Application of the Denovo Discrete Ordinates Radiation Transport Code to Large-Scale Fusion Neutronics

Katherine E. Royston,[a,*] Seth R. Johnson,[a] Thomas M. Evans,[a] Scott W. Mosher,[a]

Jonathan Naish,[b] Bor Kos,[c] and JET Contributors[†]

[a] *Oak Ridge National Laboratory, Radiation Transport Group*
P.O. Box 2008 MS6170, Oak Ridge, Tennessee 37831-6170

[b] *Culham Centre for Fusion Energy, Culham Science Centre*
Abingdon, Oxon, OX14 3DB, UK

[c] *Jožef Stefan Institute, Reactor Physics Division*
Jamova cesta 39, Ljubljana SI-1000, Slovenia

[*]Email: roystonke@ornl.gov
Phone: 865-241-7429

For submission to:

Fusion Neutronics Special Issue
Fusion Science and Technology

Number of pages: 25
Number of tables: 4
Number of figures: 10

[†] See the author list of "X. Litaudon *et al* 2017 *Nucl. Fusion* 57 102001"

Abstract

Fusion energy systems pose unique challenges to the modeling and simulation community. These challenges must be met to ensure the success of the ITER experimental fusion reactor. ITER's complex systems require detailed modeling that goes beyond the scale of comparable simulations to date. In this work, the Denovo radiation transport code was used to calculate neutron fluence and kerma for the Joint European Torus (JET) streaming benchmark. This work was performed on the Titan supercomputer at the Oak Ridge Leadership Computing Facility (OLCF). Denovo is a novel three-dimensional discrete ordinates transport code designed to be highly scalable. Sensitivity studies have been completed to examine the impact of several deterministic parameters. Results were compared against experiment as well as the Monte Carlo N-Particle (MCNP) and Shift Monte Carlo codes.

Keywords — fusion energy, JET, ITER, radiation transport, discrete ordinates

I. INTRODUCTION

Fusion energy systems pose unique challenges to the modeling and simulation community due to their size and complexity. The ITER experimental fusion reactor will be the world's largest fusion facility and will demonstrate the feasibility of fusion power. The design of the ITER facility requires detailed radiation responses beyond the scale of any comparable simulations that have been performed to date. Current simulations are limited to a combination of coarse spatial meshes and energy grids, as well as coupling of smaller, compartmentalized models. Novel implementations and methods are needed to fully resolve the radiation field for this class of problems.

Denovo is a three-dimensional discrete ordinates (S_N) transport code under active development at Oak Ridge National Laboratory (ORNL) [1]. Denovo is highly scalable, allowing it to provide high fidelity solutions across large phase spaces such as those encountered in fusion energy systems. Denovo has been applied for neutron shielding analyses of ITER problems in the past [2,3]; however, this work aims to provide a more comprehensive, systematic validation of Denovo for fusion neutronics analysis. To accomplish this goal, the neutronics behavior of the Joint European Torus (JET) facility has been modeled. The JET streaming benchmark [4] problem was run with Denovo on the Titan XK7 supercomputer at the Oak Ridge Leadership Computing Facility (OLCF), and sensitivity studies were conducted to examine the impact of several deterministic parameters. In addition to comparison with the experimental benchmark data, results were compared against solutions generated by the Monte Carlo N-Particle (MCNP) [5] and Shift [6] Monte Carlo codes.

Section II gives background on the Denovo code, while Section III details the JET streaming benchmark. Section IV describes the process used to model the benchmark problem with Denovo. Results are presented and discussed in Section V, and Section VI gives final remarks and recommendations for future work.

II. DENOVO RADIATION TRANSPORT CODE

Denovo is a massively parallel deterministic transport solver that solves the steady-state linear Boltzmann transport equation using state-of-the-art solution methods. Denovo provides multiple transport methods, but in this work we confine our use to the discrete ordinates (S_N)

solver on regular (Cartesian) grids. The numerical methods and parallel strategies used in Denovo are documented in detail in several recent studies [1,7–10]. Appendix A summarizes the discrete ordinates (S_N) equations and provides a brief overview of the solution methods in Denovo in order to motivate the parameter analysis and performance of the JET model runs described in Section V.

The spatial discretization of the S_N equations (3) can be performed using several methods. Denovo supports a family of cell-balance and discontinuous finite-element schemes, but in this work, all results are generated using a step-characteristics/slice balance (SC) method [1]. Although SC is only first-order accurate, this method is used in the JET simulations because it produces uniformly positive solutions as long as the source is positive. SC is also memory efficient for large problems because it only requires one unknown per spatial cell, while higher order finite-element methods require multiple unknowns per spatial cell and can produce negative solutions when the mean free path of the particle is not fully resolved. Higher order finite-element methods also require nonlinear iterative techniques to correct for nonphysical negative fluxes.

Denovo supports several stationary and nonstationary (Krylov) solvers with a family of preconditioners [11] that can be applied to Eq. (10). All of the results that follow use the unpreconditioned, restarted generalized minimal residual method (GMRES)(m), which has been demonstrated to be very efficient for this class of problems [1,7,12].

Denovo uses an implementation of the Koch-Baker-Alcouffe (KBA) parallel sweep algorithm [1,9,13] to invert $\mathbf{L}$ in Eq. (10). KBA is a parallel wavefront algorithm that parallelizes the space-angle components of the sweep. In this algorithm, the mesh is decomposed into domains of size (I_b, J_b, K) on each processor, and each domain is decomposed into computational blocks of size (I_b, J_b, K_b). Defining (I, J, K) as the total global number of cells in each direction and (I_b, J_b, K_b) as the block number of cells in each direction results in

$$P_I = \frac{I}{I_b}\,, \quad P_J = \frac{J}{J_b}\,, \quad B_K = \frac{K}{K_b}\,,$$

where (P_I, P_J) are the number of processors in the x and y directions, and B_K is the number of blocks in the z direction. The description above, which implies a uniform block size on every domain, is used for convenience. While Denovo does not require uniform block sizes, nonuniform block sizes can yield parallel inefficiencies.

Two enhancements were recently added to improve the efficiency of Denovo's parallel solution

methodology. One enhancement is a multilevel energy decomposition [7, 11] in which space is partitioned into blocks as described above, and energy is partitioned into sets. Each set contains the full mesh (all of the blocks), and the KBA sweeps never cross set boundaries. Supplementing this inter-domain decomposition, an on-node space-angle-energy optimization was added that utilizes an NVIDIA graphics processing unit (GPU) accelerator and introduces energy into the traditional KBA space-angle pipeline [9, 10].

The GPU algorithm has been specifically designed to map to the Titan supercomputer memory and parallelism architecture [14]. A Titan node has an AMD 16-core Interlagos processor connected by a PCIe-2 bus to a Kepler (K20x) NVIDIA GPU processor. The nodes are configured in a 3D torus via Gemini interconnects. Message Passing Interface (MPI) is used to manage internode communication. While this algorithm was developed for Titan, it will apply to any NVIDIA GPU accelerated architecture.

