CCFE
CULHAM CENTRE FOR FUSION ENERGY

UK Atomic Energy Authority

UKAEA-CCFE-PR(22)24

Alex Valentine, Thomas Berry, Steven Bradnam,
James Hagues, James Hodson

Benchmarking of emergent radiation transport codes for fusion neutronics applications

Benchmarking of emergent radiation transport codes for fusion neutronics applications

Alex Valentine, Thomas Berry, Steven Bradnam, James Hagues, James Hodson

This is a preprint of a paper submitted for publication in
Fusion Engineering and Design

Benchmarking of emergent radiation transport codes for fusion neutronics applications

A.Valentine[a], T.Berry[a], S.Bradnam[a], J.Hagues[a], J.Hodson[a]

[a]*Culham Centre for Fusion Energy, Culham Science Centre, Abingdon, OX14 3DB, Oxford, UK*

Abstract

The accurate and efficient mapping of the radiation environment in a nuclear fusion reactor requires the most advanced radiation transport tools. The Monte Carlo method has long been deployed to deal with the complexity of fusion relevant geometries, with MCNP the adopted industry standard code among the European and wider international community. However, reliance on a single code has driven explorations into alternatives to establish their capabilities and maturity for fusion analyses. It is imperative that the transport codes meet: i) stringent modelling and analysis requirements for fusion, and ii) may be used within an integrated engineering design workflow that can support ITER, DEMO, STEP analysis as well as existing experimental devices such as JET and MAST. The radiation transport codes, Serpent, OpenMC as well as the framework for allowing CAD based particle transport, DAGMC, are being being actively developed and increasingly adopted in some types of applied analysis by the user community. In this paper, we explore both experimental and computational benchmarks in order to examine the code capabilities for over a broad range of fusion relevant nuclear responses and geometries. This spans from more simple parametric models adopted in reactor scoping studies to the current ITER reference model which has been successfully translated to Monte Carlo codes other than MCNP using an open source utility, csg2csg. An assessment for both CSG and CAD based workflows has been conducted as well as a hybrid approach combining the two. The FNG HCPB, Cu and a subset of the FNS experiments were also converted to Serpent and OpenMC input files for comparison of calculation to available experimental data. Good agreement was observed across all codes for the determined tritium production rates and activation foil measurements. Potentially more efficient workflows for complex tokamak models are detailed. For a heterogeneous model of JET octant 1, an optimised CAD based model in Serpent is over 50% faster than the MCNP CSG equivalent model. In the case of the generated OpenMC model of ITER, an order of magnitude reduction in simulation time, including a model loading time of the order of minutes, is reported. Such validation and benchmarking activities should in the future be integrated as part of a testing suite for which a basic framework has been demonstrated here. To conclude, the current limitations and required development are outlined as well as identifying where each code may specialise for a particular application. Based on this future work relevant to both the developers and user community is briefly discussed.

Keywords: DAGMC, ITER, JET, MCNP, Monte Carlo, neutronics, radiation transport, Serpent, SINBAD, validation, variance reduction

1. Introduction

The history of the Monte Carlo codes for performing radiation transport calculation dates back to the post second world war era where the dawn of computing paved way for the stochastic methods which lay at the heart of the Monte Carlo simulation to be put into practice. Applications now span medicine, particle accelerators, fission, radiation shielding and nuclear fusion. The physics and geometry of each application differs significantly, with each presenting its own unique set of challenges in understanding the radiation environment. For example, in nuclear fusion, the current most widely investigated technology globally is the tokamak, the toroidal/spherical template upon which the detailed engineering structure and materials are formed is not trivial to capture to a high level of fidelity in a particle transport code. Further, the neutron energy regime is also considerably different to nuclear fission, the domain where most application and experimental benchmarking of Monte Carlo codes with underlying nuclear data have been performed.

MCNP [1], which has been developed at Los Alamos National Laboratory (LANL) for over 60 years, is to date the most widely adopted radiation transport code among the global fusion neutronics community. It is the reference code for nuclear analysis of devices such as JET, ITER and DEMO. The qualification of the code relies on its validation in the relevant physics domains which has been demonstrated over its long history. For application to nuclear fusion, a radiation transport code must be validated to perform neutron-photon coupled transport calling point-wise cross section data; capture the problem geometry in all its complexity; accommodate complex plasma neutron source definitions; allow deployment in parallel on high performance computing architectures and support acceleration techniques required for deep shielding problems. These requirements are well documented in [2] and are the basis of the investigations presented in this paper.

Several emergent Monte Carlo codes are currently being explored for application to fusion neutronics analysis, as alter-

Email address: alex.valentine@ukaea.uk (A.Valentine)

natives to MCNP. Here we have explored the capabilities of Serpent [3] and OpenMC [4]. Both the constructive solid geometry (CSG) and CAD based transport modelling through the toolkit DAGMC have been explored. Experimental data serves as the benchmark for validation and verification that the physics of a problem is being accurately captured by a particle transport code. In nuclear fusion, owing to the development timeline of current technology, there is limited data available globally in relevant geometry and energy regimes. In the absence of experimental data, the results of alternative codes are compared to MCNP. The SINBAD (A Shielding Integral Benchmark Archive and Database) database [5], available through the NEA and RSICC, contains 31 fusion relevant shielding experiments, which are a collection of data sets performed across multiple research institutions. Included in the data set are spectra measurements, activation foil measurements and time-of-flight (TOF) spectra. Since 2008, there has been a concerted effort to conduct quality reviews of the benchmarks, most of which were conducted over 20 years ago. The Frascati Neutron Generator (FNG) series of experiments are one of the highest quality benchmarks available in the database. FNG was constructed in ENEA, Frascati in 1992 and is capable of operating in both pulse and continuous operation mode.

Here, both the FNG Copper experiment [6] performed in the years 2013-2015 and the FNG helium cooled pebble bed (HCPB) experiment [7] from 2005 were selected as suitable for conducting benchmark comparisons of CSG based models. The Fusion Neutron Source (FNS) located at the Japanese Atomic Energy Authority (JAEA) has also contributed several experiments to SINBAD over the course of 35 years using its pulsed operation DT source. Measurement data is available for both TOF neutron spectra and in-situ neutron and gamma-ray measurements which have been recorded at various different collimated scattering angles from the irradiated sample of variable thickness [8].

Section 2 gives a background on each of the studied transport codes. As well as demonstrating the above mentioned capabilities inherent to the code, the usability and code stability from installation to the process of performing a simulation using high performance computing (HPC) are fundamental to code uptake by the community which is commented on in section 3. The methodology for conducting both the experimental and computational benchmarks is outlined in section 5. For tokamak relevant benchmarking, we have explored the octamak model, a sector model with homogenised representation of major reactor components and single equatorial port plug. This is used to assess the fundamental capability of each code, serving as a pre-requisite for deployment on two complex analysis fusion neutronics analysis problems: a recently developed model of JET octant one (section 6.6) and the ITER (section 6.7) reference model. This provides detailed insight on the complete workflow involved with conducting analysis using each respective code for current most challenging problems faced in fusion neutronics. The critique formed on the outcomes of the diverse range of benchmarks performed (section 6) is used to formulate, in broad scope, a capability matrix for each of the codes. Outstanding development needs (section 8) as well as where future validation efforts should focus are outlined at the end of the paper.

2. Background

The most typical analysis workflow of a fusion reactor starts from an engineering CAD model which is translated to MCNP using a CAD conversion tool which writes the input file in constructive solid geometry (CSG) format. The CAD model contains many detailed features which either can not be translated to CSG or lead to a computationally inefficient transport model. Notably, splines and off-axis tori must be removed or redrawn. The simplification effort can account for a major portion of the analysis workflow (often >50%). As fusion enters an engineering era for prototype fusion plants, there is naturally an increasingly stringent requirement on the accuracy of models used in predicting the radiation environment which feed into plant safety, maintainability, lifetime and ultimately regulatory approval. As technology advances and computing resource becomes increasingly less restrictive, such bottlenecks relating to CAD simplification and necessary approximations become increasingly prohibitive depending on the scope of the performed analysis. This serves as an additional motive for continued investment into alternative codes and workflows.

As the CAD model preparation can account for a significant portion of the workflow, UKAEA invested some effort into the development of tools utilising the SpaceClaim API [9] which automate aspects of the simplification workflow. This tool suite includes methods for redrawing complex pipe networks (often splines and tori) and automatically building the reciprocal (void) space for CAD model. Some of the tools which aid in the diagnosis of geometry problems have been used in this work.

One alternative to adopting a multi-stage process from CAD to CSG model is to transport particles directly on a faceted version of the CAD model itself. The use of unstructured mesh and unstructured surface geometries eliminates the need to remove complex geometry shapes and provides a one to one mapping between the CAD geometry and the radiation transport model to acceptable tolerances for most fusion neutronics applications. One clear advantage of this is that the CAD model acts as the reference geometry, a significantly more flexible format than a potentially modified CSG text file.

DAGMC [10] is a toolkit which interfaces with several Monte Carlo codes supporting CAD based radiation transport. Once a faceted geometry file is produced, a Monte Carlo transport simulation can be performed with any one of the codes DAG is capable of interfacing with by pointing to the same geometry file. One alternative adopting this workflow is Serpent which in 2015-2016 was extended to support transport of unstructured surface meshes in the format of STL geometry files [11]. MCNP also supports unstructured meshes which has previously been used in fusion relevant analysis. This feature is however not looked into in this paper as previously outlined development needs [12] are anticipated to be addressed in the upcoming release of MCNP6.3 [13].

The Serpent Monte Carlo code has been developed out of a PhD project completed by Jaakko Leppänen of VTT Technical

Research Centre, Finland. The scope of the code has increased significantly since 2004 when the original purpose was group constant generation in nuclear fission problems. The capabilities of the code now extend to multi-physics and photon transport with a neutron-photon coupled transport mode added in 2015 [14]. As with MCNP and OpenMC, the particle tracking routine relies on a universe based CSG geometry description. With the availability of photon transport, several fusion relevant studies have been conducted demonstrating the potential of Serpent for fusion neutronics [15][16][17].

OpenMC is an open source Monte Carlo code which has been developed since 2012 at both Argonne National Laboratory (ANL) and Massachusetts institute of Technology (MIT). It is maintained via a publicly available git repository using modern day best practices for software versioning and quality control. In general, the open access nature of the code facilitates rapid integration of developments enabling deployability for fusion neutronics problems on a timescale which would be difficult to achieve with licenced software.

OpenMC is based on an underlying C++ solver with user input driven by extensible markup language files (.XML) files. A python API can be used to manage all aspects of the simulation. The completeness of the API should be stressed - from writing the input files to the post-processing of large data sets, this inherent code interface provides the analyst with a rich set of tools for streamlining the complete analysis cycle. The geometry .XML file describes the geometry in native CSG format and more recently, CAD-filled universes are supported using DAG-OpenMC (v0.13.0). Recent developments include coupled neutron-photon transport and the capability to track particles on tori. Both are fundamental for application to fusion neutronics.

A key benefit to the Monte Carlo method is that it is very well suited to parallelisation on high performance (HPC) computing architectures. MCNP, Serpent and OpenMC are capable of running in both MPI and OpenMP parallelism. One advantage of using threading (OpenMP parallelism) within a node is that little additional memory is required per thread (shared memory). Further, as data does not need to be broadcast and received between individual processes at certain rendezvous points, there is a time saving. In MPI processing, the allocated memory is duplicated between each process which can lead to bottlenecks for particularly memory intensive applications. All calculations have been performed using internal UKAEA Intel Xeon E5-2665 computing cluster with 32 CPU cores and two sockets per physical node.

As well as computing resource, variance reduction techniques are often a requirement in the Monte Carlo method to converge to a solution in deeply shielded regions. The techniques explored in this paper are all based on weight windows which are one example of population control methods based on a conceptual probability model. In any of these methods, the basic idea is that the simulation remains unbiased given that an assigned manipulated quantity termed the statistical 'weight' is preserved. In this way, very rare events which may have a significant influence on the calculated nuclear response of interest can be sampled with high frequency without biasing the results. Weight windows can be produced as a function of either the geometrical location or energy in the problem phase space, or indeed both. Further detail is beyond the scope of this paper and can be found in [18].

A capability to generate weight windows based on the response matrix method was recently developed in Serpent [19]. This approach uses an adaptive mesh which is split recursively until a user supplied density criterion is met. An importance function, which is the contribution of a given region of phase space to a particular tally is calculated on this mesh by conducting the adjoint transport problem. The neutron importance is inversely proportional to the weight window boundaries of the output weight window. This capability is only available for neutrons. In an alternative approach, ADVANTG [20] relies on the deterministic code DENOVO to determine the importance function. ADVANTG is one of the most commonly used programs for the generation of weight windows in MCNP. One final method we have explored is WWITER [21], a UKAEA developed method based on using the information from successive iterations to uniformly populate the geometry through optimisation of weight window boundaries. To assess the effectiveness of weight windows a quantity called the figure of merit (FOM) is calculated. This gives an indication of the computational efficiency through factoring the run time and the magnitude of uncertainty as FOM=$\frac{1}{\sigma^2 T}$, where σ is the variance and T the computing time in minutes.

3. Installation and code use

3.1. MCNP

The most recent code release of MCNP, MCNP6.2[1], available for a licence fee through RSICC and NEA data bank, has been used in this analysis. The cost is elevated for those requiring source code access. The export control of the code leads to restricted access for many on the grounds of security. Nonetheless, globally there are over 10,000 users with over 400 person years of invested development providing a wealth of experience and support.

MCNP is granted as a single user licence which must be administered within the licence conditions to handle permissions if installed on a computing cluster. User development of the code often requires specific knowledge of the code which often lacks modern programming practices. Many patches tailored towards application of the code to analysis of fusion systems have been developed by the neutronics community, often independently without independent verification or a clear route to unify contributions through the codes focal point at LANL to make them openly available. MCNP6.3 will include an update to modern Fortran 2018 standards for the majority of the code which will facilitate improved user development.

3.2. DAGMC

DAGMC relies on commercial software Cubit [22] for the core of its workflow. Preparation of the CAD file including any simplification and fixing can be performed in SpaceClaim or other CAD software capable of exporting an ACIS file. As

detailed in Figure 1, there are several specific steps performed in Cubit prior to exporting the DAG geometry in H5M format. The graveyard is a shell encasing all the geometry of the problem to mark its outer bounds, where particles are terminated. Material information is captured by assigning every part of the geometry to a group according to material. The imprinting and merging of surfaces allows the DAG geometry engine to work in the most optimal way capitilising on indexing in the structure tree as opposed to having perform a binary search for neighbouring facets. The process of imprinting creates a common surface interface between touching volumes. Merging then takes these surfaces and combines them into a single surface. If boundary conditions, for example, reflective boundaries in a sector model are required, they should also be assigned in Cubit. Caution must be taken to ensure that every surface that is part of the boundary is identified. The geometry is then exported with a user specified faceting tolerance specifying the maximum distance from a facet to the surface in the CAD model.

To check and resolve holes in the faceted geometry file, the make watertight and check watertight commands must be run. Materials are assigned according to the University of Winsconsin Unified Workflow (UW^2) whereby a material library is incorporated with the watertight geometry according the tags created via grouping the geometry in Cubit. This approach facilitates the creation of a H5M geometry file which includes material data, enabling all geometry and material data to be independent to the choice of radiation transport engine used.

Aside from issues related specifically to fixing the CAD model, the most potentially time consuming drawback of this workflow is if changes are required to the geometry. Once exported from Cubit, the geometry can not be modified. The complete problem geometry must also be converted together in a single H5M file, as opposed to in a modular fashion, by component for example. In the event that errors are uncovered when assigning materials or checking if the geometry is watertight, it is necessary to return to the CAD in Cubit or even SpaceClaim if major revision is required. The user must therefore be very attentive in performing each of the specific steps in Cubit. In some cases this overhead can be limited through the built in scripting functionality in Cubit to automate many of the steps.

