ANS MC2015 - Joint International Conference on Mathematics and Computation (M&C), Supercomputing in Nuclear Applications (SNA) and the Monte Carlo (MC) Method · Nashville, Tennessee · April 19–23, 2015, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2015)

ACCELERATING FUSION REACTOR NEUTRONICS MODELING BY AUTOMATIC COUPLING OF HYBRID MONTE CARLO/DETERMINISTIC TRANSPORT ON CAD GEOMETRY

Elliott Biondo
University of Wisconsin at Madison
1500 Engineering Drive
Madison, WI 53706
biondo@wisc.edu

Ahmad M. Ibrahim, Scott W. Mosher, and Robert E. Grove
Oak Ridge National Laboratory *
1 Bethel Valley Road
Oak Ridge, TN 37831
ibrahimam@ornl.gov; moshersw@ornl.gov; grovere@ornl.gov

ABSTRACT

Detailed radiation transport calculations are necessary for many aspects of the design of fusion energy systems (FES) such as ensuring occupational safety, assessing the activation of system components for waste disposal, and maintaining cryogenic temperatures within superconducting magnets. Hybrid Monte Carlo (MC)/deterministic techniques are necessary for this analysis because FES are large, heavily shielded, and contain streaming paths that can only be resolved with MC. The tremendous complexity of FES necessitates the use of CAD geometry for design and analysis. Previous ITER analysis has required the translation of CAD geometry to MCNP5 form in order to use the AutomateD VAriaNce reducTion Generator (ADVANTG) for hybrid MC/deterministic transport. In this work, ADVANTG was modified to support CAD geometry, allowing hybrid (MC)/deterministic transport to be done automatically and eliminating the need for this translation step. This was done by adding a new ray tracing routine to ADVANTG for CAD geometries using the Direct Accelerated Geometry Monte Carlo (DAGMC) software library. This new capability is demonstrated with a prompt dose rate calculation for an ITER computational benchmark problem using both the Consistent Adjoint Driven Importance Sampling (CADIS) method an the Forward Weighted (FW)-CADIS method. The variance reduction parameters produced by ADVANTG are shown to be the same using CAD geometry and standard MCNP5 geometry. Significant speedups were observed for both neutrons (as high as a factor of 7.1) and photons (as high as a factor of 59.6).

Key Words: Hybrid Monte Carlo/Deterministic, CAD-Based Monte Carlo, Variance Reduction

1 INTRODUCTION

Radiation transport simulations are instrumental in the design of fusion energy systems (FES) to ensure safe and effective operation. In FES, high-energy neutrons activate system components. The nuclear inventory of these components must be quantified for waste disposal. The ITER experimental fusion device uses superconducting magnets to confine the plasma. The nuclear heating within these magnets must be estimated in order to design systems to maintain these magnets at cryogenic temperatures. In addition, the dose rate resulting from neutrons and photons during operation and photons after shutdown must be quantified to ensure occupational safety.

The high degree of accuracy required for these calculations necessitates the use of Monte Carlo (MC) radiation transport, which allows for high-fidelity representations of geometry, accurate resolution of particle streaming, and continuous-energy treatment of particle interactions. The deep-penetration shielding problems encountered in FES—which tend to be large and heavily shielded (blanket, vacuum vessel, bioshield)—are prohibitively computationally expensive without the use of MC variance reduction. Variance reduction techniques bias the probability distribution functions that describe particle behavior in order to preferentially sample behavior that results in tally scores, accelerating the convergence of results.

In contrast, deterministic codes can provide computationally inexpensive results with a sacrifice in accuracy due to spatial, angular, and energy discretization. Hybrid MC/deterministic transport methods use deterministic estimates of the forward and/or adjoint flux to generate variance reduction parameters for MC transport. The Consistent Adjoint Driven Importance Sampling (CADIS) method [1] uses a deterministic estimate of the adjoint flux to generate MC weight windows and source biasing parameters to optimize MC transport with respect to a tally. The Forward-Weighted (FW)-CADIS method [2] uses an additional deterministic estimate of the forward flux to generate an adjoint source to be used with the CADIS method to simultaneously optimize multiple tallies. These methods have been shown to drastically improve the efficiency of MC simulations for large-scale fusion applications [3].