In order to achieve optimal efficiency, the GPU algorithm limits the choice of available solution algorithms that can be applied to solve Eq. (10). Specifically, to attain maximum occupancy on the GPU, parallelization is required over energy. This requires using GMRES(m) over all independent variables—space, angle, and energy—during the solution. The energy groups used in the JET problems described herein contain no upscattering, which is where a neutron *gains* energy from a collision with a neutron in thermal motion. Without upscattering, the scattering operator $\mathbf{S}$ is lower triangular, and Gauss-Seidel will converge in a single iteration. However, the improved parallel performance obtained on the GPU compensates for the poorer iterative behavior of GMRES(m) for this class of problems [10].

In summary, all of the problems described in Section V use Denovo with SC spatial differencing. GMRES(m) Krylov iteration is used over space, angle, and energy to solve Eq. (10). Energy is decomposed in sets, and space is decomposed into blocks on each node using the standard KBA decomposition. The sweep is performed using the KBA algorithm, with on-node space-angle-energy sweeps being performed on the GPUs.

III. JET STREAMING BENCHMARK

The JET facility located at the Culham Centre of Fusion Energy (CCFE) is the largest tokamak in the world. During deuterium (DD) operations, JET produces neutron emission rates

in excess of 10^{16} n/s, and an upcoming deuterium-tritium (DT) campaign anticipates neutron emission rates over 10^{18} n/s [4]. As the currently operating fusion device closest to ITER, JET is uniquely positioned to provide experimental data that are extremely relevant to ITER [15]. In particular, experiments performed at JET are vital to the validation of the neutronics codes and nuclear data to be used in the design and operation of the ITER facility.

In the JET streaming benchmark, the neutron fluence and air kerma through penetrations present in the JET torus hall were measured with thermoluminescent dosimeters (TLDs). This analysis compares against measurements made during the 2015–2016 JET campaign with calibration factors calculated at the Frascati Neutron Generator (FNG) in 2017. The benchmark provides an MCNP input file, as well as the measured neutron fluence and air kerma, at 16 detector locations identified as A1–A8 and B1–B8. Figure 1 shows a visualization of the JET torus MCNP model that was created with the Radiant ray-trace–based three-dimensional image generator [16]. The MCNP model contains a detailed representation of the JET torus, but beyond the machine, the model is simplified and excludes many structures.

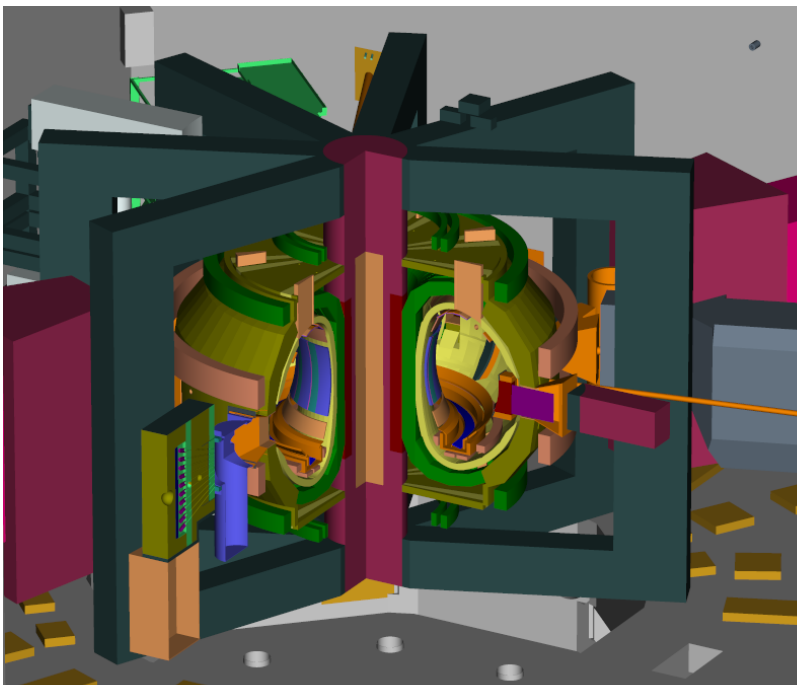

Fig. 1. Visualization of the MCNP model of the JET torus with the northeast quadrant cut away.

Figure 2 indicates the x-y positions of the detectors placed for the benchmark on a slice

through the model geometry at $z = -400$ cm. Note that these detectors are located at different elevations. As shown in Figure 2, detectors A1, B1, and A8 are located near the torus, detectors A2–A7 and B8 are located in the southwest labyrinth area, detectors B2–B5 are located in the southeast chimney, and detectors B6 and B7 are located in the basement. During the experiment, several TLDs were placed in a high density polyethylene cylindrical moderator at each detector location. To detect any directional effects, the TLDs were grouped into two orientations: a circular arrangement of horizontal TLDs, and a rectangular arrangement of vertical TLDs. The measurements from the two orientations were then averaged. The B5 measurement is not included due to placement issues during the experiment and the B7 and B8 measurements are excluded because they are near background and as a result have very large uncertainties.

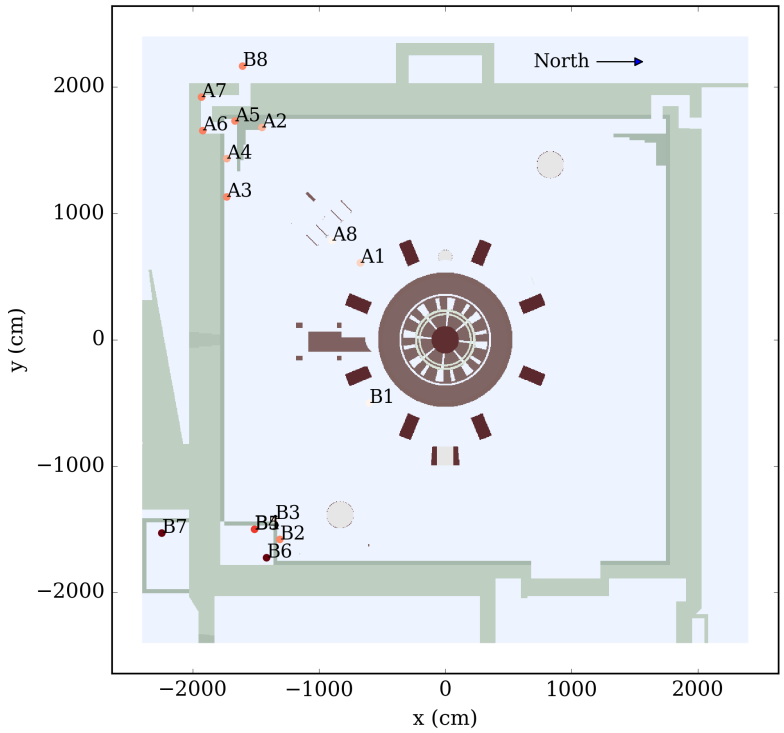

Fig. 2. Slice through the MCNP model of the JET Torus Hall at z = -400 cm, with detector locations for the 2015–2016 campaigns indicated. Note that detector locations are projected to the same elevation with darker points indicating lower elevations.