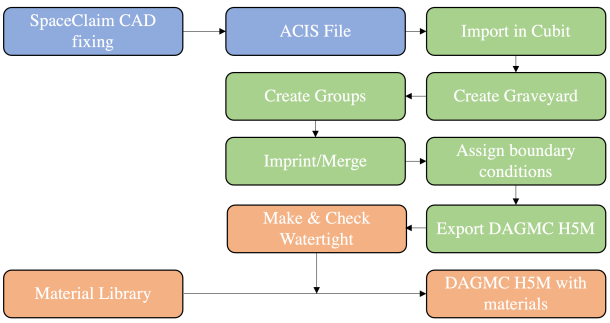

Figure 1: DAGMC workflow detailing the steps from CAD model to generation of the h5m file required for neutron transport.

It is has been previously documented that the installation procedure of DAGMC is not trivial. The large number of interconnecting code dependencies, which often require specific release versions, often leads to a time consuming compilation process. To arrive to a working DAG-MCNP executable required almost 1 week of effort with the help of one of the lead DAG developers. Additional effort is required to enable DAG-OpenMC. Some of the difficulties can be circumnavigated using docker or having root access if attempting to install on a HPC cluster. It is hoped that this process can be streamlined to make DAG more accessible to the community; many of these issues could be resoled with improved documentation and HPC compilation examples.

3.3. *Serpent*

Serpent is currently available through the NEA and RSICC with the current version used in this analysis, 2.1.32, available from February 2021. This is still a beta version of the code. There are many similarities to MCNP in the use of a single input file capturing all of the geometry, materials and simulation data. Also akin to MCNP, the code calls point-wise continuous energy cross section stored in A Compact ENDF (ACE) format.

One notable difference to MCNP is the codes use of delta tracking as opposed to the more conventional surface tracking. This can prove optimal in problems where the mean free path is large compared to the geometry dimensions. It is possible in Serpent to switch between surface and delta tracking [23] in Serpent, the most optimal choice of which is geometry dependent. One consequence of using delta tracking is that tallies must be recorded using a collision flux estimator (cfe) based on counting the number of physical collisions (analog) and/or additionally virtual (implicit) collisions in the response region. This generally performs well in dense material regions such as a fission reactor core (where particle collisions are frequent). Both MCNP and OpenMC instead use a track length estimator that integrates the length of all particle tracks through the tally region regardless of whether or not a collision occurs. A sensitivity study has been performed for these different parameters controlling the tracking routines in Serpent (section 6.5).

3.4. *OpenMC*

As an open source transport code, the OpenMC development process is transparent with complete visibility of the source code. This serves as a useful diagnostic for the novice user. In order to acquire the latest code features the development branch is needed, which is where current developments of the code are merged in the absence of an official code release. Static releases of the code are available however some of the code developments needed were made over the course of this research (features of versions 0.12.2 and 0.13.0 have been used). In collaborations this can make consistency in code versions difficult and also from a validation perspective, there should be means for demonstrating that each revision does not impact on the physics or otherwise incur compatibility issues with previously run problems. The same is also partially true for Serpent which also does not have an official release although the distribution method is more analogous to MCNP (static versions).

While OpenMC is in the development stages, a set of fusion relevant benchmarks such as those discussed in this paper should be automated to run in validation code performance and consistency prior to being used for production calculations. The majority of existing validation of OpenMC has been focused to the nuclear fission domain [24][25][26][27][28][29]. Section 4 outlines a starting framework for this which together with other validation efforts such as the development of the JADE [30], currently focused towards nuclear data validation, can provide a method to integrate automated code testing in actively developed codes.

To perform a transport simulation, a minimum of 3 '.XML' files are required which independently contain information on the system geometry, materials and settings. A path must be specified to the the cross section data which unlike Serpent and MCNP, OpenMC adopts HDF5 format. The python API can be used to convert between the two. The boundary, equivalent to the graveyard in MCNP where all particles are killed is defined on a surface in OpenMC which can be assigned as reflecting, transmissive or vacuum. Here the vaccuum boundary terminates all particles intersecting it, equivalent to an importnace zero region on the opposite side of the surface.

Contained within the settings.XML file is the execution settings for the problem, which must include number of particles and a definition of the source as mandatory arguments. For any nuclear responses in the problem, a separate 'tallies.XML' file is required. This is capable of capturing a broad range of responses including 'heating', 'H3-production' and 'damage-energy' which calculated displacements per atom (DPA). These are predefined to point to the relevant reaction MT numbers that are called explicitly in Serpent and MCNP. As for the geometry, materials and settings, the tallies XML file describing the nuclear responses can be produced through the Python API.

4. Automated validation

Transport codes that are under constant development must be validated to give confidence that any updates to the code have not retain its accuracy. The user should be able to perform this in an automated way and importantly, these benchmarks should be fusion relevant and represent the vast energy regime and breadth of shielding configurations unique to this application.

A framework has been established for the development of such a validation suite. The starting point for this is a code referred to as MCNP file writer. This python package parameterises the generation of CSG models, with the an MCNP input file constructed automatically based on the users definition of the problem. There are methods to write the geometry, materials, tallies and many of the physics captured in the data section of MCNP. The tool allows for the straightforward study on thousands of iterations of a problem.

Taking the MCNP output file(s) produced with MCNP file writer, a series of processing scripts automate the process of generating the different transport code input files with csg2csg and performing the transport calculation on a HPC system. This is performed over all tested materials/isotopes for each transport code for different nuclear data libraries. The output data is tabulated in CSV format and for visualisation, is also displayed graphically as in Figure 2. In this example, a 14 MeV point source is at the centre of a hollow sphere with an outer tallying region recording the neutron flux. The problem was described completely using the MCNP file writer and the comparison automated to give the ratio against the MCNP result. This scans all 190 isotopes in the FENDL-3.2a [31] nuclear data library released in 2021.

	1001	1002	1003	2003	2004	3006	3007	4009	5010	5011	6012	6013
Serpent	0.968677	1.00173	1.003532	0.997742	1.002246	1.00383	1.004571	0.990068	0.995985	0.99017	1.004884	0.992408
OpenMC	1.002276	1.000321	0.999323	0.999212	1.000469	1.000994	1.000122	0.999165	0.999821	1.000215	0.999481	0.999417
	7014	7015	8016	8017	8018	9019	10020	10021	10022	11023	12024	12025
Serpent	0.994832	1.001021	1.001995	0.999322	0.998087	0.998902	1.002055	1.002199	0.998686	0.997373	0.993831	1.002015
OpenMC	0.999539	0.99929	0.999382	0.999037	1.000609	0.999826	0.999604	0.99964	0.999622	1.00028	0.999729	1.000189
	12026	13027	14028	14029	14030	15031	16032	16033	16034	16036	17035	17037
Serpent	0.998281	0.996941	1.006474	1.006348	0.998538	1.002302	1.005428	1.003777	0.996771	0.989793	0.997278	0.999304
OpenMC	1.000172	0.999964	0.999522	1.000022	0.999918	1.000071	0.999586	1.00015	0.999884	0.999194	1.000037	0.999358
	18036	18038	18040	19039	19040	19041	20040	20042	20043	20044	20046	20048
Serpent	0.996239	0.995174	1.000789	1.001602	0.992829	1.004245	1.000469	0.994476	0.996744	0.995411	1.000943	0.996807
OpenMC	0.999683	0.999782	0.99979	0.999741	0.999809	0.999667	0.99973	0.999602	0.999115	0.999269	0.999297	0.99935
	21045	22046	22047	22048	22049	22050	23050	23051	24050	24052	24053	24054
Serpent	0.993507	0.995653	1.003106	1.002575	0.997075	0.997078	1.006689	0.99185	0.994272	1.003122	1.000804	1.00164
OpenMC	0.999532	0.999611	0.999137	0.999421	0.999476	0.99927	0.99911	0.999327	0.999468	0.999521	0.999749	0.999339
	25055	26054	26056	26057	26058	27059	28058	28060	28061	28062	28064	29063
Serpent	0.998516	0.993434	0.996945	0.995035	1.003451	1.006936	0.998371	1.00403	1.007424	0.996305	0.999578	1.001037
OpenMC	0.999273	0.999604	1.000003	0.999787	0.998879	0.999262	0.999484	0.999687	0.999211	0.999532	0.998965	0.999572
	29065	30064	30066	30067	30068	30070	31069	31071	32070	32072	32073	32074
Serpent	0.995049	1.006625	1.006652	0.991578	0.997629	0.996811	1.006775	0.992807	1.004636	1.000116	0.993928	1.000669
OpenMC	0.999118	0.999512	0.999664	0.999931	0.999583	1.000045	0.999272	0.999275	0.999865	0.999263	1.000036	0.999946
	32076	35079	35081	39089	40090	40091	40092	40094	40096	41093	42092	42094
Serpent	0.999499	0.998782	0.995895	0.996774	1.002046	0.996191	1.001873	1.003539	1.003275	1.002878	0.998837	1.000421
OpenMC	1.00021	0.999593	0.999647	0.999529	0.999821	0.999968	1.000173	1.000127	1.000303	1.000286	1.000097	1.000497
	42095	42096	42097	42098	42100	45103	47107	47109	48106	48108	48110	48111
Serpent	0.996499	0.999237	0.998851	1.001835	0.99758	0.997173	0.987845	1.000377	1.003218	1.003046	0.989773	1.001133
OpenMC	1.000127	1.000639	1.000186	1.00005	0.999712	0.999858	1.000501	1.000232	0.999779	0.999577	1.000112	0.999934
	48112	48113	48114	48116	50112	50114	50115	50116	50117	50118	50119	50120
Serpent	1.004925	1.003414	1.002826	0.997169	1.000934	1.004563	1.000216	1.002504	0.996062	0.993304	1.00513	1.006419
OpenMC	1.000345	1.000251	1.000115	1.00041	0.999866	1.000234	1.000338	1.000461	1.000509	1.000499	1.000707	1.000617
	50122	50124	51121	51123	53127	55133	56130	56132	56134	56135	56136	56137
Serpent	1.002061	0.997107	1.001107	0.997801	1.00885	1.000425	0.99739	0.999791	1.003129	1.002568	1.000976	0.993218
OpenMC	1.000735	1.000648	0.999607	0.999448	1.000189	1.000212	0.99928	0.999676	1.000047	1.000322	1.000321	1.000611
	56138	57138	57139	58136	58138	58140	58142	62144	62147	62148	62149	62150
Serpent	1.001447	1.001437	1.001879	1.003183	1.004913	1.003503	0.998101	0.997868	0.999122	1.003506	0.997127	0.98919
OpenMC	1.000592	1.00057	1.000365	1.000433	1.000393	1.000819	1.000618	0.999876	1.00047	1.000698	1.000319	1.000339
	62152	62154	64152	64154	64155	64156	64157	64158	64160	62148	62149	62150
Serpent	1.000646	1.000505	0.99662	0.999126	1.000408	1.000363	1.005971	1.011194	1.006801	1.003506	0.997127	0.98919
OpenMC	0.999986	1.00028	1.000325	1.000017	1.000185	1.000046	1.000135	1.00019	1.00023	1.000698	1.000319	1.000339
	62152	62154	64152	64154	64155	64156	64157	64158	64160	68162	68164	68166
Serpent	1.000646	1.000505	0.99662	0.999126	1.000408	1.000363	1.005971	1.011194	1.006801	0.999093	1.000217	0.998284
OpenMC	0.999986	1.00028	1.000325	1.000017	1.000185	1.000046	1.000135	1.00019	1.00023	0.999602	1.000159	1.000435
	68167	68168	68170	71175	71176	72174	72176	72177	72178	72179	72180	73180
Serpent	0.998564	1.007916	0.991749	1.003144	1.004012	1.001955	1.000364	0.992779	0.999583	0.995657	0.999738	1.003148
OpenMC	1.000436	1.000682	1.000473	1.000099	1.000299	0.999973	1.00021	1.000562	1.000074	1.000227	1.000105	1.000377
	73181	74180	74182	74183	74184	74186	75185	75187	78190	78192	78194	78195
Serpent	1.005452	0.997728	1.003631	0.998174	0.999363	1.008651	1.003742	1.007023	0.992395	1.00444	0.996176	0.994709
OpenMC	1.000778	1.000339	1.000144	1.000751	1.000243	1.000558	1.000295	1.000463	0.999622	0.999623	1.000184	1.000099
	78196	78198	79197	82204	82206	82207	82208	83209	90232	92234	92235	92238
Serpent	1.002263	0.993897	0.997348	0.988491	0.992858	1.004875	0.995167	0.993385	1.003249	0.991614	0.992631	0.996187
OpenMC	1.000098	0.999877	1.000066	0.999786	0.999935	0.999912	0.999994	0.999875	1.00136	1.000399	1.000989	1.000309

Figure 2: Example output of the post-processed data for a sphere leakage test, tallying the neutron flux calculated by MCNP, OpenMC and Serpent for all isotopes present in the FENDL-3.2a library. The ratio of the result in Serpent and OpenMC to MCNP is given for each isotope. The colour is scaled based on minimum and maximum differences to MCNP (red).

So that the test is fair across isotopes, the density of the spheres was scaled based on the flux in ^{184}W. This is equivalent to altering the thickness of the sphere such that it is roughly equal to the same number of neutron mean free paths in each material. In the near future, thus will be extended to include different geometries and experimental data sets. At present there is functionality to automate any geometry and nuclear response input in MCNP file writer. This front end will be substituted by a database of benchmark experiments. There is also capability to call different nuclear data libraries. This gives an efficient and accurate way to validate both transport code and nuclear data libraries.

5. Methodology

5.1. Conversion between transport models

The Serpent and OpenMC input models in CSG format have all been produced using csg2csg [32], which translates the geometry and materials between the transport codes MCNP, Serpent, OpenMC, PHITS and FLUKA. Recent developments to this tool permit translation of macrobodies, cone surfaces and processing of duplicate surfaces allowing conversion of complex geometries. The converted files do not include a data section i.e. source terms, tallies and miscellaneous physics data which must be input by the user.

5.2. Octamak

The octamak contains a homogenised representation of major tokamak components including the first wall, blankets, divertor, vacuum vessel, port plug and superconducting magnets. The model spans 45° with lateral reflecting boundary conditions as an approximation to eight-fold toroidal symmetry of the reactor. Two variants of the octamak were studied. A baseline model with the above described components and material compositions reflecting those in JET and ITER, and a more detailed model with a revised equatorial port region in order to accommodate a detailed port plug geometry. This second model variant is shown in Figure 3.

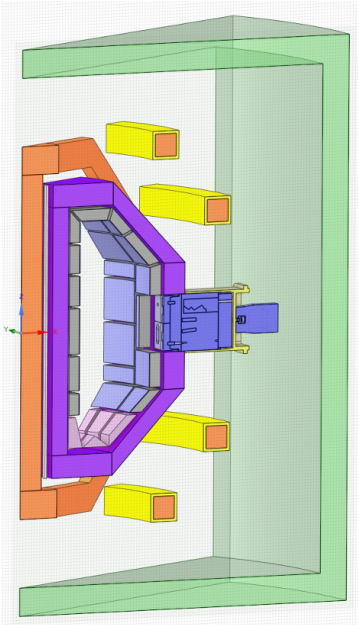

Figure 3: CAD model of Octamak 45° sector model with detailed port geometry as detailed in Figure 4

The purpose of this second model is to test the ability of the codes to handle complex geometries akin to a reactor design, and provide a more accurate comparison of the effort required to produce a simulation ready transport model from an engineering CAD model. The port plug is a representative ITER dummy port plug with three drawers occupied by diagnostic shield modules (DSM) and various channels, mirrors and optical components for plasma viewing systems. Preparation of the radiation transport model for conversion to CSG required 4-5 days of simplification effort. Most of the optical pathways are spline surfaces which must be redrawn to conform to a mathematical surface definition supported in CSG whilst also retaining an accurate representation of the radiation streaming pathway. Figure 4 shows the unmodified engineering port plug CAD model. Highlighted in figure is all of the surfaces in the geometry that require redrawing (splines or higher order problematic surfaces), identified with one tool developed with the SpaceClaim API. Not shown are the many regions of interference that also had to be resolved.