The AutomateD VAriaNce reducTion Generator (ADVANTG) [4] automates the process of performing the CADIS and FW-CADIS methods for the MCNP5 Monte Carlo radiation transport code [5]. ADVANTG reads standard MCNP5 input files and uses ray tracing to discretize the materials and tallies onto a 3D mesh. ADVANTG then uses the Denovo [6] 3D Cartesian S_N code to calculate deterministic estimates of the forward and/or adjoint flux. CADIS or FW-CADIS is then applied to produce a spatial- and energy- dependent weight window mesh in the form of an MCNP5 WWINP (weight window input) file and MCNP5 SB (source bias) cards.

In MCNP5, a text-based combinatorial solid geometry (CSG) language is used to define geometry cells and surfaces. However, due to the tremendous complexity of FES, geometries are typically created using Computer Aided Design (CAD) software. In many cases, CAD models already exist for fluid mechanics, heat transfer, or structural mechanics calculations, so it is desirable to use these same models for neutronics analysis as well. The process of translating CAD models into MCNP5 CSG by hand requires extraordinary human effort. Tools such as McCad [7] have

been created to automate this process. However, MCNP5 does not support surfaces beyond order 2 (aside from axis-aligned tori), so modifications to CAD geometry may be required for translation. In addition, translated geometries often have many more geometry cells then a standard MCNP5 model, which can lead to performance issues. It is therefore highly favorable to perform analysis directly on CAD geometry.

Radiation transport can be performed directly on CAD geometry using DAG-MCNP5, a modified version of MCNP5 that uses the Direct Accelerated Geometry Monte Carlo (DAGMC) [8] software library, a component of the Mesh Oriented dataABase (MOAB) toolkit [9]. DAG-MCNP5 has been used for accurate and high-resolution analysis of FES such as ARIES, HAPL, and ITER [10]. Until now, ADVANTG has only supported CSG geometry representations. Previous use of ADVANTG for ITER analysis required the use of McCad to translate CAD models into MCNP5 input files [3].

In this work, ADVANTG has been modified to support DAG-MCNP5. This allows for the automatic generation of variance reduction parameters, using the CADIS and FW-CADIS methods, directly from CAD geometry. This is done by adding a new ray tracing routine to ADVANTG, which queries a DAG-MCNP5 representation of a geometry in order to perform material and tally discretization. This new capability is demonstrated with a prompt dose rate calculation for an ITER computational benchmark problem. The discretized material meshes generated directly from CAD geometry, as well as the resulting weight window meshes, are shown to be consistent with those produced using MCNP5 CSG geometry with some minor discrepancies, which can be explained by the differences in the geometry representation. Weight windows and source biasing parameters generated using the CADIS and FW-CADIS methods directly from CAD geometry are demonstrated to accelerate the MC simulations and results are shown to be statistically equivalent to results obtained without variance reduction.

2 METHODOLOGY

2.1 Implementation

ADVANTG accomplishes the automatic coupling of MC and deterministic transport by discretizing the MC geometry onto a superimposed Cartesian mesh, resulting in a mesh of materials, sources, and tallies that can be used to run Denovo. In order to do this, the volume fraction of each MC geometry cell within each volume element of the superimposed mesh must be estimated. This is done using a stochastic ray tracing method [4]. Rays are randomly fired down rows of mesh volume elements (in any or all of the x, y, z directions, as specified by the user). For each ray (i), the track length through each MC geometry cell (j) is tallied within each mesh volume element (k) in the mesh row. The track length of i through j within k is $l_{i,j,k}$. The total track length of i through k is $l_{i,k}$. The ratio