The MCNP model of the JET torus hall provided with the benchmark files does not model the TLDs or polyethylene moderators. Instead, the benchmark includes separate attenuation factors for fluence and air kerma as well as flux-to-air kerma conversion factors from ICRP-74 [17]. The

attenuation factors for the neutron air kerma are given in Table I and were estimated by calculating the reduction in air kerma when the polyethylene moderators were present for a few detector positions with MCNP. Energy-dependent attenuation factors were provided for the neutron fluence calculation.

TABLE I
Neutron air kerma attenuation factors for polyethylene moderators [4]

Detector position	Neutron kerma factor
A1, B1	0.101
A2–A8, B2–B8	0.081

IV. DENOVO PROBLEM SETUP

Modeling a continuous-energy, constructive solid geometry model using a multigroup discrete ordinates code requires discretizing the problem in space and energy. The Omnibus [18] interface uses the particle transport routines in the Shift Monte Carlo code to discretize the geometry. Because these routines require additional libraries and codes (viz. the MCNP5 executable to generate a `runtpe` file and the Lava library [19] to perform the transport), we convert the MCNP representation of the JET problem into an explicitly meshed representation as a preprocessing step so that one discretization can be reused in multiple problems run on Titan. This preprocessing step is not run on Titan, but rather on an institutional computer cluster, typically with 256 compute cores at up to a half hour of wall time. Sections IV.A and IV.B discuss the material and energy discretization.

IV.A. Material discretization

The KBA spatial decomposition discussed in Section II partitions the problem across a configurable number of computational domains, and each domain locally generates a table of material volume fractions in each spatial cell by tracing a configurable number of rays orthogonally through the problem. Tracing more rays improves the estimate of the volume fractions at the cost of additional run time for the trace; 64 rays per external Cartesian grid cell face were traced for this model. Increasing the number of rays also makes it more likely to encounter complex regions of the model where particles can be lost: error-correcting routines had to be added to the ray

tracer in order to successfully trace the JET model.

Since the memory usage requirements of the runtime-calculated cross sections scale with the number of distinct material mixtures, it is desirable to reduce the size of the volume fraction tables. These mix tables are collapsed by declaring spatial cells with similar components as equivalent, where the similarity is determined by comparing the relative volume fractions using a user specified tolerance; this analysis used a tolerance of 1%. For example, two cells containing a wall-air interface may have a slightly different volume fraction of air, but after collapsing the mix table, they may share a material identifier. This introduces additional discretization error into the Denovo approximation of the model.

At this point in the discretization process, the material identifiers and volume fractions are all local to a compute domain and based on the spatial decomposition used. The spatial domains on the institutional cluster used for this preprocessing will be considerably larger than those used on Titan since the number of processors in use is orders of magnitude smaller. Therefore, a *global* mix table and material ID field must be constructed from this local information to provide an input for use on Titan. A global scatter of the number of distinct mixtures per domain enables each domain to shift its local IDs to a global ID.

The domain-local data are written to two separate HDF5 fields using collective I/O operations. One three-dimensional integer field stores the provisional material IDs, and an array of volume fractions stores the sparse mix table in coordinate format.

After the parallel discretization is complete, a serial tool reads in the global IDs and mix tables, collapses the mix tables using an $O(N \log N)$ fuzzy sorting algorithm [20], and writes the globalized IDs and mix tables to a new HDF5 file. Collapsing reduces the number of mixtures by a factor of two in the discretization with almost a billion mesh cells, yielding a proportional memory reduction in the cross section storage that allows the problem to fit within the memory constraints of Titan.

IV.B. Energy discretization

To examine the solution's sensitivity to energy group structure, multigroup cross sections with 64, 128, and 256 energy groups were generated from the FENDL-3.1b library [21] with group boundaries based on the 175-group Vitamin J library. In practice, since the experimental

benchmark used a D-D plasma, all energy groups above the D-D source energy can be omitted in their entirety from the Denovo solve. However, maintaining the number of energy groups as a multiple of two improves load balancing on the GPUs. This resulted in reducing the three group structures to 32, 88, and 216 energy groups.

The MCNP source term is described as a separable function of space and energy, so it is only necessary to discretize the source once per mesh resolution rather than at each combination of mesh grid and multigroup library. To discretize the source, independent MCNP history source particles are sampled in parallel using the Lava library. Their accumulated positions on the mesh grid are summed over all domains to obtain the global source strength. Because of the stochastic nature of the sampling, the Denovo source term has some statistical sampling error in addition to the discretization error. However, the error introduced by the two billion source samples is no different than the equivalent Monte Carlo source sampling error.

The energy spectrum is separately and analytically discretized onto the multigroup energy bounds. This spectrum is combined with the spatial discretization and material discretization to define the final Denovo input.

V. RESULTS

These studies were performed on the Titan Cray XK7 supercomputer at the OLCF. Titan has 18,688 compute nodes, each with a 16-core AMD Opteron CPU, an NVIDIA Kepler GPU, and 32 GB of RAM. The number of nodes used for each simulation was driven by the level of decomposition needed to fit the problem on the cluster.

Results were generated from an MCNP model of the JET Torus Hall that was provided along with benchmark data. MCNP and Shift results are also provided as a computational reference. The MCNP results were calculated by the Jožef Stefan Institute (JSI) [22] using variance reduction parameters generated by ADVANTG [19] and the analog Shift results were calculated by ORNL on Titan. The Denovo and MCNP calculations used the FENDL-3.1b library [21], while Shift used the ENDF/B-VII library [23].

V.A. Sensitivity Study

The impact of several deterministic parameters on the Denovo solution were examined, including spatial mesh, quadrature set, energy groups, and anisotropic scattering order (P_N). The Denovo discretization parameters that were considered are summarized in Tables II and III. The spatial meshes were uniformly spaced and quadruple range product quadratures [24] were used for all cases as they tend to exhibit fewer ray effects than level-symmetric quadratures. All calculations used the step characteristics spatial differencing scheme described in Section II. This scheme was chosen because it is both linear and maintains positivity, regardless of the discretization parameters chosen, provided that the source is positive. As described in Section II, all energy groups are solved at once and so a brief study was conducted to examine the convergence of the flux spectrum and determined that requiring 200 solver iterations for all cases would ensure convergence.

TABLE II
Deterministic Parameters for JET Streaming Benchmark

Quadrature sets	36 angles/octant 128 angles/octant 256 angles/octant
Energy groups	32 88 216
Scattering orders	P_1 (4 moments) P_3 (16 moments)

TABLE III
Spatial Meshes for JET Streaming Benchmark

Case name	# Cells	Mesh size [x, y, z] (cm)
1m	1 572 864	[62.50, 72.92, 54.69]
10m	12 582 912	[31.25, 36.46, 27.34]
100m	88 080 384	[15.63, 15.63, 18.23]
1b	935 854 080	[7.35, 7.29, 7.81]

The thermal neutron flux calculated by Denovo on 2 048 nodes with 100 million spatial cells, 32 energy groups, 256 angles per octant, and P_3 scattering with a wall-time of 4.6 hours is shown in Figure 3 through the center of the machine ($z = 0$) with the xy detector locations indicated.