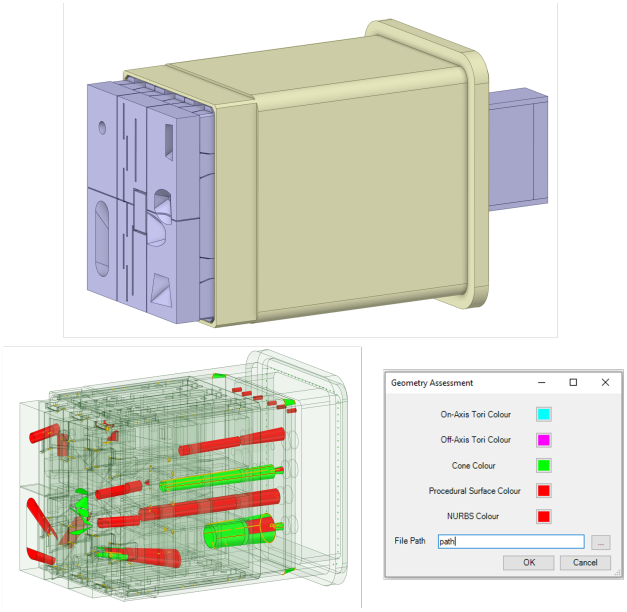

Figure 4: CAD model of original ITER dummy port plug model. The bottom image shows complex surfaces that have to be simplified prior to translation to CSG as identified by an API developed tool.

With the codes adopting a modular structure through universes, it is possible to create a hyrbid CSG-CAD input model. This can be advantageous for many reasons, for example, a model may be largely homogenised comprising simple shapes in which case a CSG representation is more sensible, however, a specific study may be required in a particular region of the tokamak for which a high fidelity model is needed. Such a workflow is also useful in closed loop engineering design analysis where components are designed, assessed and improved. This approach was tested through creating a hybrid CSG (major reactor components) and CAD (detailed port plug model) in Serpent. A universe envelope was created to conform to the port plug dimensions and the CAD model exported as STL.

The CAD model as in Figure 3 was also converted to a DAG-MCNP model. Starting from the engineering CAD model, the time taken to produce a standalone model of the port plug in Serpent STL format was 1-2 hours. At least half of this time is in the removal of interference's which is a mandatory requirement also for conversion to CSG. Additional time (1 hour) was required to create a geometry that could be faceted in Cubit, largely related to very small gaps in the geometry resolved through merging smaller bodies into single larger bodies. Con-

trary to preparation of a model for conversion to CSG, large complex bodies can be faceted more optimally for computational performance than small simple geometrical shapes. One example of a common issue that was encountered is shown in Figure 5. This was reported as a 'degenerate facet' in Serpent and also caused several meshing issues in Cubit. Where a degenerate facet occurs, two of the three points defining the triangular facets coincide, inducing a fatal geometry error which must be fixed in the CAD file. Once identified, the issue could be resolved simply by nudging the surfaces by 0.01 cm (negligible effect on the transport). Note that the above reported model preparation times correspond to a transport model that passes Serpent STL geometry inspection methods ('checkstl') and in the case of DAGMC is watertight, examined through the 'check watertight command'. Simulation of 10^8 particles with all materials assigned as void yields no lost particles for point isotropic source. The preparation of faceted CAD models is elaborated on in section 6.6 for the more complex case of the JET sector model.

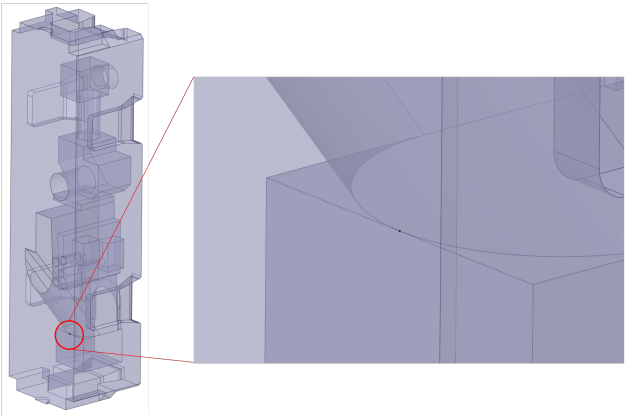

Figure 5: Example of geometry problematic for both Cubit and Serpent. This is reported as a 'degenerate facet' due to three of the triangular vertices coinciding.

The baseline model was likewise validated via the same method to be free from geometry errors. The assessed nuclear responses cover those typically calculated in current applied fusion reactor nuclear analysis. This includes the on load flux spectrum, nuclear heating and displacements per atom (DPA). Each quantity was evaluated for the 40 blanket modules and 16 divertor tiles in the octamak model. The spectra and nuclear heating values were calculated for both neutrons and photons in blanket cells adjacent to the port, through running each code in coupled neutron-photon transport mode.

Both the assessed baseline and modified port variation of the octamak contain significant shielding such that variance reduction is required to converge ex-vessel nuclear responses. There is no opening through the vacuum vessel at the upper or lower level and the equatorial port plug has few streaming channels. Several methods have been explored for both the baseline and revised geometry including ADVANTG, the Serpent response matrix method and WWITER as reported in section 6.5. The response of the poloidal field coils (PFC) and two cell tallies in the port interspace, between the rear of the port plug and the bioshield, are used as a metric for the effectiveness of these methods. In the absence of this capability in OpenMC at the time of analysis, it was excluded from the assessment of ex-vessel responses.

5.3. Parametric reactor design

Through the development of the 'paramak' [33] tool, it is possible to automate the generation of paramaterised 3D CAD models and scan a broad range of a reactor design space. With input of a set of engineering/physics driven set of parameters, the user can define the complete constraints of the generated CAD model which can be output in STP or STL format for subsequent transport analysis, facilitating a purely code-driven analysis workflow. The package is built around CadQuery2 and is a completely open source project [34]. One application of the paramak is in the pre-conceptual design of STEP, a prototype reactor being developed at UKAEA with operations scheduled for 2040.

This workflow has been validated for a generated DAG-OpenMC CAD based transport model and compared to a DAG-MCNP model by examining in-vessel responses, namely, tritium breeding ratio (TBR), flux and nuclear heating. The neutron spectra in 175 energy groups were also calculated in each component. Figure 6 shows one of the several iterations of generated 3D models containing a central column with shielding and a representative breeder blanket structure. The radial build profile is all driven by paramak input parameters.

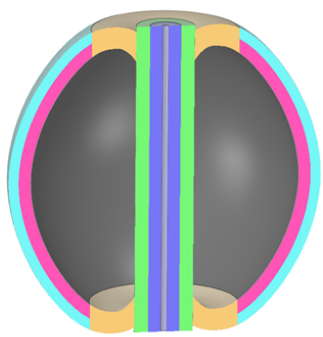

Figure 6: Vertical slice through the centre of a 3D reactor CAD geometry generated with the paramak.

For this assessment, OpenMC 0.12.1 results are compared to MCNPv6.2. All calculations were performed to 1×10^7 histories, sufficient for converged responses in this simple model. DAG-OpenMC and DAG-MCNP simulations were performed in neutron-photon mode while MCNP. Cubit converts the CAD geometries into faceted mesh geometries using specified faceting and merge tolerances. A faceting tolerance of 1×10^{-2} and a merge tolerance of 1×10^{-4} were used. Material definitions were incorporated within the DAGMC geometry file using the workflow as described in section 1. As a result, an identical DAGMC geometry file which included material definitions was used for both the DAG-OpenMC and DAG-MCNP calculations.

The faceting tolerance dictates the accuracy to which the faceted geometry conforms to the 'real' geometry and its impact on results was investigated in a separate study using paramak. Two separate geometries were used for this assessment - the single-null helium cooled pebble bed (HCPB) ball reactor and single-Null HCPB ball reactor with 8 equatorial ports of radius 100cm. A 2D plot of an example faceted geometry file which is produced through paramak is shown in Figure 7.

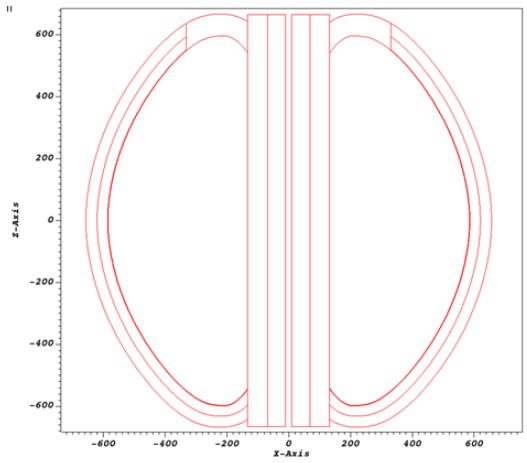

Figure 7: Paramak generated faceted mesh (1 cm tolerance) for the single-null HCPB ball reactor.

The faceting tolerances were changed between 5×10^{-5} and 1.0. A comparison was made with an MCNP CSG representation of the geometry generated with SuperMC [35] which acts as the benchmark (see section 6.4.2).

5.4. SINBAD database

5.4.1. FNG HCPB

The FNG HCPB experimental configuration is shown in Figure 8. The total tritium production was determined in a series of Li_2CO_3 pellet stacks at increasing distance from the source, nominally referred to as ENEA 2, 4, 6 and respectively. Through the centre of the mock up, a series of activation foils are used for determining the following reaction rates: ^{27}Al(n,α), ^{197}Au(n,γ), ^{93}Nb(n,2n) and ^{58}Ni(n,p).

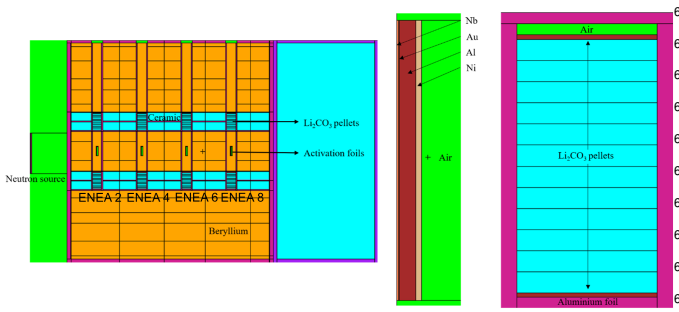

Figure 8: MCNP model of the FNG HCPB mock-up. The activation foil configuration (centre) and the arrangement of Li_2CO_3 in the breeder pellet stacks (right) is shown.

A standard methodology to validate a converted model is to perform a stochastic volume calculation for all material cells in the geometry. In MCNP, a spherical source is constructed enclosing the geometry with particle weight equal to the πr^2 and a flux F4 tally with all material cells. Both Serpent and OpenMC contain methods called on the command line (*-checkvolumes* and *openmc.VolumeCalculation* respectively) to determine both cell and material volumes. As this is performed stochastically, an error value is provided for each reported volume. It was verified in the case of the HCPB geometry that no single cell deviated greater than 3 standard deviations from the MCNP reported volume.

Several versions of the FNG source term are available that can be called in MCNP, both in standard SDEF format and as a programmed source routine that is called in the input file. Other Monte Carlo codes do not at present contain as advanced capability as the SDEF card in MCNP, capable of capturing complex energy, angular and time dependency of source terms. To produce a comparable source term for Serpent and OpenMC, it is necessary to either re-write the routine in the native code language or create an external library which wraps the source code. A version of the source has been re-written in C for use in Serpent which was used for some of the results presented in the FNG Cu case. However a validated FNG source for OpenMC was not available. For this reason, we have prioritised consistency between codes by using a 14 MeV point source positioned at the centre of target cell in the assembly. The approximation to the experimental set up inevitably introduces a systematic deviation from experimental responses which has been characterised. The nuclear data libraries FENDL3.1d [31] and JEFF3.3 [36] were used in all SINBAD benchmark experiments for neutron transport. In cases where OpenMC is included in the comparison, the transport cross sections are also used in determination of response functions for consistency between codes - difficulties were encountered in creating a mixed HDF5 cross section library and pointing to specific cross sections which could relate to specific reaction channels. It is understood that the resulting absolute responses will be less accurate, however again, the emphasis is on consistent validation between transport codes. All simulations were performed to 1×10^9 histories with no application of variance reduction.

5.4.2. FNG Cu

As for the HCPB mock up, the input files in Serpent and OpenMC format were automatically converted through csg2csg using the MCNP input distributed with SINBAD. The conversion process was validated by cross volume comparison as for the HCPB assembly. The irradiated copper block consists of $60 \times 60 \times 70$ cm^3 oxygen free copper with average density of 8.7982 g cm^{-3}. Cylindrical copper rods holding the activation foils are inserted at 8 positions in the block. ^{93}Nb(n,2n) and ^{186}W(n,γ) reaction rates have been determined at each of the foil positions as shown in Figure 9.

The tungsten foils are 25 μm thick and 18 mm diameter, while niobium foils are 1 mm thickness and also 18 mm diameter. It was necessary to use variance reduction in Serpent simulations in order to reduce the uncertainty on the recorded reac-

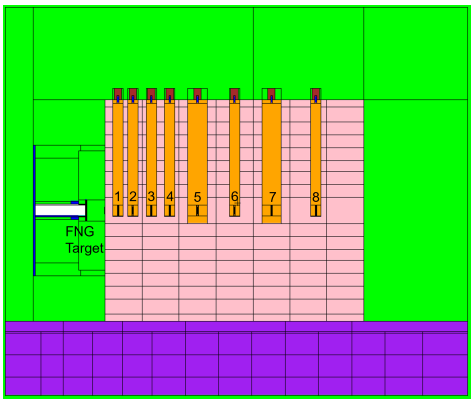

Figure 9: MCNP model of the FNG Cu experimental assembly at X=0 with the 8 foil capsules held in copper rods at increasing distance from the source.

tion rate for the Tungsten foils. For both MCNP and OpenMC, analogue calculations were acceptable with integral reaction rates recording <5% relative error.

5.4.3. *FNS experiments*

The FNS experimental set up can be seen in Figure 10. The collimators are positioned at 0°, 12.2°, 24.9°, 41.8° and 66.8° with respect to the slab, with the measurements recorded at a distance of 723 cm from the sample material. The DT source is positioned 20 cm in front of the first surface of material.

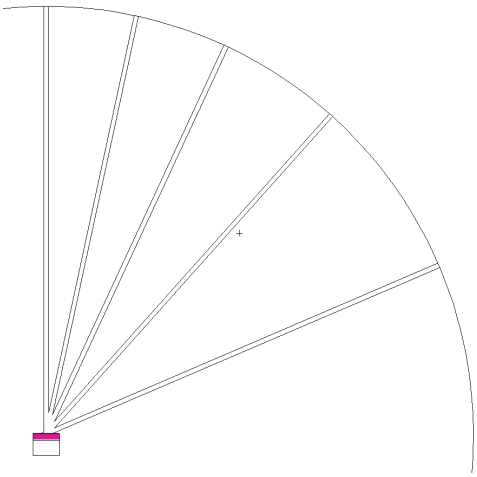

Figure 10: Plot of the FNS MCNP model consisting of a material sample (purple) and series of collimators at different angles

For this experiment, the comparison is made between OpenMC and and MCNP only. With the validated csg2csg converted input files, one difficulty was encountered in representing the importance zero region that is assigned in MCNP between each of the collimators ('infinite absorber'). As boundary conditions are assigned to surfaces in OpenMC this led to an issue at the opening of each collimator tube whereby all particles were being terminated. The geometry was modified to bound the opening of the collimator tubes by a spherical surface.

In this analysis, beryllium and iron were selected for the material block. The source distribution was taken from the SDEF card distributed with the FNS benchmark. The emission probability in each energy bin was written as a source term in OpenMC with isotropic spatial distribution. Cell tallies were created in each of the spherical cells representing the detectors, located in each of the collimator tubes to score neutron spectrum in 175 energy groups at each scattering angle individually.

OpenMC calculations were run to 1×10^9 histories. In this analysis, the thickness of material was selected to be 50 mm - larger block sizes require application of variance reduction techniques.

6. Results and Analysis

6.1. *FNG HCPB*

The determined tritium activity calling the FENDL3.1d library in each of the pellet stacks is given in Figure 11. In the evaluation of reaction rates, OpenMC folds in the material density automatically. The cell tally volume needed for the track length estimator must however be included in post-processing.