$$x_{i,j,k} = \frac{l_{i,j,k}}{l_{i,k}} \tag{1}$$

constitutes an independent random sample of the volume fraction of geometry cell j within mesh volume element k. This volume fraction is denoted by $v_{j,k}$. After firing N randomly sampled rays, $v_{j,k}$ can be estimated by

$$v_{j,k} = \frac{1}{N} \sum_{i=0}^{N-1} x_{i,j,k}, \tag{2}$$

with the associated relative error (for large N)

$$R_{j,k} = \sqrt{\frac{\sum_{i=0}^{N-1} x_{i,j,k}^2}{(\sum_{i=0}^{N-1} x_{i,j,k})^2} - \frac{1}{N}}. \tag{3}$$

In ADVANTG, samples of $x_{i,j,k}$ are obtained by firing rays on the MCNP5 CSG geometry using an MCNP5 wrapper library known as Lava [4]. Material fractions can then be obtained by applying the known mapping of geometry cells to materials.

The DAGMC geometry representation is fundamentally different from CSG. As a preprocessing step, DAGMC decomposes the CAD representation of the geometry into a collection of triangular facets using the Common Geometry Module (CGM) [11]. This process is approximate for curved surfaces, which are faceted using a user-supplied faceting tolerance. An example of this faceted representation is shown in Figure 2 in Section 3. Facets are arranged in a hierarchical tree of bounding boxes. DAGMC provides a library of functions that can be used for ray tracing. The functions perform geometry queries by traversing this tree of bounding boxes.

In this work, a new ray tracer was incorporated into ADVANTG that obtains samples of $x_{i,j,k}$ by firing rays using these DAGMC functions. An option to use this functionality was then incorporated into the ADVANTG user interface.

2.2 Assessment of Results

In order to test the CAD geometry discretization capability, results produced from the CAD geometry were compared to those from MCNP5 CSG. The average discrepancy in material fraction produced by these two methods is defined as

$$d_{mat} = \frac{1}{N_{mat}} \sum_{m=0}^{N_{mat}-1} |V_{m,\,\text{DAG-MCNP5}} - V_{m,\,\text{CSG}}|, \tag{4}$$

where N_{mat} is the number of materials within a given mesh volume element, $V_{m,\,\text{DAG-MCNP5}}$ is the volume fraction of material m as reported by ADVANTG using the DAG-MCNP5 geometry, and

$V_{m,\,\text{CSG}}$ is the volume fraction of material m using MCNP5 CSG geometry. A d_{mat} of 0 denotes equivalent mixtures. Likewise, the relative discrepancy in the resulting weight window lower bounds that is produced using CADIS or FW-CADIS is defined as

$$d_w = \frac{|w_{\text{CSG}} - w_{\text{DAG-MCNP5}}|}{w_{\text{CSG}}}, \tag{5}$$

where w denotes the weight window lower bounds for weight windows produced from the MCNP5 CSG geometry and the DAG-MCNP5 geometry, for a given mesh volume element and energy group.

In order to assess the efficacy of variance reduction parameters produced from CAD geometry, the MC Figure of Merit (FOM) is used, as defined by

$$\text{FOM} = \frac{1}{R^2 T}, \tag{6}$$

where R is the relative uncertainty of a tallied result and T is the execution time. More efficient MC simulations result in higher FOM. The *speedup* describes how much faster an MC simulation can be run with variance reduction compared to analog to achieve some fixed R. Here *analog* refers to transport performed without any variance reduction except implicit capture. The speedup is defined as

$$\text{speedup} = \frac{\text{FOM}_{\text{VR}}}{\text{FOM}_{\text{analog}}}, \tag{7}$$

where the subscript VR denotes the FOM resulting from the use of variance reduction.