The neutron fluence spectra for detectors A1, A2, A6, and B8 are given in Figure 4 for the Denovo solution with 10 million spatial cells, 32 energy groups, 256 angles per octant, and P_3 scattering. The neutron fluence drops about 8 orders of magnitude from detector A1 to B8.

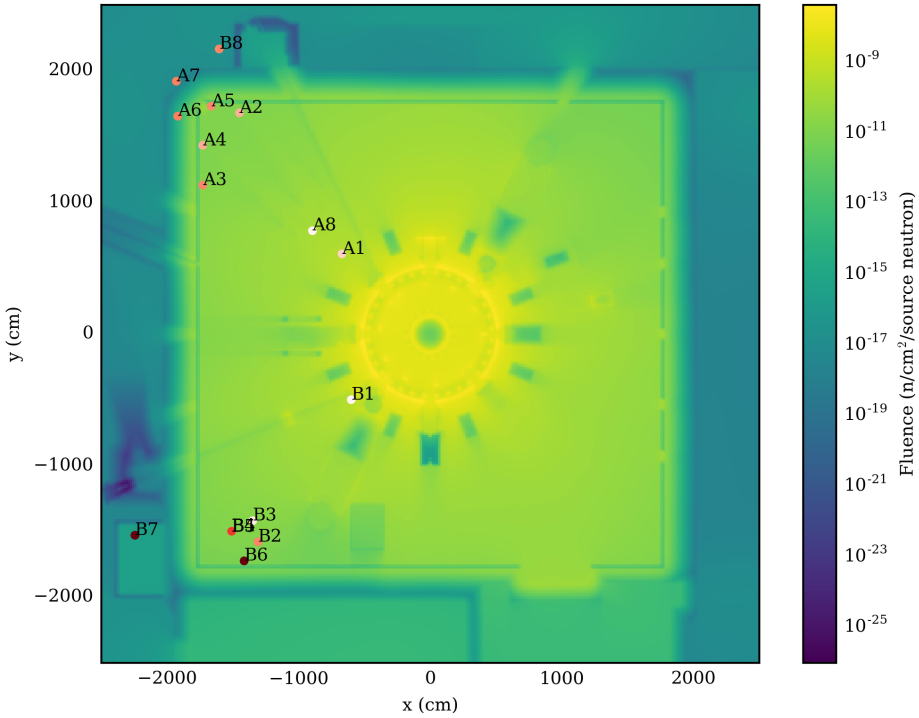

Fig. 3. Thermal neutron fluence (10^{-5}-10^{-1} eV) at $z = 0$ for 100 million spatial cells, 32 energy groups, 256 angles per octant, and P_3 scattering with detector locations projected to this elevation.

The Denovo solution to the JET streaming benchmark proved to be largely insensitive to the deterministic parameters examined, excepting spatial mesh. Varying the number of angles per octant, the number of energy groups, and the scattering order did not result in a significant change in the calculated neutron fluence or kerma at the detector locations. Instead, spatial mesh size had the dominant effect on the solution, indicating the importance of resolving the narrow streaming paths through the labyrinth and chimney. With further mesh refinement in these areas, it is expected that the impact of the quadrature set would become more apparent.

The uniformly spaced meshes that were considered are described in Table III. Figure 5 gives the neutron fluence at each detector location for varying numbers of spatial cells. The neutron fluence calculated using MCNP with ADVANTG is also given for comparison, and all MCNP relative errors are less than 2%. In general, the Denovo solution under predicts the MCNP solution

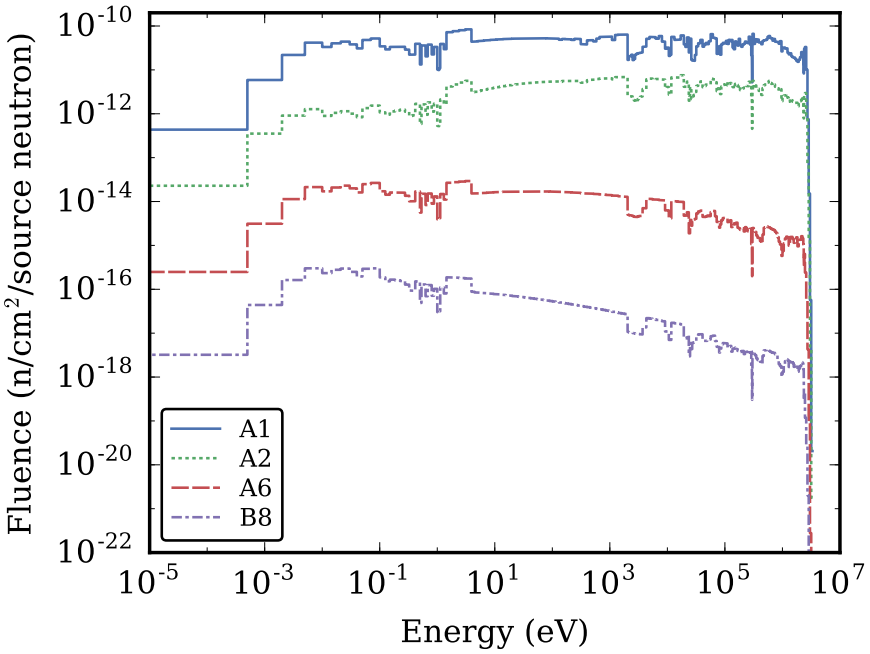

Fig. 4. Neutron fluence spectra calculated by Denovo with 10 million spatial cells, 216 energy groups, 256 angles per octant, and P_3 scattering at the locations of detectors A1, A2, A6, and B6.

and improves with mesh refinement. The Denovo and MCNP solutions match best near the source and larger differences are observed as the detector locations move farther from the machine. More studies with localized mesh refinement are needed to resolve these differences. The solution for detector B8, which is located beyond the southwest labyrinth, is inaccurate with the coarser 1 million and 10 million cell meshes and MCNP data was not available for this location.

Figure 6 illustrates the neutron kerma calculated by Denovo at each detector location for varying numbers of spatial cells. MCNP results for the neutron kerma were not available so Shift results are provided as a computational reference. As with the neutron fluence, the Denovo solution improves with mesh refinement and larger differences are observed in detectors located farther from the source.

The wall times and nodes used for several Denovo simulations are given in Table IV. The number of nodes used was generally determined by the problem size, but the smallest cases were also impacted by wall time limits on Titan that restrict jobs below certain node counts.