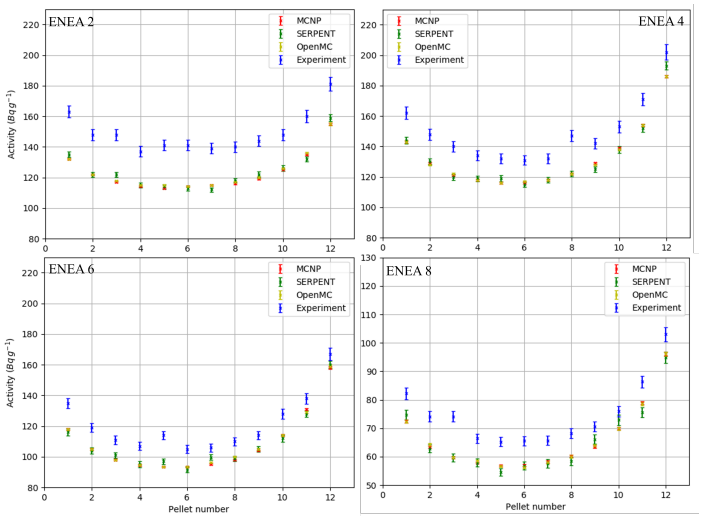

Figure 11: Comparison of Serpent, OpenMC and MCNP evaluations of the activity of tritium in Li_2CO_3 pellets in the HCPB mock up

Across each of the 12 pellets for ENEA 2,4,6,8, excellent agreement is seen between MCNP, Serpent and OpenMC with all calculated results within the statistical error. The under prediction of the experimental results is consistent with what has been previously reported with MCNP with a more representative source term. The serpent results can likely be improved by including a track length estimator tally in order to reduce the uncertainty on the result.

In Figure 12, the ratio of the OpenMC, Serpent and MCNP results against experiment is plotted for each foil through the experimental assembly. Differences to experiment stem from both the source term approximation and the use of transport cross sections to determine reaction rates. Where possible, dosimetry libraries such as IRDFFv1.05 [37] and IRDFF-II [38] should be used. However, consistency is observed across transport codes for each reaction at each depth through the assembly. For the aluminium and nickel foils furthest from the source,

there is some discrepancy with the Serpent calculated response - once more, enforcing the track length estimator would likely correct for this relatively high uncertainty result.

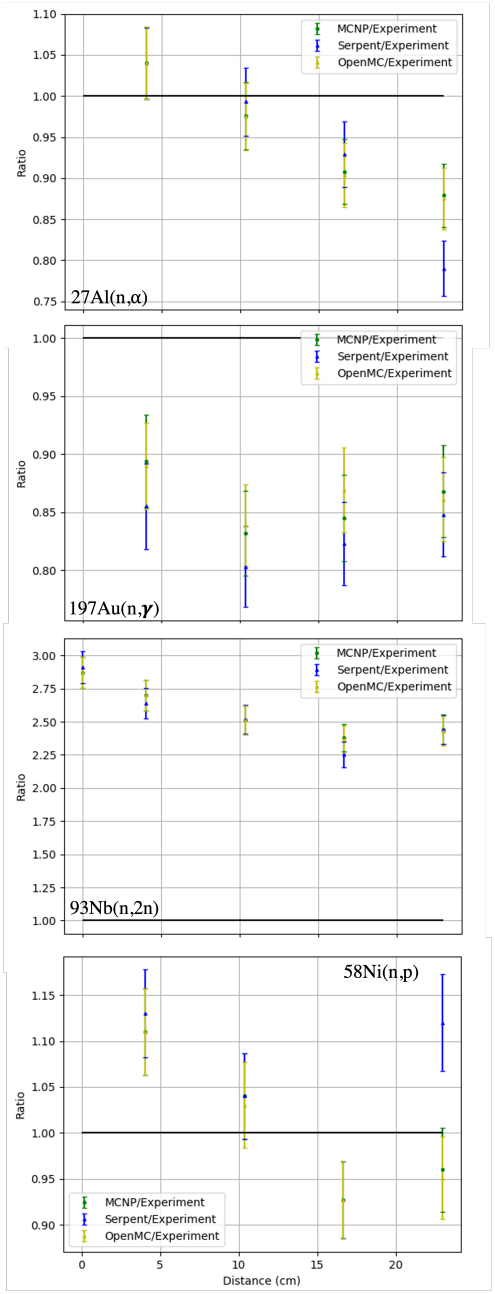

Figure 12: Ratio of the calculated reaction rates in FNG HCPB for Al, Au, Nb and Ni, comparing Serpent, OpenMC and MCNP to experiment.

The neutron flux over the entire experimental assembly was recorded in $2 \times 2 \times 2$ cm^3 mesh voxels. For Serpent and OpenMC the agreement with MCNP over all mesh voxels in the experimental region is very good (Figure 13). The relative error in each voxel is below 5% in all mesh voxels in MCNP and OpenMC however this increases to 25% at the rear of the mock up behind the experimental assembly.

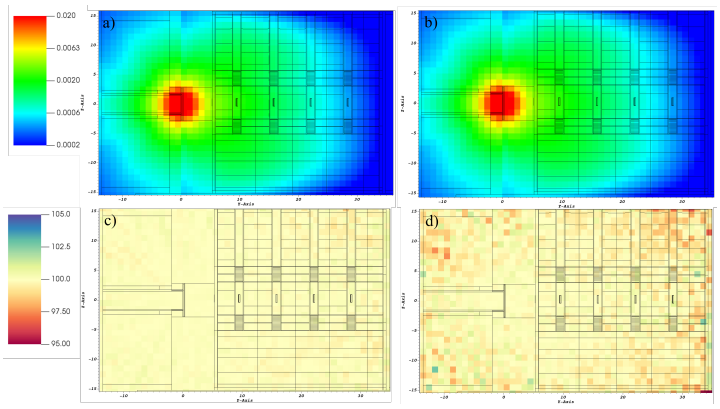

Figure 13: Neutron flux (cm^{-2} per source particle) over the FNG HCPB assembly for OpenMC (a) and Serpent (b). The ratio expressed as a percentage of OpenMC/MCNP (c) and Serpent/MCNP (d) is also given.

6.2. FNG Cu

For the available source terms: SDEF ENEA, SDEF JSI and the source code routine, the calculated reaction rates for ^{186}W(n,γ) and ^{92}Nb(n,2n) are shown in Figure 14. Good agreement is seen for niobium with some discrepancy for deeper tungsten foils with higher associated error.

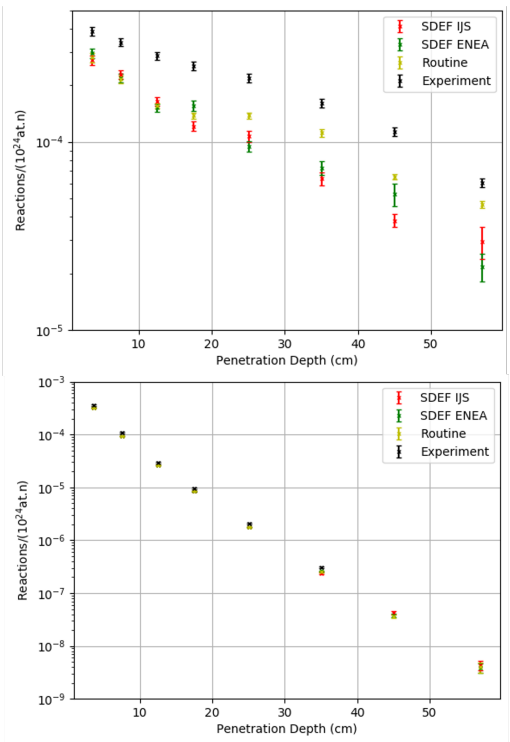

Figure 14: Reaction rate for ^{186}W(n,γ) (top) and ^{92}Nb(n,2n) (bottom) for different FNG source terms in MCNP. Reactions are given in units of number of reactions per unit volume/(10^{24}*source neutrons).

In Serpent calculations both global and targeted weight windows for specific foil responses have been used. With application of the global weight window, the relative error on the tungsten results for which the error is highest is reduced from 27% to 7% in the foil furthest from the source. The adaptive

mesh method in the global scheme proved optimal for convergence of results in all foils. Taking the global weight window, a weight window was generated pointing to the furthest tungsten foil, still further reducing the relative error to 2%. In this instance the track length estimator was also enforced, which accounted for ˜few % reduction in the error. To inspect the weight window, the neutron importance's over a Cartesian mesh were plotted, shown for both the global and targeted scheme in Figure 15.

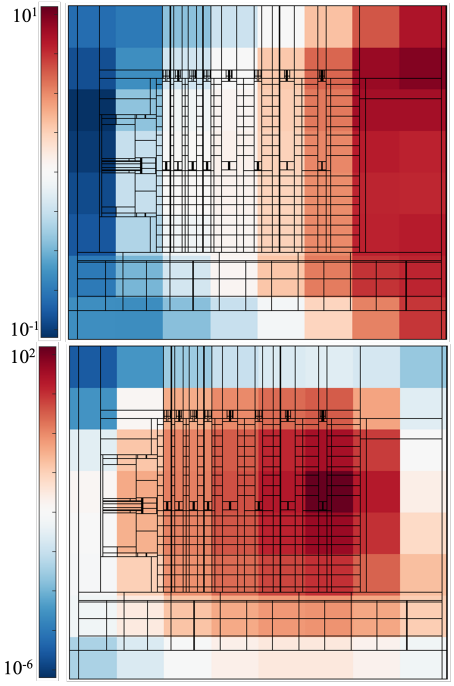

Figure 15: Plot of the logarithmic neutron importance for a Serpent generated global (top) and targeted (bottom) weight window for the foils furthest from the source in FNG Cu. The global weight window uses an adaptive mesh.

The calculated reaction rates for both the analog simulation and with the applied global weight window in Serpent, together with MCNP analog results for the SDEF ENEA source routine are plotted in Figure 16. The rewritten source routine for Serpent was used. The comparison is presented for FENDL3.1d transport data and IRDFFv1.05 for the reaction rate cross sections.

The C/E values for niobium vary from 0.75 to 0.86 with all calculated results within one standard deviation. For tungsten, the C/E values determined by Serpent with a global weight window vary from 0.4 (foil 8) to 0.8 (foil 1). As demonstrated for other SINBAD benchmarks, the reaction channel cross sections are well characterised and it is unlikely that there is an error in experimental measurement. The differences likely originate from the copper cross section data which require re-evaluation. The results have also prompted further characterisation of the density of the tungsten foils in the experiment which is currently ongoing. There is strong dependence of self-shielding effects on this quantity and implicitly the reaction rate.

MCNP calculations were repeated with a point source and compared to OpenMC calculated reaction rates as plotted in

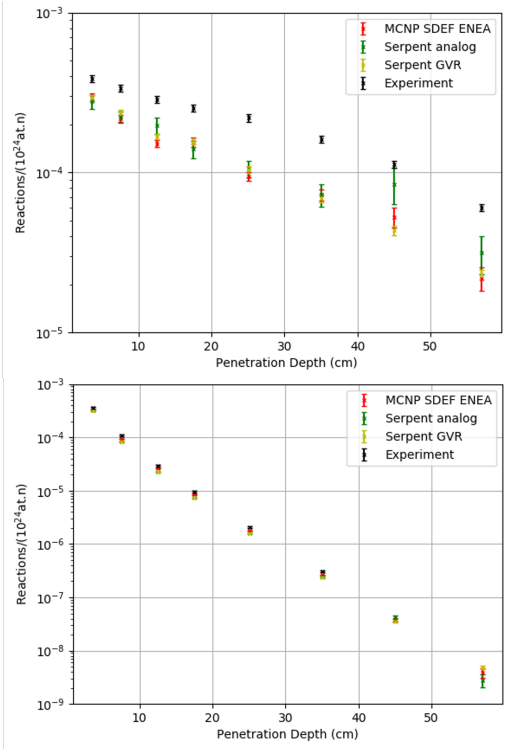

Figure 16: Reaction rate for ^{186}W(n,γ) (top) and ^{92}Nb(n,2n) (bottom) comparing Serpent in both analog and non-analog modes against MCNP with SDEF ENEA source. Reactions are given in units of number of reactions per unit volume/(10^{24}*source neutrons)

Figure 17 for both JEFF3.3 and FENDL3.1d data libraries. In general the effect of using the point source is on average a 5% under prediction of the Nb and W reaction rates. Nonetheless, the important comparison here is between codes which are consistent for both nuclear data libraries in under predicting the experimental data as reported above.

6.3. FNS

The calculated neutron spectra in MCNP and OpenMC as a function of collimator angle in the FNS experimental set up is shown in Figure 18 and Figure 19 for Beryllium and Iron respectively. The ratio of calculated to experimental result and between calculated results is also shown.

For beryllium, the most significant deviation is between the OpenMC results and experiment at large scattering angles. At both 41.8° and 66.8° there is a large amount of fluctuation in the neutron spectrum usually related to high statistical error. In Iron, there is also large differences for OpenMC at 66.8°, however in this case the calculation is more consistent in under predicting the measured flux spectra. Interestingly, the agreement is much better at 44.8° in iron, while there is on the other hand a clear over prediction at both 24.9° and 0°. In spite of differences in the absolute data, the trend is in general consistent - a harder neutron spectrum is observed in beryllium owing to its lower angular scattering cross section and the resonances present for inelastic scattering in iron above 1 MeV. The differences in OpenMC are almost certainly the result of the modification that was required to the geometry at the mouth of each

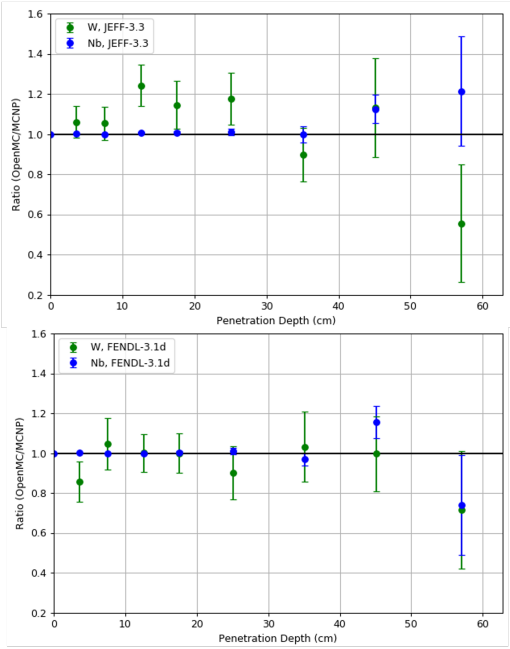

Figure 17: Ratio of OpenMC to MCNP calculated reaction rates for W and Nb in FNG Cu. This comparison is shown for JEFF3.3 (top) and FENDL3.1d (bottom) used for both transport and reaction rate evaluations.

collimator to handle the issue related to the definition of boundary conditions. The MCNP calculated results in general are very consistent with the experimental data for each scattering angle in both beryllium and iron.

6.4. Parametric CAD model

6.4.1. Nuclear responses

In general, very good agreement is found between the DAG-OpenMC and the DAG-MCNP results. The results for TBR and integral neutron flux are within 0.5%. The calculated neutron spectra broadly fall within one standard deviation. Results have been normalised assuming a 1004 MW reactor. Neutron spectra for the first wall, breeder blanket, divertor and central column are given in Figure 20.

One important aspect of calculating the nuclear heating due to photons in OpenMC is the decoupling between photons and electrons. For a result comparable to MCNP and Serpent, the individual contributions of photons, electrons and positrons must be summed. In this case, tallying only the contribution from photons accounted for a 30-70% difference in the nuclear heating. The neutron and total nuclear heating as tabulated in Table 1 and Table 2 demonstrate excellent agreement.

6.4.2. Study of faceting tolerances

For varying the faceting tolerance between 5×10^{-5} and 1, it was seen that there is negligible impact on the TBR with a maximum deviation of 0.007%. A difference is however noted in the total nuclear heating in the first wall where 7% difference is found. The impact is less (< 1%) in the blanket and rear blanket heating.