In order to assess the reliability of dose rates estimated from transport with variance reduction (D_{VR}) in respect to dose rates from analog transport (D_{analog}), the standard error of the difference of the dose rates, S_{diff}, is computed. This is shown in Equation 8:

$$S_{\text{diff}} = S_{D_{\text{analog}} - D_{\text{VR}}} = \sqrt{S^2_{D_{\text{analog}}} + S^2_{D_{\text{VR}}}}, \tag{8}$$

where $S_{D_{\text{analog}}}$ and $S_{D_{\text{VR}}}$ are the standard errors in the dose rates from transport in analog and with variance reduction, respectively. If D_{analog} and D_{VR} are statistically equivalent, $D_{\text{analog}} - D_{\text{VR}}$ should be less than $1S_{\text{diff}}$ 68.3% of the time; $2S_{\text{diff}}$ 95.4% of the time; and $3S_{\text{diff}}$ 99.7% of the time, (i.e. a Gaussian distribution).

3 TEST PROBLEM

The ITER shutdown dose rate (SDDR) benchmark problem was designed to emulate the irradiation of an ITER diagnostic port [12]. The problem consists of an isotropic 14 MeV neutron source defined uniformly on a disk in vacuum. Four tally cells reside on the other end of a cylindrical shielding region, as shown in Figure 1.

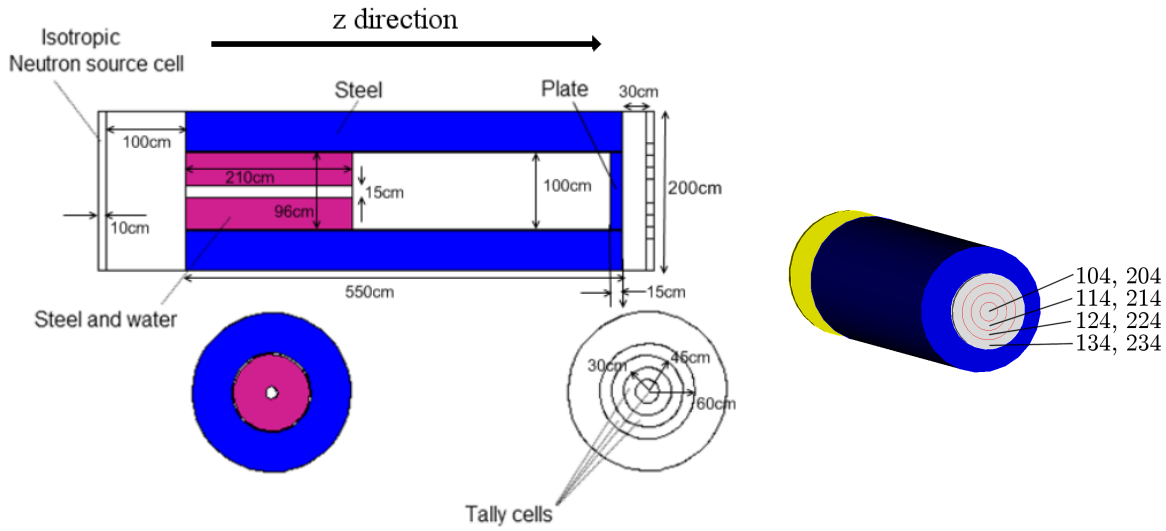

(a) Original geometry specification

(b) Tally assignments

Figure 1. Original geometry specification, taken from [12], and tally assignments for neutron (104, 114, 124, 134) and photon (204, 214, 224, 234) tallies used within this work.

Though the benchmark was intended to be used as a shutdown dose rate problem, in this work the prompt dose rate from neutrons and photons was estimated during the first segment of the irradiation scenario specified in the benchmark. A standard MCNP5 CSG model of this geometry was created, and a CAD model was created in CUBIT [13]. The CAD geometry was then faceted with a faceting tolerance of 1×10^{-4} cm. CAD and faceted models are shown in Figure 2. An MCNP5 source was made to match the problem description. This source was divided into four axial bins and 16 radial bins (so that biasing could be applied by ADVANTG).