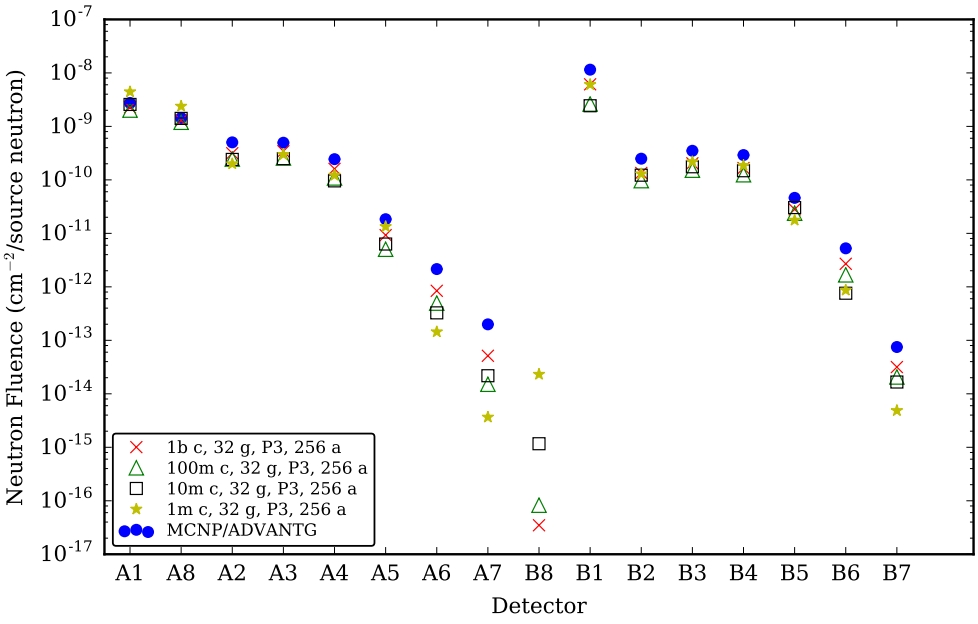

Fig. 5. Neutron fluence calculated values for different mesh sizes in Denovo.

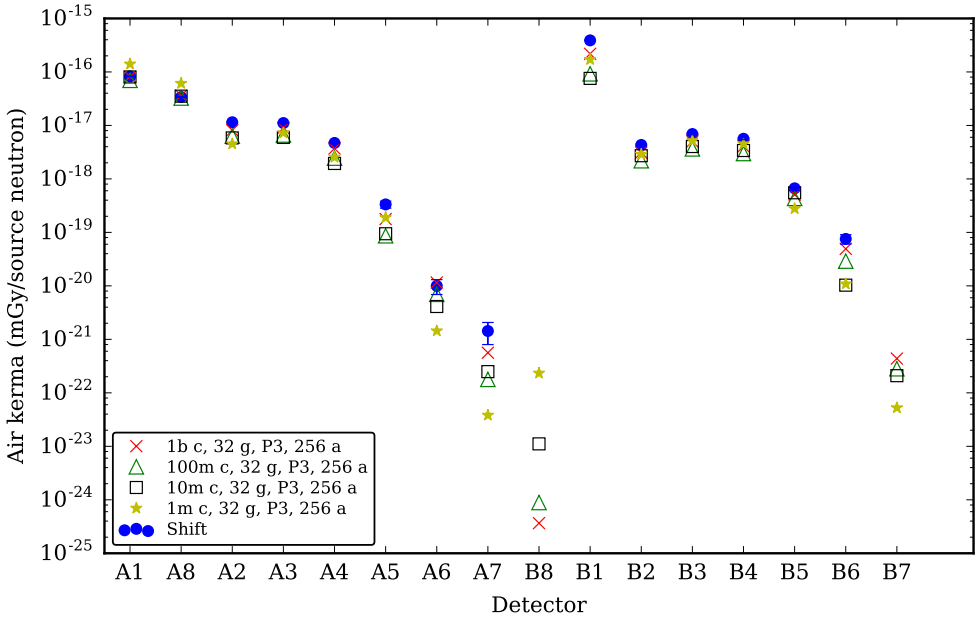

Fig. 6. Neutron air kerma calculated values for different mesh sizes in Denovo.

TABLE IV
Wall times and node usage for selected cases.

Case name	Nodes used (cores[a])	Wall time
1m c, 32 g, P_3, 256 a	128 (3 840)	1.3 hrs
10m c, 32 g, P_3, 256 a	512 (15 360)	2.5 hrs
10m c, 88 g, P_3, 256 a	1024 (30 720)	1.7 hrs
10m c, 216 g, P_3, 256 a	2048 (61 440)	1.5 hrs
100m c, 32 g, P_3, 256 a	2 048 (61 440)	4.6 hrs
1b c, 32 g, P_1, 256 a	3 840 (115 200)	10.0 hrs

[a] The OLCF defines each Titan XK7 node as possessing 30 total cores for the purposes of resource accounting.

V.B. Benchmark Comparison

Measurements from the 2015-2016 DD campaign at JET with calibration factors calculated at FNG in 2017 were used for comparison with Denovo. Figure 7 gives the neutron fluence that was measured at each detector and calculated by Denovo, MCNP with ADVANTG, and Shift. The neutron air kerma measured by the experiment and calculated by Denovo and Shift at each detector is given in Figure 8. The Denovo results calculated with the 1b (935 854 080) cell spatial mesh, 32 energy groups, 256 angles per octant, and P_3 scattering were used for comparison with experiment results.

The measurements and calculations show good general agreement. The neutron fluence at detectors B7 and B8 is near background levels and has not been included due to large experimental uncertainties. The calculated neutron fluence and air kerma underestimate the experiment at detector A1, which is located near the source. This was also observed in Ref. [4], which notes that uncertainty in the exact detector position or shielding equipment not represented in the model may be the cause. The MCNP and Shift solutions agree well, with some differences likely due to the different cross-section libraries used.

The ratios of calculated-to-experimental (C/E) neutron fluence and air kerma for Denovo, MCNP with ADVANTG, and Shift are given in Figure 9 and Figure 10, respectively. In Figure 9, all three calculations show a tendency to overestimate the neutron fluence with increasing distance from the machine, which was also observed in a previous comparison with MCNP [4]. The Denovo neutron air kerma C/E is shown to underestimate the experiment at all detector locations. The accuracy of the simulation may be limited by the air kerma attenuation factors described in

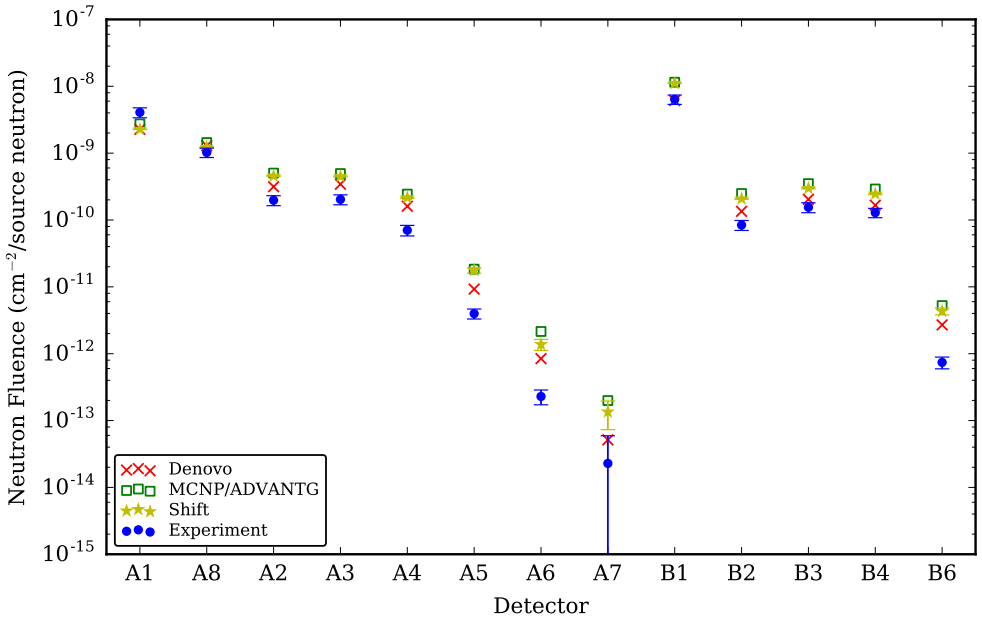

Fig. 7. Neutron fluence measured and calculated by Denovo, MCNP with ADVANTG, and Shift.