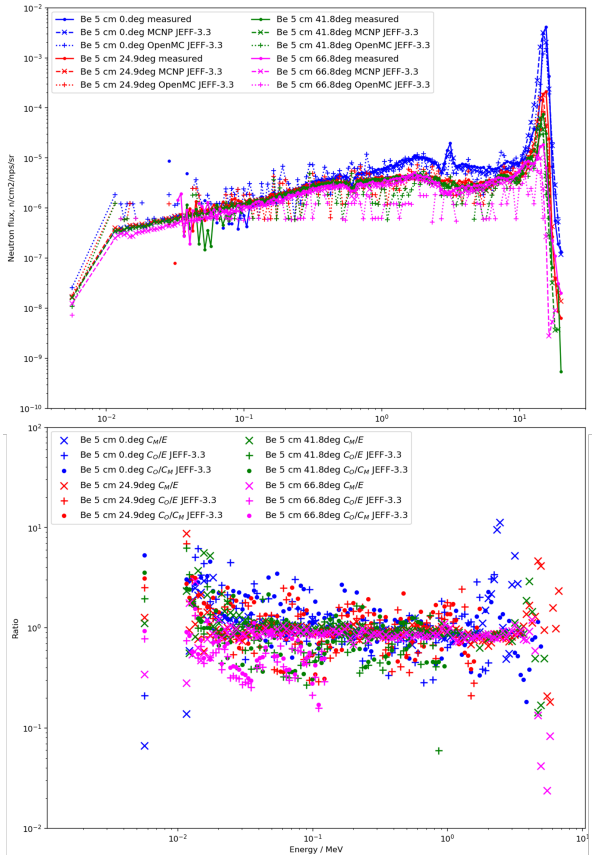

Figure 18: Neutron spectrum at each collimator angle for MCNP, OpenMC and measured data for the **beryllium** sample (top) and ratio of results (bottom). C_M, C_O and E refer the calculated results by MCNP, OpenMC and the experimental result respectively.

As evidenced in Figure 21, which plots the variation in TBR and first wall, breeder zone and rear blanket heating, there is a clear step change in the front wall heating between a faceting tolerance of 1×10^{-4} and 5×10^{-5}. This is seen for both baseline and importantly, a different model including ports. The reason for this step is not clear as a finer faceting tolerance is expected to better approximate the actual geometry. Examining individual particle track data may provide some insight. Further investigation is required from these results, however for such a low-fidelity model the faceting tolerance appears to make little difference until the result degrades with a step change as the model moves to a finer tolerance between 1×10^{-4} and 5×10^{-5}.

Region	OpenMC		MCNP		Ratio
Blanket	313	± 0.03%	313	± 0.02%	1
Rear Blanket	48	± 0.08%	48	± 0.04%	1
Centre Column Shield	24.7	± 0.04%	24.7	± 0.08%	0.999
Divertor	9.35	± 0.10%	9.35	± 0.11%	1
First Wall	11.9	± 0.06%	11.9	± 0.02%	1
Inboard TF Coils	0.02	± 3.97%	0.01	± 3.49%	1.05

Table 1: Neutron heating (MW) in reactor components calculated with DAG-MCNP and DAG-OpenMC for the paramak generated spherical reactor geometry.

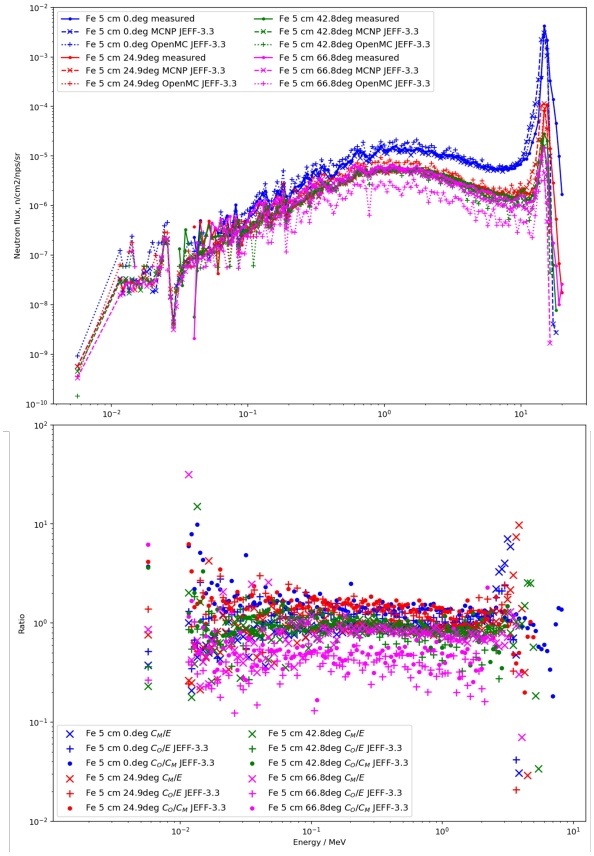

Figure 19: Neutron spectrum at each collimator angle for MCNP, OpenMC and measured data for the **iron** sample (top) and ratio of results (bottom). C_M, C_O and E refer the calculated results by MCNP, OpenMC and the experimental result respectively. [add error bars]

6.5. Octamak

6.5.1. Nuclear responses

Neutron-photon coupled simulations were performed for the MCNP, OpenMC and Serpent CSG models as well as DAG-MCNP and Serpent STL CAD based models of the octamak baseline. The coupled OpenMC simulations were performed with ENDF/B-VIII.0 [39] due to the aforementioned issues relating to cross sections which were found in the converted photon data for FENDL3.1d. Neutron results were in all cases for the FENDL3.1d.

For the octamak baseline geometry, the flux and heating are presented in Figure 22. The consistency with MCNP is general very good and within the statistical error across the 40 blanket

Region	OpenMC		MCNP		Ratio
Blanket	566	± 0.02%	566	± 0.02%	1
Rear Blanket	72.2	± 0.10%	72.1	± 0.06%	1.002
Central Column Shield	180	± 0.05%	178	± 0.06%	1.009
Divertor	62	± 0.13%	62.2	± 0.14%	0.997
First Wall	38.8	± 0.08%	38.8	± 0.05%	0.999
Inboard TF Coils	0.09	± 1.14%	0.08	± 2.37%	1.038

Table 2: Total nuclear heating in reactor components calculated with DAG-MCNP and DAG-OpenMC for the paramak generated spherical reactor geometry.

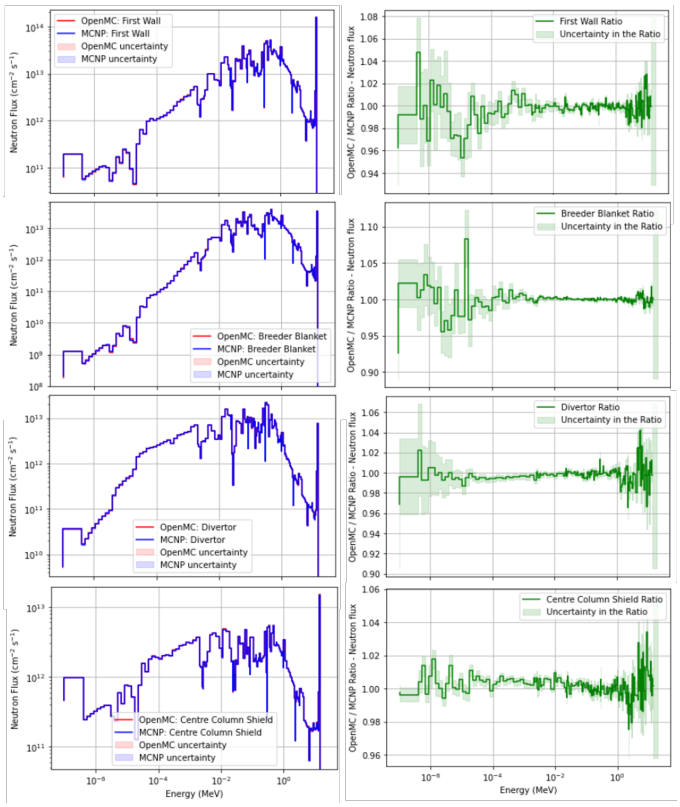

Figure 20: Comparison of calculated neutron spectra in DAG-MCNP and DAG-OpenMC for the first wall, breeder blanket, divertor and central column in a spherical paramak generated model.

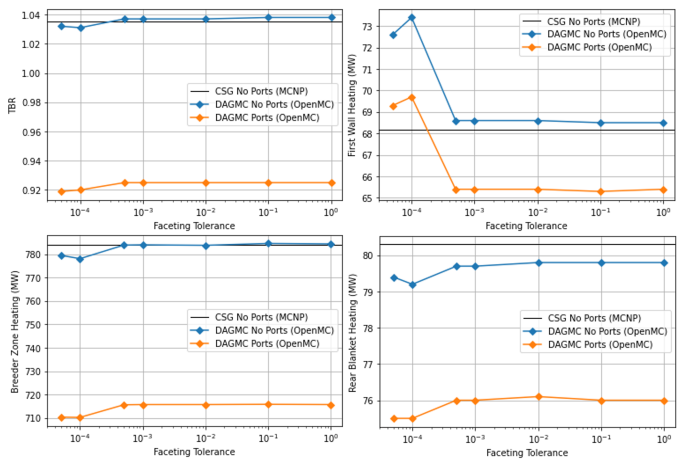

Figure 21: TBR and heating in different components of the blanket as a function of the faceting tolerance as defined in Cubit.

modules. As can be expected, DAG-MCNP gives almost identical results to MCNP CSG in all cases. Serpent and OpenMC appear to consistently under predict the neutron nuclear heating. The nuclear heating evaluations are based on KERMA coefficients that are extracted from the NJOY [40] HEATR module. Across all isotopes in the FENDL3.1d, 5 isotopes were found to have negative cross sections for the average heating numbers. Such issues have been reported , specifically related to negative KERMA factors in other transport libraries [41].

The method for handling incorrect negative data differs between transport codes. In this case, it may account for the observed differences to MCNP. It is not expected every single value to lie within one standard deviation - aside from systematic deviation relating to cross sections, the maximum difference across the responses is of the order of 2%. The neutron physics implementation between the codes is very similar however there are differences in the photon physics. The different cross section used in OpenMC accounts for a few % difference as observed. The comparison for the calculated DPA value in each of the divertor tiles is given in Figure 23. Inspection of MT 444 which stores damage energy data also uncovered 4 isotopes with negative cross sections for the FENDL library which may explain some of the observed differences, which are in any case within 1.2%.

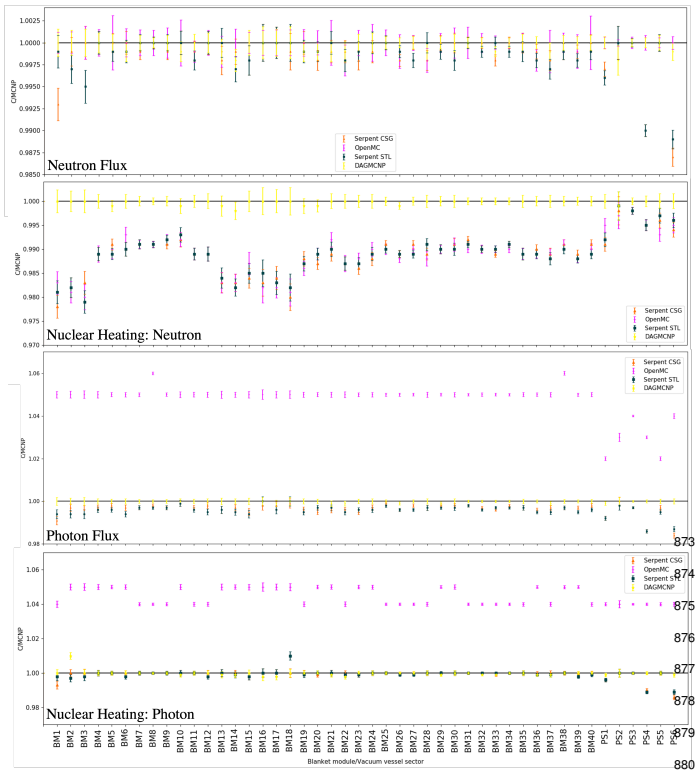

Figure 22: Ratio of the calculated flux and heating values different transport codes against MCNP in each of the blanket modules and vacuum vessel sectors for the octamak baseline geometry.

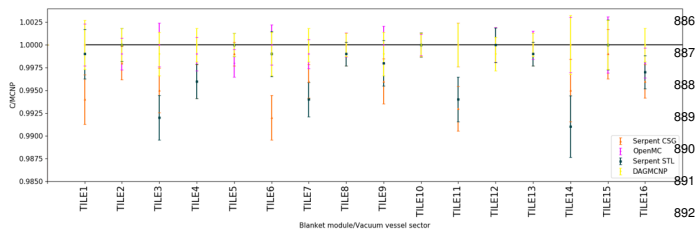

Figure 23: Ratio of the calculated DPA in each of the divertor tiles for different transport codes against MCNP for the octamak baseline geometry.

Initially, large errors and relative differences to other codes were encountered in Serpent, particularly in the case of photon flux spectra. This was understood to be the result of the collision flux estimator tally, the effect of which is more pronounced when scoring in energy bins. It is possible in Serpent to define a minimum mean distance for scoring collisions. The default value is 20 cm which was not changed for the Serpent STL calculation. Better agreement was found with the minimum distance, defined by the 'set cfe' parameter, set equal to 2 cm as plotted in Figure 24 which shows the photon spectra in 175 energy groups for a blanket cell adjacent to the port plug. In the thermal energy range, the poor statistics explains the observed differences.

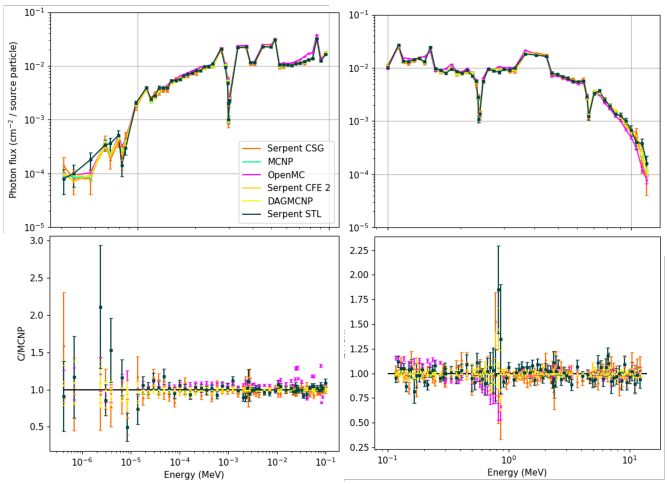

Figure 24: Comparison of the photon spectrum in an outboard blanket cell for both CSG and CAD based transport codes. The ratio to the MCNP result is shown at the bottom for each energy range.

The geometry also presents an interesting case for whether delta tracking as used in Serpent, that is theoretically more efficient where large void regions (large mean free path compared to problem dimensions) are present. By varying the 'set dt' parameter in Serpent, the ratio between the level of surface tracking and delta tracking was varied. The default value is 0.9. It was found that the shortest run time occurs for a value of 0.2 which was 3% faster than the longest run time for dt equal to 0.6. Although the octamak contains large void regions, around the blanket and vessel are regions where the neutron mean free path is short relative to the component dimensions hence in this case the optimal dt value is a balance between the two. The small difference does however suggest this is not significant.

As a related study on the effect of the collision flux estimator parameter, two different thicknesses of a tungsten block placed in nitrogen were studied with the neutron flux recorded in a cell of constant size on the opposite side of the tungsten cell. Although decreasing the value of cfe decreases the relative error on the result, the simulation time is increased. This results in a peak value of the figure of merit as shown in Figure 25. Evidently the optimal cfe value is strongly coupled to the number of particle collisions and must be selected on an individual geometry basis. The other option to improve tally convergence is to enforce a track length estimator for the tally. This was tested suing the default cfe value and gave a comparable figure

	Wall time (minutes)				
	MCNP	**DAGMC**	**S STL**	**S CSG**	**OpenMC**
N	58	110	344	180	110
NP	125	251	832	335	621

Table 3: Total elapsed simulation time for the octamak baseline geometry comparing both neutron only (N) and coupled neutron-photon (NP) modes.

of merit value to the optimal cfe case suggesting that if available, it is recommended to use this option.

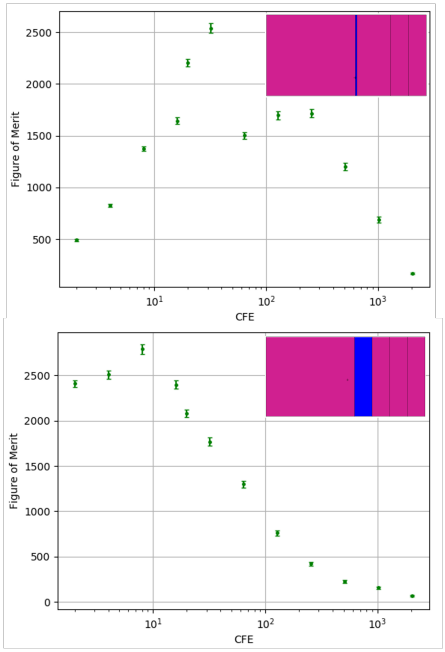

Figure 25: Study conducted in Serpent showing the figure of merit as function of the 'cfe' parameter. Two different thicknesses of tungsten (blue) were investigated.