MCNP5 track length (F4) tallies were created in the tally cells. These were numbered 104, 114, 124, 134 for neutrons and 204, 214, 224, 234 for photons. The tally assignments are shown in Figure 1. These tallies were normalized by a 1.0714×10^{17} n/s source strength (the first pulse in the scenario described in the benchmark) and modified by ANSI/ANS-6.1.1-1977 [14] flux-to-dose-rate conversion factors. Two additional tallies were also defined for the purpose of serving as a CADIS adjoint source. These tallies (144, neutron; 244, photon) were MCNP5 FMESH4 tallies, with a single mesh volume element forming a bounding box around the outermost tally volume (the tally 134/234 volume).

ADVANTG was then used to generate weight windows and source bias cards using the CADIS

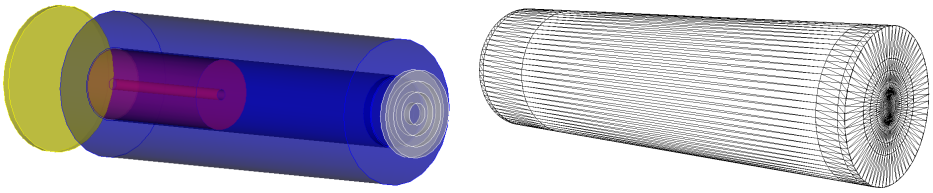

(a) CAD geometry (b) Faceted geometry

Figure 2. CAD and faceted versions of the ITER SDDR benchmark geometry. Note that the faceted version pictured was faceted with a faceting tolerance of 0.01 for the purpose of visualization only.

and FW-CADIS methods. This was done using both the MCNP5 CSG geometry and also the CAD model. For CADIS, the bounding box tallies 144 and 244 were used as an adjoint source. For FW-CADIS, the eight tallies in Figure 1 were used. ADVANTG was run using a P_3 Legendre order, a quadruple range quadrature set of order 16, and a convergence tolerance of 1×10^{-5} for within-group iterations. ENDF/B-VII.0 multigroup transport cross sections were used, with 27 neutron groups and 19 photon groups [15]. A 58×58×72 Cartesian mesh was used.

Simultaneous neutron and photon transport was then performed using DAG-MCNP5. This was first done in analog with 1×10^{10} histories on 160 cores using MPI on a Linux cluster. Transport was then done using the weight windows and source bias cards produced by the CADIS and FW-CADIS methods using the DAG-MCNP5 geometry. This was done with 80 simulations (each with a different random number seed) for both CADIS and FW-CADIS, each with 1000 minutes of processor time. The results from these simulations were then combined using appropriate statistical methods.

4 RESULTS AND DISCUSSION

4.1 ADVANTG Output from DAG-MCNP5 Geometry and Standard MCNP5 Geometry

Ray tracing different geometry representations can result in minor systematic differences in mesh material fractions, as shown in Figure 3 (a). This figure shows the average discrepancy in volume fractions of all materials within each mesh volume element as defined in Equation 4. Figure 3 (a) shows two continuous lines of mesh volume elements that lie on a cylindrical cell boundary with d_{mat} values in the 0.05 – 0.10 range. These lines span the geometry axially (in the z direction), which suggests that this discrepancy results from rays fired in the z direction.

This is confirmed in Figure 3 (b), which shows d_{mat} when rays are only fired in the x and y directions using the same number of rays per face (10, the default value) and same random number seed. An explanation for how firing rays in the z direction yielded these discrepancies is shown in Figure 4. In addition to only firing rays in the x and y directions, another experiment was done by increasing the number of rays per face to 100 and firing rays in the x, y, and z directions. This also resulted in the elimination of the larger discrepancies in d_{mat}.