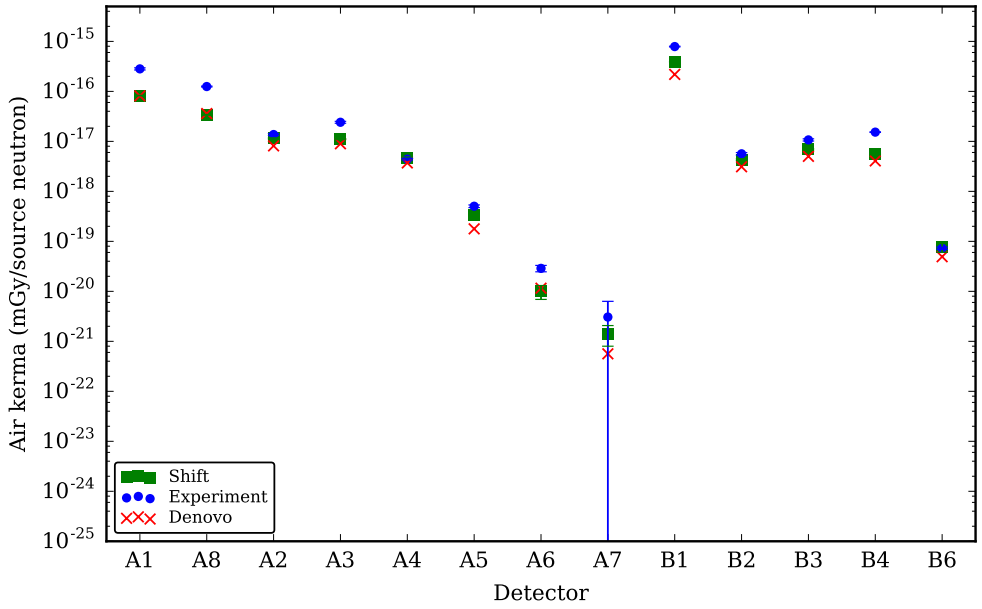

Fig. 8. Neutron air kerma measured and calculated by Denovo and Shift.

Section III and further work is needed to fully understand some of these differences. In particular, future work should examine the impact of localized mesh refinement. It is also important to note that there are several sources of uncertainty in these ratios, in addition to the experimental errors, statistical errors in Monte Carlo calculations, and discretization errors in the deterministic calculations. Other sources of error include uncertainty in the attenuation coefficients for the polyethylene cylinders, approximations in the MCNP model, and inaccurate material compositions for concrete.

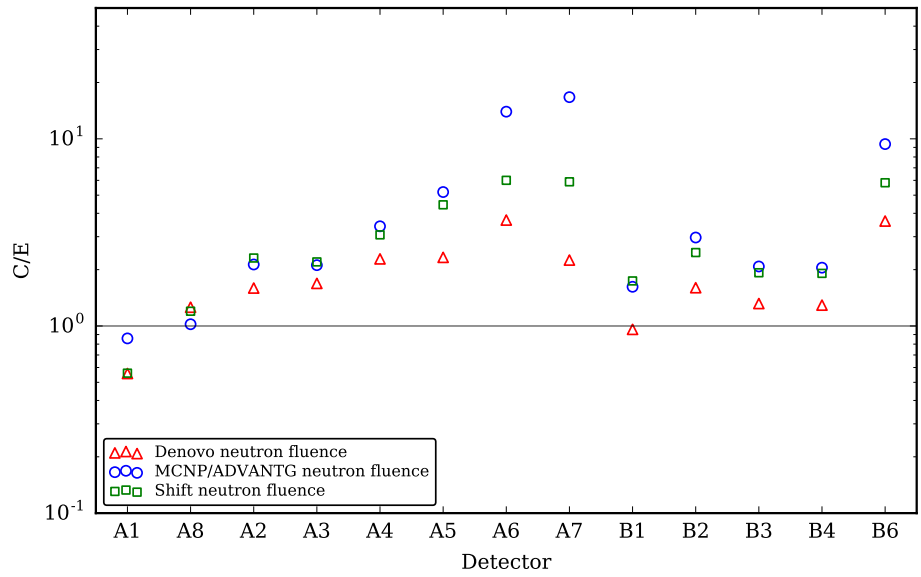

Fig. 9. Ratio of calculated (C) over experiment (E) neutron fluence for Denovo, MCNP with ADVANTG, and Shift.

VI. CONCLUSIONS

The highly scalable Denovo discrete ordinates code successfully modeled the JET streaming benchmark at an unprecedented level of fidelity, taking 10 hours of wall time on 3 840 nodes of the Titan computer at OLCF to solve a discretization with 1 billion mesh cells, P_1 scattering, 32 energy groups, and 36 angles per octant on GPUs. A parameter study determined the sensitivity of the computed results as a function of the discretization parameters; of these, the spatial mesh size had the strongest effect on the calculated detector responses.

The Denovo-calculated neutron fluence and air kerma showed good general agreement with

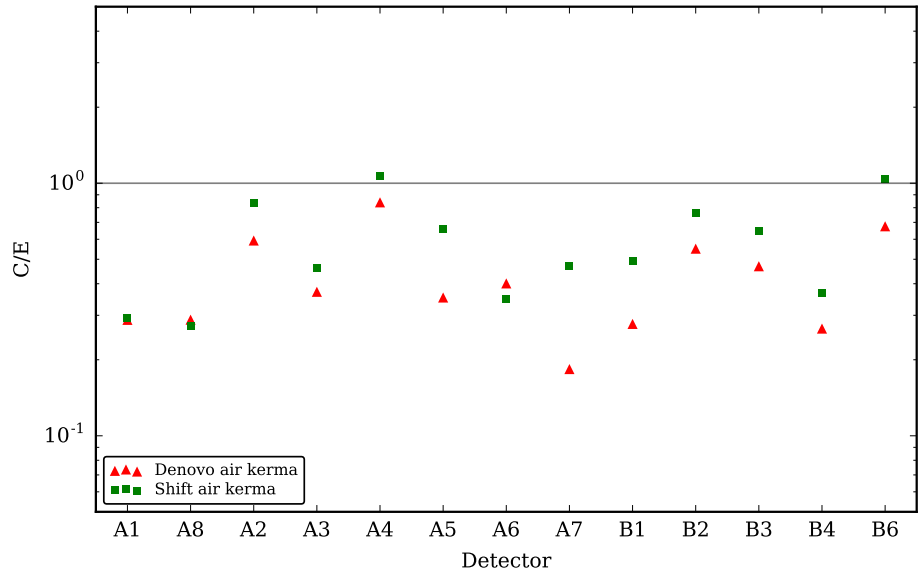

Fig. 10. Ratio of calculated (C) over experiment (E) neutron air kerma for Denovo and Shift.

the experimental results and previously simulated MCNP neutron fluence results. Further work is needed to fully explain and resolve some observed differences. Additional results generated by the massively parallel Shift Monte Carlo code on Titan provided further confidence in the ability to accurately model this complex fusion system using the deterministic Denovo radiation transport code. The Denovo-calculated flux distribution would also be effective for variance reduction of Monte Carlo simulations, and coarser meshes could be used for this application.