All simulations were performed using 32 CPU cores in either MPI, OpenMP or hybrid MPI-OpenMP. This translates to the available resource on a single node of the UKAEA cumulus HPC cluster. In the hybrid case, it was possible to run 2 MPI processes across the two sockets within a node together with 32 threads. The optimal choice of these parallel methods varies between codes - only for Serpent was improvement observed in the hybrid case. For OpenMC, pure OpenMP recorded the shortest run time while for MCNP, pure MPI was optimal. Table 3 presents the shortest simulation time for 10^8 particles in both neutron transport only and coupled neutron-photon transport modes.

For this geometry the MCNP CSG model runs in the shortest time. Large transport times are found for the Serpent STL geometry, particularly in the coupled neutron photon scheme. It is likely that optimisation of the geometry can significantly reduce this as highlighted in section 6.6. Relative to the neutron only simulations, the largest increase in run time when running

Tolerance	# Facets	Memory (kbytes)	CPU minutes
0.00001	3523128	1503216	141592
0.0005	488010	524892	28978
0.001	347232	479508	23362
0.005	156970	416900	12904
0.01	110885	400076	12153
0.1	37159	375252	7972

Table 4: Effect of faceting tolerance on allocated memory and simulation time for DAG-MCNP octamak baseline model.

in a coupled mode is found with OpenMC (82% increase). The performance of Serpent and OpenMC CSG geometries against MCNP is likely related to the geometry structure. The baseline geometry includes only a top level universe containing all components. For Serpent and OpenMC, nested universes are more efficient for the tracking routines which is opposite to MCNP. There are caveats associated with this data presented in Table 3 however. The faceting tolerance of the DAGMC geometry has significant weighting on model run time (and memory use). The run time reported is for a tolerance of 10^{-2}, containing in total 400076 facets. Table 4 demonstrates how the faceting tolerance impacts the run time and allocated memory requirement for the octamak baseline geometry. This is a factor 3 times larger than the number of facets in the Serpent STL geometry. The DAGMC run time which is comparable to OpenMC is partly due the recently developments to the tracking routine which include the addition of the double-down interface [42] to Embree, an Intel developed ray tracing kernel [43].

Figure 26 gives the neutron flux in the analog simulation for the Serpent STL geometry and OpenMC. The impact of the collision flux estimator tally in Serpent is evident with much lower scoring outside of the vacuum vessel relative to OpenMC. The agreement with MCNP is shown to be consistent for in-vessel regions with only results having <50% included.

6.5.2. *Variance Reduction*

The most optimal weight window for a specific problem can take some trial and error to achieve. For Serpent, the response matrix method was used to generate a global weight in order to uniformly populate the geometry. This was performed for the variant model of the octamak comprising hybrid CSG with an integrated STL port model. A series of iterations are performed in which new data is collected in order to extend the mesh into deeper geometry regions. Originally, a 10 iteration cycle was trialled in Serpent however this was found to be excessive for a $5 \times 5 \times 5$ cm^3 neutron flux mesh tally. Figure 27 shows the neutron flux and neutron importance at each cycle of generating the weight window, demonstrating that after 3 iterations sufficient convergence is achieved. In each iteration, $5{\times}10^6$ particles were simulated, which is increased sufficiently to allow Serpent to converge on a solution to the adjoint transport problem. The adjoint flux is representative of an importance function and hence used to derive the weight window boundaries written to the weight window file. The particle importance spans several

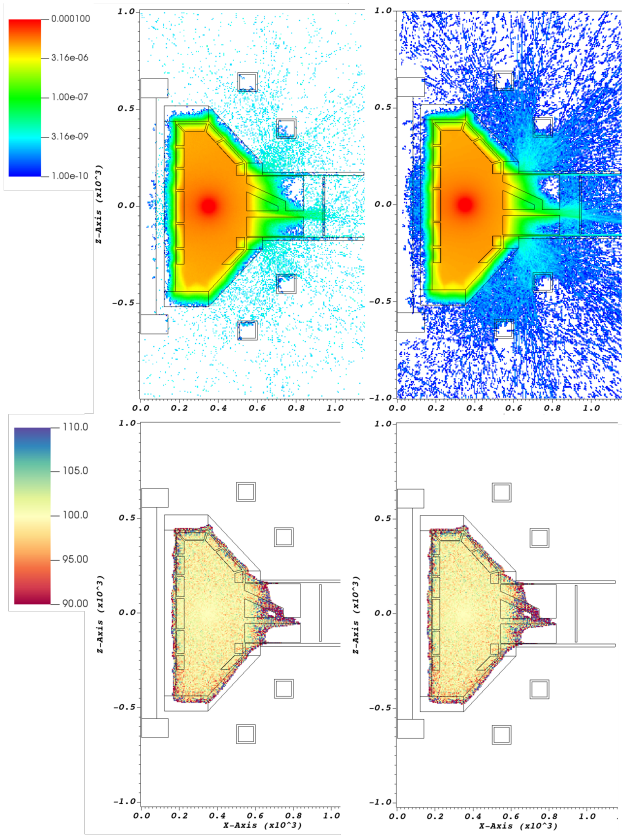

Figure 26: Neutron flux (cm^{-2} per source particle) map for the octamak baseline geometry with for Serpent STL (left) and OpenMC CSG geometry (right). Below each plot is a ratio to the MCNP result for each voxel. White represents a zero flux result. Only those values with statistical error <50% are plotted in the map of statistical error.

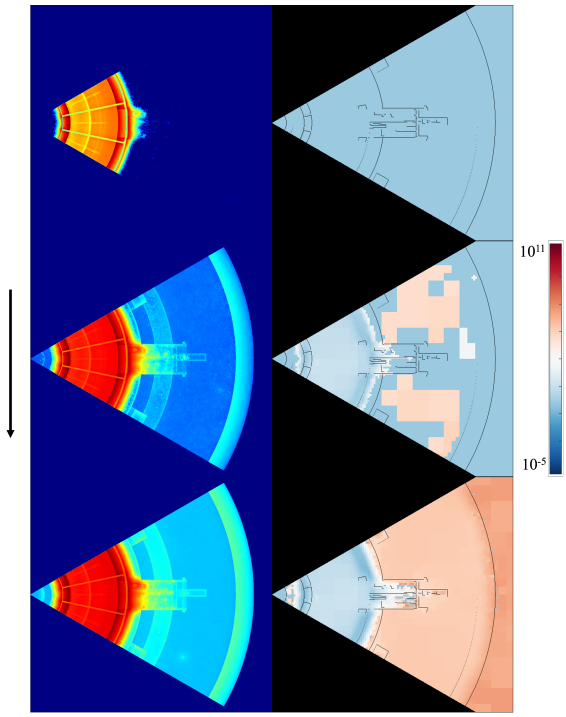

Figure 27: Plot of neutron flux (left) in each cycle of weight window generation for the Serpent hybrid CSG baseline with integrated STL model of the equatorial port. The corresponding logarithmic neutron importance is shown on the right with high importance after 3 iterations assigned globally across the ex-vessel region. Note that flux colour scale on the min and max values therefore the results between cycles are not similarly normalised.

orders of magnitude as anticipated through the shielding of the blanket and the vacuum vessel.

One penalty in performing a transport simulation with a weight window is an increase in runtime for the same number of particle histories - this increased 45% relative to the analog simulation. It is important to note that this was recorded with the 'set bala 1' option in order to mitigate issues related to significant fluctuations that occur in CPU usage due to excessive particle splitting performed by the weight window. Nonetheless, in spite of increased run time, the statistical improvement relative to the analog simulation significantly improves the figure of merit across all voxels by a factor 816. The error across the majority of voxels is reduced below 5% which is demonstrated in the plot of neutron flux displaying significantly less statistical noise (Figure 28) than that in the analog simulation as was presented in Figure 26.

A global weight window was also generated for the baseline geometry with equally good improvement relative to the analog simulation. Using the baseline geometry, a comparison between different weight window methods could be performed. For PFC 1-4, with 1 adjacent to the upper port, the nuclear heating has been used as a metric for the relative efficiency gain for each methodology. The integral neutron flux is also recorded in two identical volume cell tallies in the port interspace, positive (+y) and negative (-y) with respect to the y-axis. Table 5 gives the factor of figure of merit improvement relative to the analog simulation. In all cases there is clear improvement and importantly, good agreement between the calculated nuclear heating and neutron flux values. In the analog case, the relative error on all PFC value exceeded 5% which is a typical threshold for credible results. Both MCNP based methods proved less effective at reducing the error for the two interspace tallies while both ADVANTG and the Serpent weight window demonstrate the most improvement across the coils. To mitigate problems associated with long histories resulting in a considerable reduction in run time.

MCNP reports ten statistical tests that are extremely valuable when applying any form of variance reduction in which the simulation is purposefully biased. Each test is reported as pass or fail. While focus is often given on reducing the relative error, the behaviour of each statistical tests should be monitored as identification of possibly precise but inaccurate tally results is critical. Of course, precision is only one requirement for a good Monte Carlo calculation. There are many sources of error in a given calculation which if not minimised can indicate that even a zero variance calculation will not accurately predict natural behaviour. For both the ADVANTG and WWITER calculations the number of reported test failures is equal to 2, reduced from 5 in the analog case. However, it should be noted that the presence of any failures does not necessarily mean that the tally has not converged - it is ultimately down to the user to

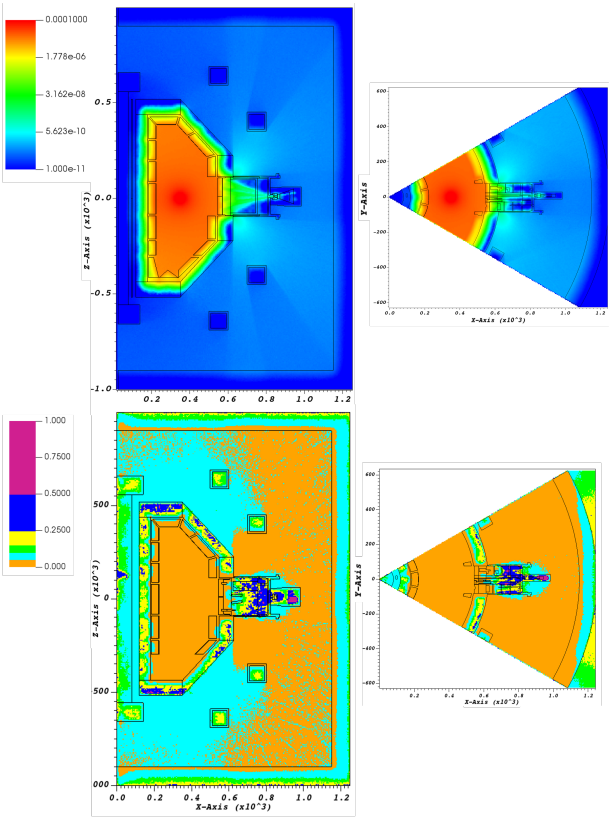

Figure 28: Neutron flux (cm^{-2} per source particle) in the Serpent hybrid model using a global weight window for a vertical (left) and horizontal slice (right). The associated map of statistical error is shown below for each slice

	Result	Rel. Error	FOM
Serpent			
PFC 1	1.14E-06	0.0038	314
PFC 2	5.02E-06	0.0015	422
PFC 3	5.27E-06	0.0016	467
PFC 4	1.88E-06	0.0042	167
Interspace +Y	3.14E-10	0.0117	48
Interspace -Y	3.21E-10	0.0119	31
MCNP + WWITER			
PFC 1	1.20E-06	0.0262	36
PFC 2	5.19E-06	0.0130	36
PFC 3	5.57E-06	0.0122	40
PFC 4	1.84E-06	0.0177	54
Interspace +Y	3.08E-10	0.0379	5
Interspace -Y	3.17E-10	0.0398	12
MCNP + ADVANTG			
PFC 1	1.20E-06	0.0043	893
PFC 2	5.24E-06	0.0032	397
PFC 3	5.48E-06	0.0034	347
PFC 4	1.95E-06	0.0047	527
Interspace +Y	3.28E-10	0.0835	1.1
Interspace -Y	3.48E-10	0.1080	1.1

Table 5: Comparison of different variance reduction techniques for the octamak baseline geometry. Results are reported for the nuclear heating (MeV per source particle) in each PFC and the neutron flux (cm^{-2} per source particle) in two port interspace tallies.

make a sound judgement.

There are statistical tests in Serpent which have been published to validate the neutron photon transport [44]. However, there is nothing yet analogous to the above statistical checks. The Serpent development team openly acknowledges that this is poorly documented and there has been very little work in this area over recent years. To the authors knowledge, no additional statistical tests are reported in OpenMC.

6.6. JET analysis model

6.6.1. Model development

Nuclear analysis on JET has been performed for many years based on 360° MCNP reference models and individual sector models of octants 1 and 2. These models are built on several assumptions and approximations to components which are significant to radiation transport. Due to the lack of documentation during the construction of the device, a significant number of unknowns are still present to this day, particularly related to understanding the isotopic composition of each material. Nonetheless, JET has operated as one of the worlds most successful tokamaks for decades, providing extensive and unique experimental data in a DT operating regime, which has proven extremely valuable in the validation of neutronics codes and methods [45][46][47].

JET will come to the end of its operations in 2023, concluding over 100,000 plasma pulses. A large scale re-purposing and decommissioning effort will be needed, with many components not having been replaced since the commissioning of the device in the early 1980's. As part of this project, the waste in JET will need to be accurately characterised and managed appropriately. In 2021, an updated MCNP model of octant 1 was developed for this purpose. The model includes heterogeneous representation of all in-vessel components including the ICRH systems, limiters, divertor, ITER-like wall, cooling channels and auxiliary systems/diagnostics. The simplification process of the CAD model, which represents a 45° sector of the tokamak, required around **3 person months** of effort. The most time-consuming tasks included the removal of spline and off axis tori, most notably in the toroidal field coils and in-vessel heating systems. The simplified CAD model was translated to MCNP CSG format using SuperMC is shown in Figure 29.

The MCNP model consists of 6069 cells and 21675 surfaces. An analog neutron only simulation for 1×10^8 particles takes 5 hours 33 minutes (wall time) with 32 CPU cores using the JET parametric plasma source routine. The highest statistical error occurs for the inboard cells; a weight window was thus employed to reduce the variance. In this instance, WWITER was the adopted method which successfully converged results on the inboard to <5% statistical error on $10 \times 10 \times 10$ cm^3 neutron flux mesh in over 98% of voxels.

6.6.2. Development of a mesh based model

Given the complexity of the CAD model, the updated JET model was used as a further test case for the CAD based work-

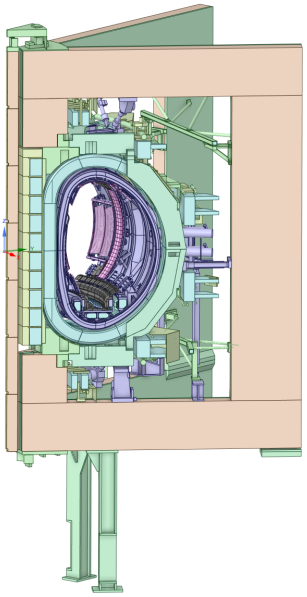

Figure 29: Updated JET octant 1 CAD model simplified for conversion to MCNP

flows. A Serpent STL model was created using the simplified CAD geometry. In total, this took ˜1 day of effort in order to produce a running model which records no lost particles. This is not strictly a fair comparison given that the model was already simplified to a level suitable for conversion to a CSG model. However, the most time consuming part of the simplification process is in removing splines and tori which would *not* have been required. Furthermore, the MCNP model loses several hundred particles in $1{\times}10^8$ source particles (for a point source). The cleaning of the model undertaken in preparation of the Serpent STL model was largely the removal of small gaps and non-physical artefacts in the geometry, that would likely resolve the geometry problems in the MCNP CSG model.