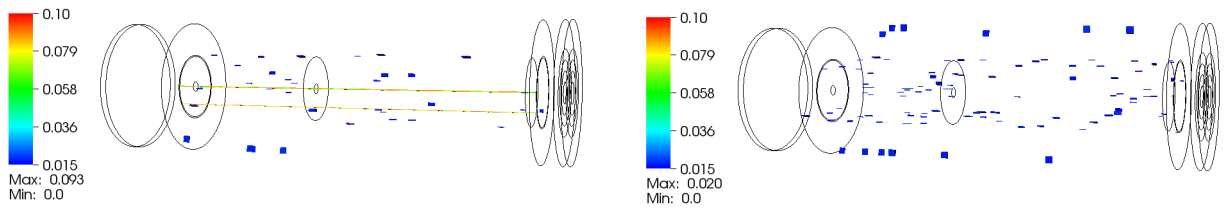

(a) Rays fired in x, y, and z directions.

(b) Rays fired in x and y directions only.

Figure 3. Plots of d_{mat} (defined in Equation 4). Values of d_{mat} below 0.015 are transparent.

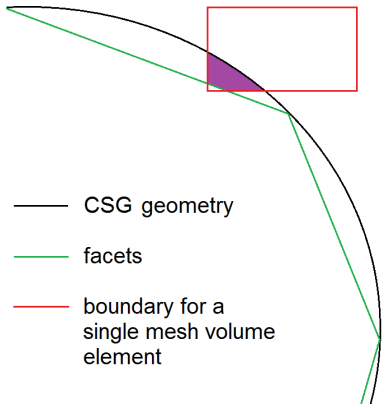

Figure 4. The source of discrepancies as seen in Figure 3. If a mesh row represented by the red rectangle is being ray traced, rays that pass through the violet region are tallied as outside the cylinder in the faceted case and inside the cylinder in the CSG case, resulting in discrepancies in material fractions.

ADVANTG combines materials with similar compositions in order to save memory. The tolerance for combining materials is governed by a user-specified parameter: `mcnp_mix_tolerance`. Ray tracing differences between faceted and CSG geometries in one area of a geometry may affect other areas because they may determine whether a new mixture is created or if a preexisting mixture is used. This results in minor differences in material composition, on the order of the `mcnp_mix_tolerance`, which has a default value of 0.01.

The discrepancy in mesh-based materials yields different transport results and ultimately different variance reduction parameters. Figure 5 shows the maximum relative discrepancy in weight window lower bounds across all energy groups as defined in Equation 5. Figure 5 illustrates that discrepancies in material volume fractions have a non-local impact on the calculated weight windows as a result of the transport calculation. These effects are more significant with the FW-CADIS method, which involves two transport steps. It should be emphasized that Figure 5 shows maximum discrepancies across all energy groups. Since the discrepancies of material fractions happened to occur on the boundary of a void region where streaming dominates, high-energy groups were more strongly affected. For low-energy particles, the average discrepancy is considerably smaller.

The minor discrepancies in weight window lower bound are not expected to have any apprecia-

