) are launched from each triangle and checked for intersections with the shadow mesh.

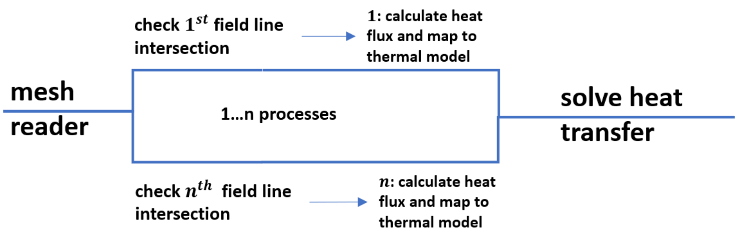

Figure 5: Partition of field line tracing.

The parallelisation currently works on CPUs, the goal is to port it to GPUs next. The cases in the L2G module are run with 1 OpenMP thread by default.

A benchmark case, named *Inres1* was run using 1, 2, 4, 8 and 16 OpenMP threads. For performance tests on huge meshes the target and shadow geometries from the Inres1 case were remeshed (using the SMESH module in SALOME) to element sizes of approximately 10, 3, 2 and 1 mm. The finest element size resulted in a target mesh of 3,046,416 and a shadow mesh of 25,309,694 triangles. This means that the algorithm launches 3 million traces and in each step each trace is checked for an intersection with 25 million triangles. However, with the algorithmic improvements described earlier the required number of checks is reduced significantly.

Figure 6 shows total execution times depending on the mesh size and number of OpenMP threads used. The results show that many threads execution speeds up field-line tracing on million size meshes for a factor more than 2 with the use of 1 MPI process with 16 threads.

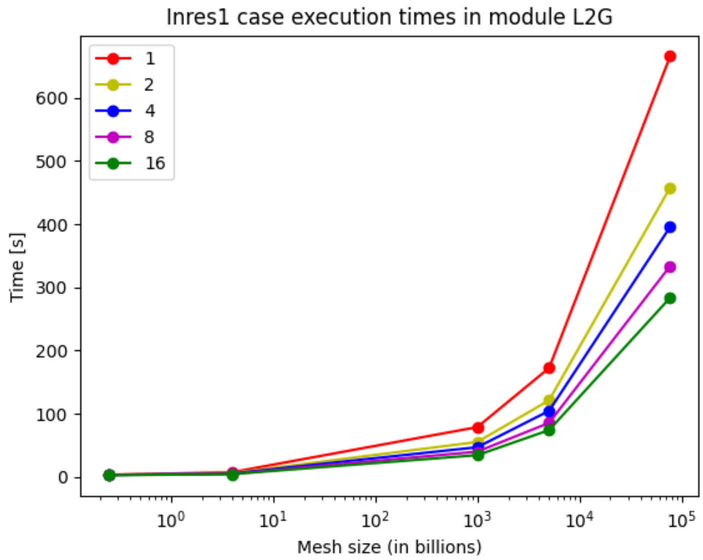

Figure 6: A comparison of execution times for different mesh sizes (triangular mesh). The legend refers to number of threads.

For the L2G code the output meshes and results are stored in Med file format (based on HDF5 library), which is the standard binary file format for storing results in SALOME open source simulation environment and used at UL extensively for pre and post processing [5].

For thermal modelling, the OpenFOAM package is used. The customised solver which uses the OpenFOAM kernel was written using the Swak4foam library. The solver maps the heat flux to the plasma facing surface. At other surfaces an adiabatic boundary condition is prescribed and the transient 3D diffusion equation is solved to obtain the surface temperatures.

A small-scale setup was prepared based on the ITER temporary limiter taking into account the workflow described above. One block of the temporary limiter was used for analysis discretised using 300k tetrahedrons (used in the thermal modelling simulation) and 230k triangles (used in the field line tracing simulation as target mesh) on the plasma facing component. The geometry to calculate intersections for the field line tracing (shadow mesh) consists of 18 segments lying in the R-Z plane. In ϕ-direction these segments are distributed at 4 locations (see Fig. 7). The shadow mesh has 16500k triangles. These 16500k triangles were checked for the intersection with each field line. The calculated heat flux was mapped to the 3D model consisting of 300k tetrahedrons and the heat equation was solved to obtain temperatures on the surface. IR camera signals were then calculated for a camera resolution of 640×512 pixels, field of view $21.7° \times 17.5°$ and pixel size of $15\mu m$ (see Figure 7). The case consists of 48 different slices with a time step size of $0.04s$.

The results of each simulation step are shown in Figure 6. First, on the left one block of the temporary limiter is shown along with a field line trace. In the middle, computed heat fluxes and corresponding temperatures on the surface are shown. At this stage, a simple ray-tracing simulation was performed with Raysect running on 64 cores. The pixelated image as seen by the camera is shown on the right, compared to the initial temperature distribution predicted by the thermal model. Finally, the reconstructed temperature is compared along the black cross-section line.