ACKNOWLEDGMENTS

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

The authors would like to express their thanks to Charles Daily of the Reactor and Nuclear Systems Division of Oak Ridge National Laboratory for generating the multigroup cross-section libraries used in this work.

REFERENCES

[1] T. M. Evans, A. S. Stafford, R. N. Slaybaugh, and K. T. Clarno, "Denovo: A New Three-Dimensional Parallel Discrete Ordinates Code in SCALE," *Nuclear Technology*, **171**, *2*, 171 (2010); 10.13182/NT171-171., URL http://dx.doi.org/10.13182/NT171-171.

[2] A. M. Ibrahim, S. W. Mosher, T. M. Evans, D. E. Peplow, M. E. Sawan, P. P. Wilson, J. C. Wagner, and T. Heltemes, "ITER Neutronics Modeling Using Hybrid Monte Carlo/Deterministic and CAD-based Monte Carlo Methods," *Nuclear Technology*, **175**, *1*, 251 (2011)URL http://www.ans.org/pubs/journals/nt/a_12296.

[3] A. Ibrahim, M. Sawan, S. Mosher, T. Evans, D. Peplow, P. Wilson, and J. Wagner, "Global evaluation of prompt dose rates in ITER using hybrid Monte Carlo/deterministic techniques," *Fusion Science and Technology*, **60**, *2*, 676 (2011).

[4] P. Batistoni, S. Conroy, S. Lilley, J. Naish, B. Obryk, S. Popovichev, I. Stamatelatos, B. Syme, T. Vasilopoulou, and JET contributors, "Benchmark experiments on neutron streaming through JET Torus Hall penetrations," *Nuclear Fusion*, **55**, *5*, 053028 (2015)URL http://stacks.iop.org/0029-5515/55/i=5/a=053028.

[5] X-5 Monte Carlo Team, "MCNP - A General Monte Carlo N-Particle Transport Code, Version 5," LA-UR-03-1987, Los Alamos National Laboratory (2008).

[6] T. M. Pandya, S. R. Johnson, T. M. Evans, G. G. Davidson, S. P. Hamilton, and A. T. Godfrey, "Implementation, capabilities, and benchmarking of Shift, a massively parallel Monte Carlo radiation transport code," *Journal of Computational Physics*, **308**, *Supplement C*, 239 (2016); https://doi.org/10.1016/j.jcp.2015.12.037., URL http://www.sciencedirect.com/science/article/pii/S0021999115008566.

[7] G. G. Davidson, T. M. Evans, J. J. Jarrell, S. P. Hamilton, T. M. Pandya, and R. N. Slaybaugh, "Massively Parallel, Three-Dimensional Transport Solutions for the k-Eigenvalue Problem," *Nuclear Science and Engineering*, **177**, *2*, 111 (2014); 10.13182/NSE12-101., URL http://dx.doi.org/10.13182/NSE12-101.

[8] J. J. Jarrell, T. M. Evans, G. G. Davidson, and A. T. Godfrey, "Full Core Reactor Analysis: Running Denovo on Jaguar," *Nuclear Science and Engineering*, **175**, *3*, 283 (2013); 10.13182/NSE12-60., URL http://dx.doi.org/10.13182/NSE12-60.

[9] C. Baker, G. Davidson, T. M. Evans, S. Hamilton, J. Jarrell, and W. Joubert, "High performance radiation transport simulations: Preparing for TITAN," *High Performance Computing, Networking, Storage and Analysis (SC), 2012 International Conference for*, 1–10 (2012); 10.1109/SC.2012.64.

[10] T. M. Evans, W. Joubert, S. P. Hamilton, S. R. Johnson, J. A. Turner, G. G. Davidson, and T. M. Pandya, *Three-dimensional discrete ordinates reactor assembly calculations on GPUs*.

[11] R. Slaybaugh and T. Evans, "Multigrid in Energy Preconditioner for Krylov Solvers," *Journal of Computational Physics*, **242**, 405 (2013).

[12] J. Warsa, T. Wareing, and J. Morel, "Krylov Iterative Methods and the Degraded Effectiveness of Diffusion Synthetic Acceleration for Multidimensional S_N Calculations in Problems with Material Discontinuities," *Nuclear Science and Engineering*, **147**, 218 (2004).

[13] S. Pautz, "An Algorithm for Parallel S_n Sweeps on Unstructured Meshes," *Nuclear Science and Engineering*, **140**, 111 (2002).

[14] O. R. L. C. Facility, "Titan Cray XK7," (2014)URL https://www.olcf.ornl.gov/computing-resources/titan-cray-xk7/.

[15] P. Batistoni, J. Likonen, N. Bekris, S. Brezinsek, P. Coad, L. Horton, G. Matthews, M. Rubel, G. Sips, B. Syme, and A. Widdowson, "The JET technology program in support of ITER," *Fusion Engineering and Design*, **89**, 7, 896 (2014); https://doi.org/10.1016/j.fusengdes.2013.12.050., URL http://www.sciencedirect.com/science/article/pii/S0920379613007643, proceedings of the 11th International Symposium on Fusion Nuclear Technology-11 (ISFNT-11) Barcelona, Spain, 15-20 September, 2013.

[16] S. Mosher, S. Wilson, A. Bevill, and J. Risner, "ADVANTG Improvements for Fusion Neutronics Applications," *XIth ITER Neutronics Meeting*, Cadarache, France (2017).

[17] ICRP, "Conversion Coefficients for use in Radiological Protection against External Radiation," 26, ICRP Publication 74 Ann. ICRP (1996).

[18] S. R. Johnson, "Omnibus: a New Front End to Denovo and Shift," *Transactions of the American Nuclear Society*, **117** (2017).

[19] S. W. Mosher, A. M. Bevill, S. R. Johnson, A. M. Ibrahim, C. R. Daily, T. M. Evans, J. C. Wagner, and J. O. Johnson, "ADVANTG - an automated variance reduction parameter generator," ORNL/TM-2013/416, Oak Ridge National Laboratory (2013).

[20] S. R. Johnson, "Fast Mix Table Construction for Material Discretization," *Proc. Intl. Conf. on Math. and Comput. Methods Applied to Nucl. Sci. Eng. (M&C 2013), Sun Valley, ID, May 5–9* (2013)[CD-ROM].