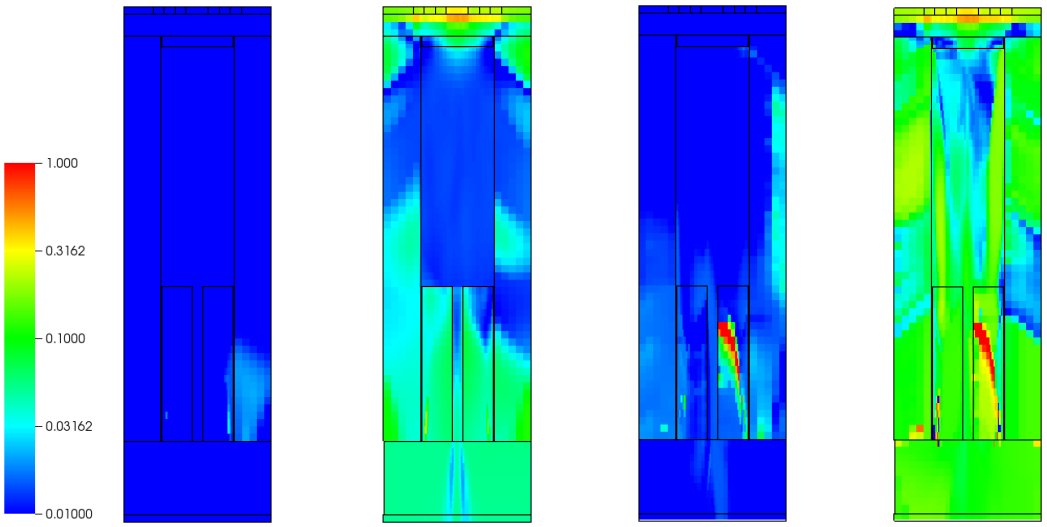

(a) Neutron CADIS (b) Neutron FW-CADIS (c) Photon CADIS (d) Photon FW-CADIS

Figure 5. Maximum relative discrepancy in weight window lower bounds across all energy groups (relative discrepancy is defined in Equation 5) for the CADIS and FW-CADIS methods for both neutrons and photons. The color scale has an upper bound of 1, though this is not the upper bound for the relative discrepancies, which in rare cases is as high as 10^2.

ble effect on performance. It is not uncommon in FES problems for weight window lower bounds to vary by 10–15 orders of magnitude within a problem, so these differences are insignificant. In order to minimize discrepancies, the number of rays fired can be increased or a smaller faceting tolerance can be used during geometry preprocessing.

4.1.1 Deterministic flux estimates

The adjoint neutron fluxes from the CADIS and FW-CADIS methods with the DAG-MCNP5 geometry are shown in Figure 6. These plots show the fluxes for energy group 10 (7.4082 – 8.1873 MeV), which was chosen arbitrarily. Note that the magnitudes of adjoint fluxes are different, as expected, because the response is different for the two methods. The color scales for the plots in Figure 6 span from the global minimum value to the global maximum value for each plot, so the shape of plots can be compared directly. Group 13 (600 – 800 keV) adjoint photon fluxes are shown in Figure 7. The deterministic forward fluxes used to calculate the FW-CADIS response are shown in Figure 8.

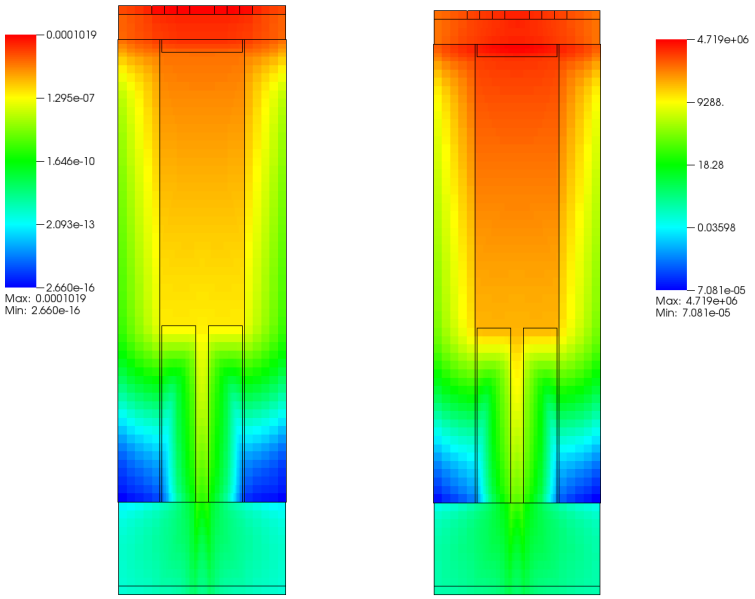

(a) CADIS neutron adjoint flux group 10

(b) FW-CADIS neutron adjoint flux group 10

Figure 6. Adjoint neutron fluxes from DAG-MCNP5 geometry.