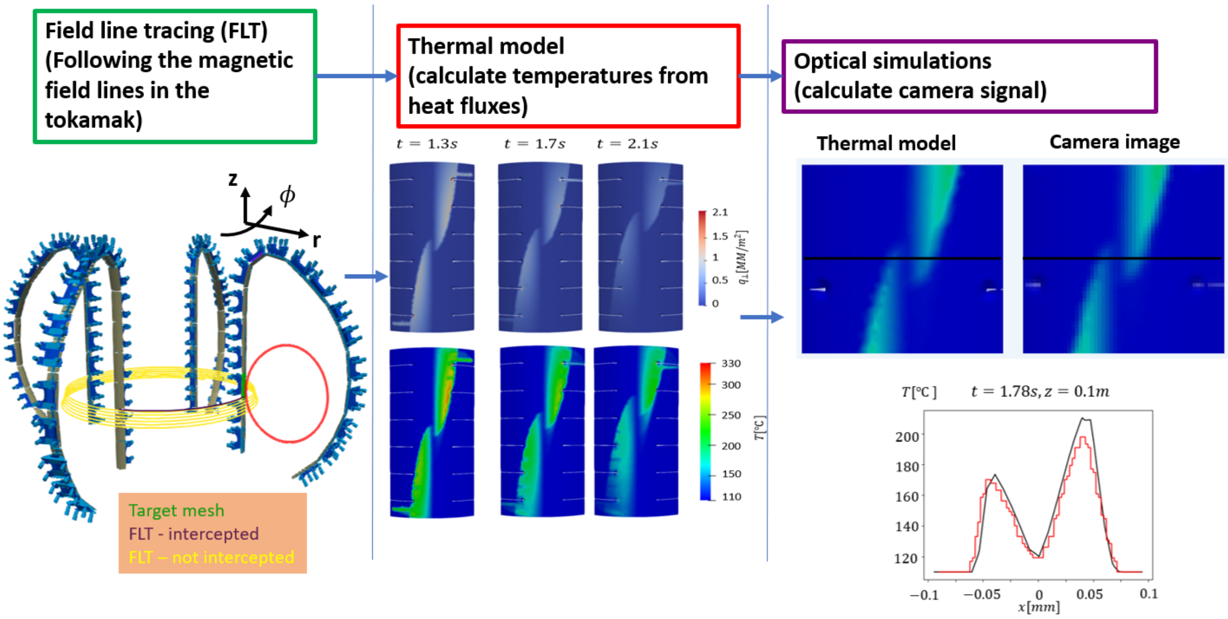

Figure 7: Small scale setup which form the basis for large scale.

As part of unit testing the field line tracing and thermal modelling, a few small cases have been prepared for benchmarking and automated testing. Access to Vega has been granted and access to LUMI, MareNostrum and Karolina is currently underway.

5 Conclusion

The goal for the second year is to increase scalability and move to more complex cases for ITER operation considering components more advanced than the temporary limiter. Currently, mesh preparation of these components for large scale is underway. The focus is on the 3D geometry of the ITER first wall which consists of circa 440 panels which cover in total around $600\ m^2$. Each panel consists of around 100 beryllium cubes, depending on the variant (position and configuration of the panel). In the current small case setup, the temporary limiter block is not actively cooled. The complex panels are actively cooled so additional heat transfer coefficient boundary conditions will be added to the case. The goal is to have a fully prepared case till the end of next year. Then in the following two years the focus will be on optimisation. In the second year, some algorithmic optimisations are planned as well. It is planned that the porting of the field line and ray tracing codes to GPUs will start. For thermal modelling, the focus is also on saving the results to HDF5 file format. In the second year, the code will be deployed to different petascale systems. UL will prepare and run the simulations by porting the workflow and moving it to large scale. Subsequently, significant workflow improvements will be pursued in order to reduce time to solution and work towards solving the real-time problem of the digital twin.

6 References

[1] "ITER - the way to new energy," ITER. Accessed: Oct. 13, 2022. [Online]. Available: http://www.iter.org
[2] Gregor Simic, "Enhancements and applications of the SMITER magnetic field line tracing and heat load mapping code package," presented at the 30th IEEE Symposium on Fusion Engineering, Oxford, UK, 7 2023.
[3] Carr Matthew and Alex Meakins, "Raysect Python Ray-tracing package." [Zenodo]. Available: https://doi.org/10.5281/zenodo.1205064
[4] IRFM CEA, "WEST tokamak." Accessed: Jan. 06, 2023. [Online]. Available: http://irfm.cea.fr/en/west/
[5] L. Kos, R. A. Pitts, G. Simič, M. Brank, H. Anand, and W. Arter, "SMITER: A field-line tracing environment for ITER," *Fusion Eng. Des.*, vol. 146, pp. 1796–1800, Sep. 2019, doi: 10.1016/j.fusengdes.2019.03.037.