The CAD model is simplified to a level such that it meets a minimum requirement of 'cleanliness', quantified by a criterion for the number of lost particles. As highlighted here, inaccuracies likely remain in the model that need addressing in the CAD based workflow to derive a model that can transport particles. These additional requirements rely on a skilled CAD analyst to firstly identify and secondly implement a reasonable approximation. Therefore, while there may be significant time saving in some aspects of the simplification process (which is of course model dependant), CAD based workflows can still require a significant investment in model preparation. It is a common misconception that adopting CAD based workflows will remove this bottleneck in the workflow entirely.

Still, in the case of the JET model the results are impressive - the reported 1 day model preparation above also includes time taken to optimise the geometry in Serpent. This process of refactoring the geometry can significantly improve the model run time which is strongly dependant on how the geometry is structured. Serpent uses an adaptive search mesh that stores pre-assigned material data. The better populated this mesh is, the less time is needed to perform highly computationally expensive ray tests on-the-fly each time a neutron enters a give mesh cell. Serpent reports the 'fill fraction' for a given search mesh. This can be increased in resolution with the aim of achieving over 90% to maximise the amount of a priori data in the initialisation of the simulation. The other aspect of geometry refactoring is to include multiple nested universes. Each one of these is assigned its own search mesh which can be optimised. Note however that the resolution of the mesh is limited by the memory footprint and the amount of processing time prior to starting the simulation (initialisation time).

In the first step, all STL files for each different material in the JET model were described in a single universe. The total elapsed simulation time for 10,000 particles on 32 CPU cores was equal to 6.48 minutes with initialisation equal to 1.163 minutes and transport, 4.74 minutes. Refining the mesh from 89% fill fraction to 91% increased the initialisation time to 8.87 minutes and a simulation time of 3.59 minutes, therefore almost doubling the total simulation time. However, only the transport time scales with the number of particles therefore at more practical number of particles needed in production calculations, this would prove optimal. More nested universes were created starting with the largest STL files and monitoring universe fill fractions, memory allocation and the timing data. Following this process, the total run time was reduced to 3.54 minutes, with 1.18 minutes initiation and 2.32 minutes of transport. This is more than a 50% reduction in run time resulting from geometry optimisation procedure. Most notably, the total run time for the MCNP CSG geometry is 6.15 minutes thus the transport in the STL geometry is 43% faster. By comparison, the Serpent CSG model runs in 3.6 minutes. Also faster than MCNP and comparable to the STL geometry.

The plotted geometry (Figure 30) can be rendered in PNG format in a few seconds compared to the 5 minutes taken to load a plot of this model in the base version MCNP6.2. The plot is however not interactive and does not display cell lines as for MCNP. A neutron flux mesh of $10 \times 10 \times 10$ cm^3 voxel size covering the extent of the geometry was used to compare between the Serpent STL and MCNP CSG model.

As comparison to a weight window generated using the WWITER approach in MCNP, a weight window was generated in the global approach for the Serpent STL geometry. This served as a test case for generation of the weight window for a geometry defined entirely in STL format. The weight window importance mesh is plotted in Figure 31.

The figure shows that the weight window mesh spatially converges on the heaviest shielded inboard regions. Indeed the outboard ex-vessel region is in general given lower neutron importance because of the open upper, lower and equatorial ports. Also visible is the adaptive mesh following the final iteration cycle. The increased depth of recursive splitting in material regions is clearly visible, with particularly high mesh voxel density on the inboard where it is particularly important to capture large gradients in the profile of neutron flux.

This weight window was found to perform less effectively than that described in section for the octamak (section 6.5.2). The importance profile spans only two orders of magnitude

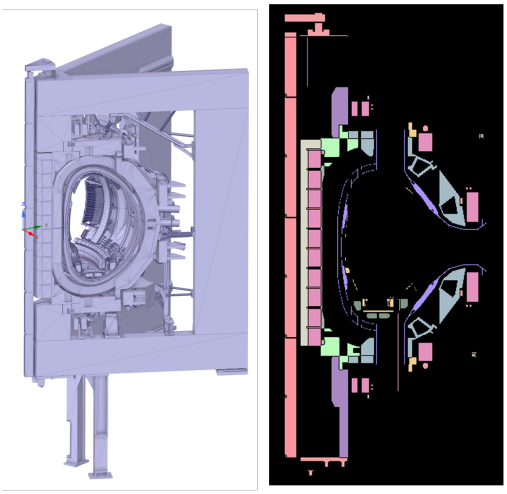

Figure 30: JET model CAD in STL format (left) and plot of radiation transport model (right) using the Serpent command line plotter

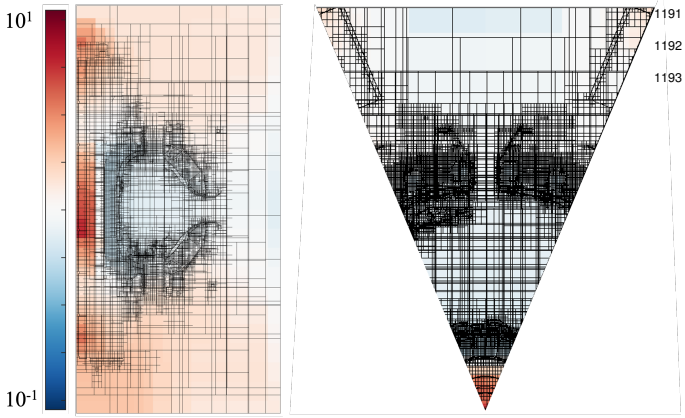

Figure 31: Logarithmic neutron importance map for global weight window generated for JET Serpent STL model. The black lines show the adaptive mesh which has been recursively split according to the spatial density profile.

compared to 10 for the octamak. It is thought that this is partly due to the combination of heavy shielded regions and large regions of particle leakage (through the ports). In this case, 65% of voxels were reduced to an error <5% compared with the analog case where this is 35%.

6.7. ITER analysis model

ITER analysis presents some of the most complex radiation transport models in fusion neutronics. The early 'A-lite' MCNP model of ITER represented a single regular 40° sector of the tokamak with all major components up to bioshield included with a simplified homogenised description. Over several years, this model was developed with ever increasing complexity through B-lite, C-lite and most recently C-Model R181031 [48], the current reference sector model for nuclear analysis of ITER. It contains a heterogeneous, as now constructed, model of the vacuum vessel and superconducting magnets. This model is fixed in CSG format as much of the 114,092 cell model geometry revisions have been made in MCNP, making it (almost) impossible to create an equivalent model that can be handled by conversion and/or CAD software.

C-Model R181031 was parsed into csg2csg with an OpenMC geometry and material XML files output after 30 hours. The most time consuming part of the conversion process is checking for duplicate surfaces and the correctly handling these in cell definitions. Using a compiled language would inevitably reduce the conversion time however this is insignificant, with the converted model serving as an ultimate testament of the capabilities of this tool.

The ITER OpenMC model is plotted in Figure 32 showing the geometry in all its complexity. Several hundred plot slices were iterated over and uncovered no geometry problems of the converted model. Each plot takes ~3 minutes to generate. Note that the MCNP file takes several hours to plot using its native plotter - the absence of a means for quick visualisation has proven restrictive for model updates and diagnosing geometry problems. Many person months of effort over several years have been spent reducing the lost particle count in the reference model which is of the order of 50 particles in 1×10^8 source particles. 90 particles were lost in OpenMC which is of the same order and therefore not immediately of concern.

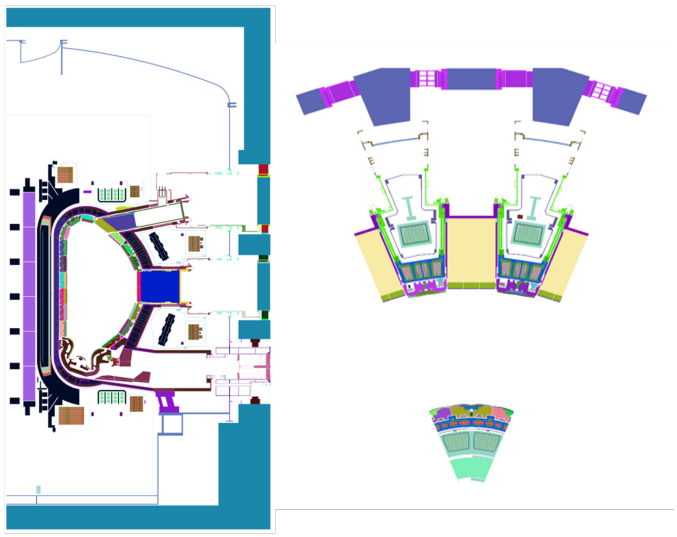

Figure 32: Plot of ITER C-Model R181031 OpenMC geometry produced using csg2csg. Two slices are shown at Y=0 (left) and Z=0 (right).

For such a complex model, there is a high memory requirement that can lead to issues particularly when tallies are included. The allocated memory in MCNP is 2.1 GB whereas in OpenMC it is almost half of this equal to 1.1 GB. No tallies are included for this comparison.

The model loading and transport times are tabulated in Table 6. The significant time to transport a single particle (134 minutes) in MCNP results from the routine responsible for loading the geometry (imcn). This processing of the geometry is also required for loading the interactive plotter. The initialisation time of this model in OpenMC is equal to 95 seconds. The single particle run was performed using a single CPU core while the result for 1×10^8 particles was run on 64 CPU cores. In both cases the simulation is performed in neutron only mode

	Wall time	
	NPS=1	**NPS=10^8**
MCNP	134 mins	9.5 hours
OpenMC	6 mins	82 hours

Table 6: Comparison of run times between MCNP and OpenMC for different number of particles (NPS) simulated in C-Model R181031.

using an isotropic 14 MeV point source. The order of magnitude difference in run times to 10^8 histories is highly significant - nuclear analysis of ITER can typically take several days of computing time alone, especially when variance reduction is needed. For this reason, large compute resource is required of the order of several hundred CPU cores. MCNP is capable of scaling up to a few hundred cores [49] beyond which there are diminishing returns. OpenMC on the other hand is highly scalable [4] - ITER analysis could be performed on hour time scales rather than days with the open deployability of the code facilitating this without restriction.

As future iterations of C-model are likely to only increase in detail as ITER construction evolves, the memory requirements may eventually become a limiting factor for future modelling. This has been observed with the recently developed E-lite model [50] which is the first 360° model of ITER. It is not possible to run this model with MCNP without any patches to the source code that optimise memory use such as those implemented in D1SUNED [51]. The alternative approaches demonstrated in this paper hold potential for this application.

7. Conclusion

In this work, we have conducted both computational and experimental benchmarks to investigate both the validity and usability of current and emergent transport codes and methods, MCNP, OpenMC, Serpent and DAGMC. Large scale fusion reactor designs provide one of the most complex and challenging particle transport problems and therefore necessitates a robust implementation of bespoke advanced code features not fundamental to other applications of Monte Carlo codes. It is only in recent years that Serpent and OpenMC have developed capabilities such as coupled neutron-photon transport and the ability to transport directly on CAD, thus broadening the scope of the codes to fusion neutronics applications.

The capability to capture reaction rates was demonstrated through comparison with selected experimental data available in the SINBAD database. Two FNG experiments, FNG HCPB and FNG Cu were performed using CSG based models in MCNP, Serpent and OpenMC, with good consistency between transport codes for each of the measured reaction rates. Tritium production rates were also calculated for the HCPB mock up in OpenMC, Serpent and MCNP with no reported statistically significant disagreement. The FNS experiment captures the physics of scattering which was compared for each code through evaluation of neutron spectra at several scattering angles. In all cases, the comparison to experimental data was limited by the unavailability of an accurate reproduction of the FNG source term in OpenMC which is under ongoing development.

The ability to compare between different transport codes would not have been possible without the csg2csg tool. It has been demonstrated to accurately convert MCNP to Serpent and OpenMC with minor modifications required by the user. A range of nuclear responses were determined in vessel using the octamak model which serves as a suitable sector tokamak model for code benchmarking. A variant of this, with a detailed ITER-like dummy port plug was integrated as an STL universe into a CSG model in Serpent. The latter provides a test of the hybrid CSG-CAD approach which the author views as the most immediate workflow that will prevail in future conceptual design phase tokamak studies. In this way, the advantages of both workflows can be harnessed with a CAD model capturing all detail of the component/detector of interest, which can be held in a universe with the rest of the geometry in CSG format.

To explore ex-vessel responses, variance reduction techniques were necessary for the octamak geometry. Three prominent methods, ADVANTG with MCNP, Serpent's deterministic response matrix solver and WWITER were explored for this deep shielding problem with interspace and PFC tallies set as the benchmark. While requiring 3.5 days to generate, the weight window generated by Serpent proved highly effective for this problem with the statistical error on a neutron flux mesh reduced below 5% for the majority of the ex-vessel with a factor 816 increase relative to the analog simulation in the figure of merit. This method worked in both the global, uniformly converging results across the geometry extend and local approach, targeting specific responses. This was demonstrated in application to the FNG Cu experiment where the relative error on the reaction was reduced from 27% to 7% for the tungsten foil positioned furthest from the source. In all applications of variance reduction techniques, a process of trial and error is required with ADVANTG and Serpent to derive the optimal weight window. The power of these methods lies in the capability to automate in the first step, the parameters defining the weight window with minimal user input. It was highlighted that outstanding in this area is the implementation of a rigorous set of statistical checks beyond the the relative error in Serpent and OpenMC, such as those given by MCNP. This is recommended to be developed in parallel with any variance reduction techniques.

While we have explored new capabilities and code features, it is equally as important to the user community to detail code limitations and areas for future development. DAGMC is not trivial to install and likely requires an experienced code user. Furthermore, in all CAD-based transport approaches, we have seen that this does not completely bypass the CAD modelling stage. Rather, in the CAD-centric approach the emphasis of the tasks performed in the CAD program shifts from simplification to the 'cleaning' of the model. The fact that OpenMC is open source holds large potential for collaboration and development is greatly aided by code transparency. While it is necessary for users to work with the development branch of the code, it is important that all versions are validated against a consistent set of fusion-relevant benchmarks. One potential limitation of

Serpent may arise if the imminent release of an official version falls under strict commercial licensing conditions.

The workflow built around the 'paramak' toolkit has also been explored. There is a clear use case for this tool in sweeping a broad design space in early pre-conceptual design studies to guide concept selection. For a simple spherical tokamak model, the workflow was demonstrated and used to create DAG-OpenMC which was subsequently validated against DAG-MCNP. The sensitivity to the faceting tolerance for the geometry was assessed, where it was seen that the response changed considerably for the blanket first wall heating between a faceting tolerance of 10^{-4} to $5{\times}10^{-4}$. This requires further investigation though in general the faceting tolerance had minimal impact on nuclear responses for this simple model. The faceting tolerance parameter also significantly effects the memory requirement and run time as was seen in an independent study of the octamak geometry. Therefore, careful selection based on the specific geometry and available computational resource is strongly encouraged.

Finally, to explore the boundaries of the transport codes current capabilities, we have looked at two of the more complex geometries in current fusion neutronics analysis. A heterogeneous JET octant 1 model developed for characterisation of nuclear waste was converted to a Serpent STL unstructured surface model. It was seen that there was significant improvement in efficiency, both in the model preparation and simulation run time following a concerted geometry refactoring effort. For ITER, the direct transition from the MCNP to a CAD based reference model is not possible owing to its complexity. However, for the first time, the current reference model is available in alternative CSG formats with some insight given to potential improved performance over the MCNP model. Notably, the run time was reduced by an order of magnitude and memory consumption by a factor 2. This is highly significant in light of the major bottlenecks in both run time and memory usage currently encountered at a time when the nuclear analysis of ITER and demonstration of safety is a top priority for the French nuclear regulator, the ASN.

8. Further work and development needs

Now that it has been demonstrated that OpenMC can handle complex tokamak reactor geometries, further investigation should be made into the parallel performance of OpenMC and how it scales at the level of thousands of processors. There is ongoing work to enable deployment of OpenMC on GPU nodes [52] which is fundamental to transition to the exascale. The throughput of a single GPU is estimated to be >144 CPUs worth of compute. At the time of writing, a capability to read in weight windows has been merged into the development branch of OpenMC (v0.13.0). The addition of weight windows was the final major outstanding requirement as outlined in the introduction to extend the scope of the code to fusion applications. In the case of the STEP reactor, where DAG-OpenMC is currently being used, it will be possible to perform more detailed ex-vessel analysis which will be a fundamental aspect of the plant licensing.