[21] R. A. Forrest, R. Capote, N. Otsuka, T. Kawano, A. J. Koning, S. Kunieda, J.-C. Sublet, and Y. Watanabe, "FENDL-3 Library - Summary Documentation," INDC(NDS)-0628, IAEA (2012)URL https://www-nds.iaea.org/fendl/.

[22] B. Kos, S. Mosher, I. Kodeli, R. Grove, J. Naish, B. Obryk, R. Villari, P. Batistoni, and JET contributors, "Application of the ADVANTG Hybrid Code on the JET3–NEXP streaming benchmark experiment," *to be submitted*.

[23] M. B. Chadwick, P. Obložinsky, M. Herman, N. M. Greene, R. D. McKnight, D. L. Smith, and e. a. P. G. Young, "ENDF/B-VII.0: Next Generation Evaluated Nuclear Data Library for Nuclear Science and Technology," *Nuclear Data Sheets*, **107** (2006).

[24] I. Abu-Shumays, "Angular quadratures for improved transport computations," *Transport Theory and Statistical Physics*, **30**, *2-3*, 169 (2001).

[25] E. Lewis and J. W.F. Miller, *Computational Methods of Neutron Transport*, American Nuclear Society, Inc., LaGrange Park, IL (1993).

A. DISCRETE ORDINATES EQUATIONS IN DENOVO

The fixed-source steady-state Boltzmann transport equation solved in Denovo is

$$\hat{\mathbf{\Omega}} \cdot \nabla \psi(\mathbf{r}, \mathbf{\Omega}, E) + \sigma(\mathbf{r}, E)\psi(\mathbf{r}, \mathbf{\Omega}, E) = \int_0^\infty \int_{4\pi} \sigma_{\mathrm{s}}(\mathbf{r}, \mathbf{\Omega}' \to \mathbf{\Omega}, E' \to E)\psi(\mathbf{r}, \mathbf{\Omega}', E')\, d\mathbf{\Omega}' dE' + q_e(\mathbf{r}, \mathbf{\Omega}, E)\,, \quad (1)$$

and it obeys a boundary condition that defines the incoming flux on problem boundaries with outgoing normal $\mathbf{n}$,

$$\psi(\mathbf{r}, \mathbf{\Omega}, E) = \Gamma\,, \quad \mathbf{r} \in \partial V\,, \quad \hat{\mathbf{\Omega}} \cdot \mathbf{n} < 0\,. \quad (2)$$

In Eqs. (1) and (2), the state is defined by the angular flux ψ, and the independent variables are $\mathbf{r} = (x, y, z)$ in cm, $\mathbf{\Omega} = (\theta, \varphi)$ in steradians (sr), and E in MeV representing space, angle, and energy respectively. Applying the discrete ordinates (S_N) angular and multigroup energy discretizations, Eq. (1) becomes

$$\hat{\mathbf{\Omega}}_a \cdot \nabla \psi_a^g(\mathbf{r}) + \sigma^g(\mathbf{r})\psi_a^g(\mathbf{r}) = \sum_{g'=1}^{G} \sum_{l=0}^{N} \sigma_{\mathrm{s}\,l}^{gg'}(\mathbf{r}) \Big[Y_{l0}^e(\mathbf{\Omega}_a)\phi_{l0}^{g'}(\mathbf{r}) + \sum_{m=1}^{l} \left(Y_{lm}^e(\mathbf{\Omega}_a)\phi_{lm}^{g'}(\mathbf{r}) + Y_{lm}^o(\mathbf{\Omega}_a)\vartheta_{lm}^{g'}(\mathbf{r}) \right) \Big] + q_e^g(\mathbf{r}, \mathbf{\Omega}_a)\,, \quad a = 1, \ldots, n\,,\ g = 1, \ldots, G\,, \quad (3)$$

where $\psi_a^g \equiv \psi^g(E^g, \mathbf{\Omega}_a)$. The functions Y_{lm}^e and Y_{lm}^o are the even and odd spherical harmonics defined to ensure a real expansion of the angular flux. In the S_N method, the angles are integrated by a quadrature rule yielding the following forms for the angular flux moments on the right-hand side of Eq. (3)

$$\phi_{lm}^g = \sum_{a=1}^{n} Y_{lm}^e(\mathbf{\Omega}_a)\psi_a^g w_a\,, \quad (4)$$

$$\vartheta_{lm}^g = \sum_{a=1}^{n} Y_{lm}^o(\mathbf{\Omega}_a)\psi_a^g w_a\,. \quad (5)$$

Finally, $\sigma_{\mathrm{s}\,l}^{gg'}$ is the lth moment of the Legendre polynomial expansion of the scattering cross section from group $E^{g'} \to E^g$. This form defines the scattering as a polynomial function in angle enabling the use of arbitrary quadratures for the same set of multigroup cross section data.

To parameterize the solution methods in Denovo, Eq. (3) is written in operator form,

$$\mathbf{L}\Psi = \mathbf{MS}\Phi + \mathbf{q}_e \,, \tag{6}$$

and the angular flux moments, Eq. (4) and (5), are written

$$\Phi = \mathbf{D}\Psi \,. \tag{7}$$

The size of the problem is determined by the number of groups (G), the number of angular flux moments required by the P_N scattering order (t), the number of quadrature angles (n), and the number of spatial cells (N_c). The number of discrete angular fluxes (Ψ) that are solved is $N_g \times n \times N_c$,

$$\Psi = \begin{pmatrix} \psi_1^1 & \psi_2^1 & \cdots & \psi_n^1 & \cdots & \psi_n^G \end{pmatrix}^T \,. \tag{8}$$

The persistent state vector is comprised of the angular flux moments (Φ), and has size $N_g \times t \times N_c$,

$$\Phi = \begin{pmatrix} \phi_{00}^1 & \phi_{10}^1 & \vartheta_{11}^1 & \phi_{11}^1 & \cdots & \vartheta_{NN}^1 & \phi_{NN}^1 & \cdots & \vartheta_{NN}^G & \phi_{NN}^G \end{pmatrix}^T \,. \tag{9}$$

Operating on Eq. (6) by $\mathbf{T} = \mathbf{DL}^{-1}$, applying Eq. (7), and rearranging terms, the fixed-source problem is

$$(\mathbf{I} - \mathbf{TMS})\Phi = \mathbf{Tq}_e \,. \tag{10}$$

This equation can be written in the standard form $\mathbf{Ax} = \mathbf{b}$, where the application of operator $\mathbf{A}$ requires the following steps:

$$\begin{aligned} \text{matrix-vector multiply}: &\quad y = \mathbf{MS}v \,, \\ \text{sweep}: &\quad z = \mathbf{DL}^{-1}y \,, \\ \text{return}: &\quad \Leftarrow v - z \,. \end{aligned}$$

The application of the transport operator, $\mathbf{L}^{-1}$ in Eq. (10), is performed by traversing through the mesh in the direction of particle travel [25]. Each so-called *sweep* requires n traversals through the mesh; therefore, every Krylov iteration requires n traversals of the mesh in the direction of

particle travel, and the solution of the transport equation is defined by the efficiency of the sweep algorithm.