As the uptake of alternative transport codes grows, priority must be given to the development of a validation suite containing benchmarks relevant to nuclear fusion. This is welcomed from both nuclear data and transport code development point of view and can build on existing efforts such as the JADE toolkit. The framework for validation as presented in this paper is under development to integrate SINBAD benchmarks and other relevant experimental data sets that will be included and be run automatically as part of a testing suite. License conditions permitting, this could be hosted on a cloud server. The availability of experimental data is fundamental to this and efforts like the Compilation of Nuclear Data Experiments for Radiation Characterisation (CoNDERC) [53] hosted by the IAEA and publicly available is an important step in this direction. If measurements of shutdown dose rates are incorporated then it will be possible to extend the scope of the validation suite beyond transport codes to include also activation/inventory codes such as FISPACT-II [54].

Related to this, the other core part of the workflow that has not been explored is the calculation of shutdown dose rates. There are many different methods currently available for coupling transport and activation calculations in order to determine a decay gamma field. These largely rely on MCNP for neutron transport however MCR2S, one of the methods developed at UKAEA has been extended to interface with Serpent and OpenMC [55]. An interesting comparison could be made against the built in depletion solver in Serpent [56] as well as OpenMC once this capability is available. The recently developed novel 1 step method (N1S) at UKAEA, which is currently implemented for MCNP should form part of these comparisons.

Acknowledgements

The views and opinions expressed herein do not necessarily reflect those of the ITER Organization. The author would like to thank Andrew Davis for their help through development of csg2csg and assistance in the installation procedure of DAGMC. Funding for this work has been provided through RCUK [grant number EP/T012250/1]. To obtain further information on the data and models underlying this paper please contact PublicationsManager@ukaea.uk. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

References

[1] C. J. Werner, J. Armstrong, F. B. Brown, J. S. Bull, L. Casswell, L. J. Cox, D. Dixon, R. A. Forster, J. T. Goorley, H. G. Hughes, J. Favorite, R. Martz, S. G. Mashnik, M. E. Rising, C. Solomon, A. Sood, J. E. Sweezy, C. J. Werner, A. Zukaitis, C. Anderson, J. S. Elson, J. W. Durkee, R. C. Johns, G. W. McKinney, G. E. McMath, J. S. Hendricks, D. B. Pelowitz, R. E. Prael, T. E. Booth, M. R. James, M. L. Fensin, T. A. Wilcox, B. C. Kiedrowski, MCNP User's Manual Code Version 6.2, Los Alamos National Laboratory (2017).

[2] R. Pampin, A. Davis, J. Izquierdo, D. Leichtle, M. J. Loughlin, J. Sanz, A. Turner, R. Villari, P. P. Wilson, Developments and needs in nuclear analysis of fusion technology, Fusion Engineering and Design 88 (2013) 454–460.

[3] J. Leppänen, E. al, The Serpent Monte Carlo code: Status, development and applications in 2013, Ann. Nucl. Energy (2015).

[4] P. K. Romano, N. E. Horelik, B. R. Herman, A. G. Nelson, B. Forget, K. Smith, Openmc: A state-of-the-art monte carlo code for research and development, Annals of Nuclear Energy 82 (2015) 90–97. Joint International Conference on Supercomputing in Nuclear Applications and Monte Carlo 2013, SNA + MC 2013. Pluri- and Trans-disciplinarity, Towards New Modeling and Numerical Simulation Paradigms.

[5] Sinbad – radiation shielding benchmark experiments, Annals of Nuclear Energy 159 (2021) 108254.

[6] I. Kodeli, M. Angelone, Fng copper benchmark evaluation for the sinbad database, Fusion Engineering and Design 146 (2019) 312–315. SI:SOFT 30.

[7] P. Batistoni, M. Angelone, L. Bettinali, P. Carconi, U. Fischer, I. Kodeli, D. Leichtle, K. Ochiai, R. Perel, M. Pillon, I. Schäfer, K. Seidel, Y. Verzilov, R. Villari, G. Zappa, Neutronics experiment on a helium cooled pebble bed (hcpb) breeder blanket mock-up, Fusion Engineering and Design 82 (2007) 2095–2104. Proceedings of the 24th Symposium on Fusion Technology.

[8] Tokai-mura, Naka-gun, Ibaraki-ken, Collection of experimental data for fusion neutronics benchmark, 1994.

[9] UKAEA, `https://github.com/ukaea/SpaceClaim_API_NeutronicsTools`, 2020.

[10] T. Tautges, P. Wilson, J. Kraftcheck, B. Smith, D. Henderson, Acceleration techniques for direct use of cad-based geometries in monte carlo radiation transport.

[11] J. Leppänen, Cad-based geometry type in serpent 2-application in fusion neutronics, in: Mathematics and Computations, Supercomputing in Nuclear Applications and Monte Carlo International Conference, Curran Associates Inc., United States, 2015, pp. 1635–1646. Mathematics and Computations, Supercomputing in Nuclear Applications and Monte Carlo International Conference.

[12] B. Colling, T. Eade, M. Gilbert, J. Naish, S. Zheng, Comparative study of neutronics analysis techniques for radioactive waste assessment, Fusion Science and Technology 74 (2018) 1–10.

[13] M. Rising, et al., Upcoming MCNP6.3(r) release: New features and improvements, 2021 MCNP(R) User Symposium, 2021-07-12/2021-07-16 (2021).

[14] T. Kaltiaisenaho, Implementing a photon physics model in serpent 2, Aalto University, Thesis (2016).

[15] A. Turner, Investigations into alternative radiation transport codes for ITER neutronics analysis, Transactions of the American Nuclear Society 116 (2017) 251–254.

[16] A. Turner, A. Burns, B. Colling, J. Leppänen, Applications of Serpent 2 Monte Carlo Code to ITER Neutronics Analysis, Fusion Science and Technology 74 (2018) 315–320.

[17] A. Valentine, B. Colling, R. Worrall, J. Lepp, Benchmarking of the Serpent 2 Monte Carlo Code for Fusion Neutronics Applications, PHYSOR 2020 conference proceedings (2020).

[18] E.J.McGrath, D.C.Irving, TECHNIQUES for EFFICIENT MONTE CARLO SIMULATION, ORNL-RSIC-38 3 (1975).

[19] J. Leppänen, Response Matrix Method–Based Importance Solver and Variance Reduction Scheme in the Serpent 2 Monte Carlo Code, Nuclear Technology (2019) 1–17.

[20] S. W. Mosher, S. R. Johnson, A. M. Bevill, A. M. Ibrahim, C. R. Daily, T. M. Evans, J. C. Wagner, J. O. Johnson, R. E. Grove, ADVANTG: An Automated Variance Reduction Parameter Generator, 2015.

[21] A. Davis, A. Turner, Application of novel global variance reduction methods to fusion radiation transport problems, International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (2011).

[22] S. N. Laboratories, Coreform cubit, `https://coreform.com/products/coreform-cubit/`, 2021.

[23] J. Leppänen, Performance of woodcock delta-tracking in lattice physics applications using the serpent monte carlo reactor physics burnup calculation code, Annals of Nuclear Energy 37 (2010) 715–722.

[24] T. J. Labossiere-Hickman, B. Forget†, Selected vera core physics benchmarks in openmc, Transactions of the American Nuclear Society 117 (2017) 1520 – 1523.

[25] K. S. Chaudri, S. M. Mirza, Burnup dependent monte carlo neutron physics calculations of iaea mtr benchmark, Progress in Nuclear Energy 81 (2015) 43 – 52.

[26] Kelly Daniel J., Aviles Brian N., Romano Paul K., Herman Bryan R., Horelik Nicholas E., Forget Benoit, Analysis of select beavrs pwr benchmark cycle 1 results using mc21 and openmc, Proceedings of the International Conference on Physics of Reactors (PHYSOR 2014) 47 (2015).

[27] Herman Bryan R., Forget Benoit, Smith Kord, Romano Paul K., Sutton Thomas M., K. D. J., Aviles Brian N., Analysis of tally correlation in large light water reactors, JAEA-Conf–2014-003 47 (2015).

[28] N. Horelik, B. Herman, B. Forget, K. Smith, Benchmark for evaluation and validation of reactor simulations (beavrs) (2013).

[29] J. A. Walsh, B. Forget, K. S. Smith, Validation of openmc reactor physics simulations with the b and w 1810 series benchmark, Transactions of the American Nuclear Society (2013) 1301 – 1304.

[30] D. Laghi, M. Fabbri, L. Isolan, M. Sumini, G. Shnabel, A. Trkov, Application of JADE v&v capabilities to the new FENDL v3.2 beta release, Nuclear Fusion 61 (2021) 116073.

[31] R. Forrest, R. Capote, N. Otsuka, T. Kawano, A. Koning, S. Kunieda, J.-C. Sublet, Y. Watanabe, Fendl-3 library–summary documentation, INDC (NDS)-628, IAEA Nuclear Data Section (2012).

[32] A. Davis, S. Lilley, J. Shimwell, csg2csg: a Tool to Assist Validation & Verification, in: European Physical Journal Web of Conferences, volume 247 of *European Physical Journal Web of Conferences*, p. 04012.

[33] J. Shimwell, J. Billingsley, R. Delaporte-Mathurin, D. Morbey, M. Bluteau, P. Shriwise, A. Davis, The paramak: automated parametric geometry construction for fusion reactor designs. [version 1; peer review: 2 approved], F1000Research 10 (2021).

[34] J. Shimwell, J. Billingsley, R. Delaporte-Mathurin, D. Morbey, M. Bluteau, P. Shriwise, A. Davis, Paramak git repository, `https://github.com/fusion-energy/paramak`, 2021.

[35] Y. Wu, J. Song, H. Zheng, G. Sun, L. Hao, P. Long, L. Hu, CAD-based Monte Carlo program for integrated simulation of nuclear system SuperMC, Annals of Nuclear Energy (2015).

[36] NEA, Joint Evaluated Fission and Fusion (JEFF) Nuclear Data Library JEFF-3.3 (2017).

[37] E.M.Zsolnay, R.Capote, H.K.Nolthenius, A.Trkov, Summary description of the new international reactor dosimetry and fusion file (irdff release 1.0), Technial report INDC(NDS)-0616 (2012).

[38] A. Trkov, P. Griffin, S. Simakov, L. Greenwood, K. Zolotarev, R. Capote, D. Aldama, V. Chechev, C. Destouches, A. Kahler, C. Konno, M. Koštál, M. Majerle, E. Malambu, M. Ohta, V. Pronyaev, V. Radulović, S. Sato, M. Schulc, E. Šimečková, I. Vavtar, J. Wagemans, M. White, H. Yashima, Irdff-ii: A new neutron metrology library, Nuclear Data Sheets 163 (2020) 1–108.

[39] D. Brown, M. Chadwick, R. Capote, A. Kahler, A. Trkov, M. Herman, A. Sonzogni, Y. Danon, A. Carlson, M. Dunn, D. Smith, G. Hale, G. Arbanas, R. Arcilla, C. Bates, B. Beck, B. Becker, F. Brown, R. Casperson, J. Conlin, D. Cullen, M.-A. Descalle, R. Firestone, T. Gaines, K. Guber, A. Hawari, J. Holmes, T. Johnson, T. Kawano, B. Kiedrowski, A. Koning, S. Kopecky, L. Leal, J. Lestone, C. Lubitz, J. Márquez Damián, C. Mattoon, E. McCutchan, S. Mughabghab, P. Navratil, D. Neudecker, G. Nobre, G. Noguere, M. Paris, M. Pigni, A. Plompen, B. Prytychenko, V. Pronyaev, D. Roubtsov, D. Rochman, P. Romano, P. S. beeckx, S. Simakov, M. Sin, I. Sirakov, B. Sleaford, V. Sobes, E. Soukhovitskii, I. Stetcu, P. Talou, I. Thompson, S. van der Marck, L. Welser-Sherrill, D. Wiarda, M. White, J. Wormald, R. Wright, M. Zerkle, G. Žerovnik, Y. Zhu, Endf/b-viii.0: The 8th major release of the nuclear reaction data library with cielo-project cross sections, new standards and thermal scattering data, Nuclear Data Sheets 148 (2018) 1–142. Special Issue on Nuclear Reaction Data.

[40] R. Macfarlane, D. W. Muir, R. M. Boicourt, A. C. Kahler, III, J. L. Conlin, The njoy nuclear data processing system, version 2016 (2017).

[41] C. Konno, K. Ochiai, K. Takakura, S. Sato, Remarks on kerma factors in ace files, Nuclear Data Sheets 118 (2014) 450–452.

[42] P. Shriwise, J. Shimwell, A. Davis, `https://github.com/pshriwise/double-down`, 2021.

[43] P. Shriwise, P. Wilson, A. Davis, Leveraging intel's embree in the dagmc toolkit, Transactions of the American Nuclear Society 113 (2015) 717–720.

[44] T. Kaltiaisenaho, Statistical Tests and the Underestimation of Variance in Serpent 2, number VTT-R-00371-14 in VTT Research Report, VTT Technical Research Centre of Finland, Finland, 2014. Project code: 77464.

[45] Iter oriented neutronics benchmark experiments on neutron streaming and shutdown dose rate at jet, Fusion Engineering and Design 123 (2017) 171–176. Proceedings of the 29th Symposium on Fusion Technology (SOFT-29) Prague, Czech Republic, September 5-9, 2016.

[46] J. D. Naish, P. Batistoni, B. Kos, B. Obryk, R. Villari, T. Vasilopoulou, I. Stamatelatos, Comparison of neutron flux streaming calculations to the 2019-2020 jet experimental deuterium-deuterium results1, Fusion Engineering and Design 170 (2021) 112538.

[47] N. Fonnesu, R. Villari, D. Flammini, P. Batistoni, U. Fischer, P. Pereslavtsev, Shutdown dose rate studies for the dte2 campaign at jet, Fusion Engineering and Design 161 (2020) 112009.

[48] D. Leichtle, B. Colling, M. Fabbri, R. Juarez, M. Loughlin, R. Pampin, E. Polunovskiy, A. Serikov, A. Turner, L. Bertalot, The ITER tokamak neutronics reference model c-model, Fusion Engineering and Design 136 (2018) 742–746. Special Issue: Proceedings of the 13th International Symposium on Fusion Nuclear Technology (ISFNT-13).

[49] T. S. Grieve, Mcnp6.2 benchmarks: Scaling and hardware considerations, 2021 MCNP User Symposium, LA-UR-21-26855 (2021).

[50] R. Juarez, G. Pedroche, M. Loughlin, R. Pampin, P. Martinez, M. Pietri, J. Alguacil, F. Ogando, P. Sauvan, A. Lopez-Revelles, A. Kolšek, E. Polunovskiy, M. Fabbri, J. Sanz Recio, A full and heterogeneous model of the ITER tokamak for comprehensive nuclear analyses, Nature Energy 6 (2021) 1–8.

[51] P. Sauvan, et al., D1SUNED system for the determination of decay photon related quantities, Fusion Engineering and Design 151 (2020) 111399.

[52] J. R. Tramm, P. K. Romano, J. Doerfert, A. L. Lund, P. C. Shriwise, A. R. Siegel, G. Ridlely, Toward Portable GPU acceleration of the OpenMC Monte Carlo Particle Transport code, Physor 2022 (2022).

[53] J. Sublet, Consultant meeting on the on compilation of nuclear data experiments for radiation characterization (conderc), `https://www-nds.iaea.org/conderc/doc`, 2021.

[54] J.-C. Sublet, J. Eastwood, J. Morgan, M. Gilbert, M. Fleming, W. Arter, FISPACT-II: An advanced simulation system for activation, transmutation and material modelling, Nuclear Data Sheets 139 (2017) 77–137. Special Issue on Nuclear Reaction Data.

[55] T. Eade, B. Colling, J. Naish, L. W. Packer, A. Valentine, Shutdown dose rate benchmarking using modern particle transport codes, Nuclear Fusion 60 (2020).

[56] J. Leppänen, T. Kaltiaisenaho, Expanding the use of Serpent 2 to fusion applications: Shut-down dose rate calculations, Physics of Reactors 2016, PHYSOR 2016: Unifying Theory and Experiments in the 21st Century 5 (2016) 2872–2883.