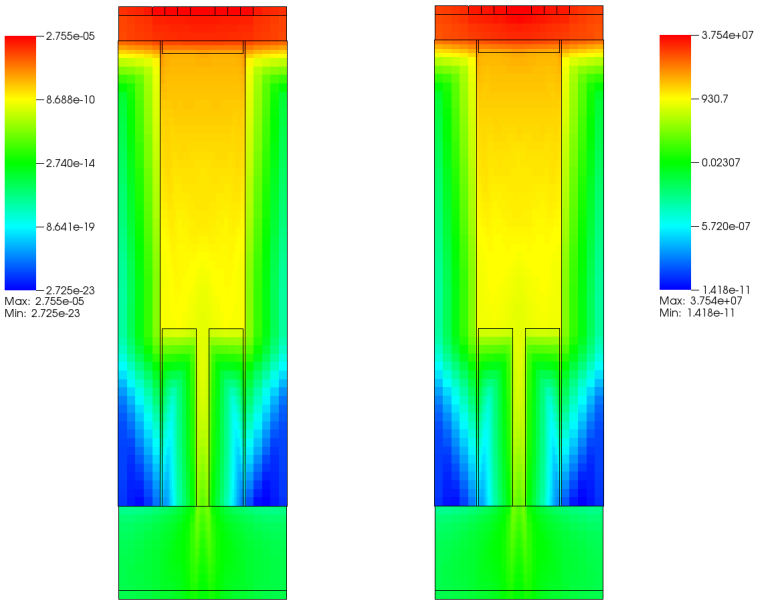

(a) CADIS photon adjoint flux group 13

(b) FW-CADIS photon adjoint flux group 13

Figure 7. Adjoint photon fluxes from DAG-MCNP5 geometry.

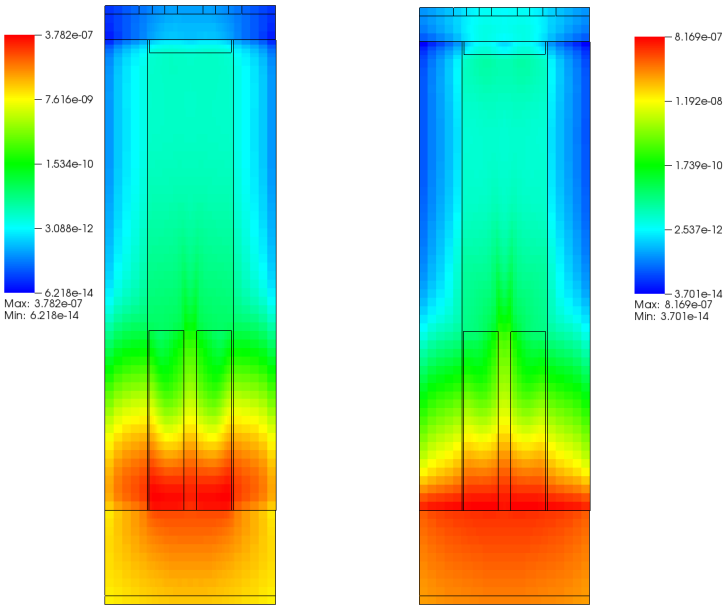

(a) Forward neutron group 10 (b) Forward photon group 13

Figure 8. Deterministic forward fluxes on DAG-MCNP5 geometry for FW-CADIS method.

4.2 MC Transport Results with Variance Reduction

Tables I, II, III show the results of particle transport in analog, with CADIS variance reduction, and with FW-CADIS variance reduction, respectively, all produced from CAD geometry. In the case of the three innermost neutron tallies (104, 114, 124) a modest speedup is observed. The dose rate in these tallies is dominated by particles that stream through the central streaming path, which is apparent from the adjoint flux distribution in Figure 6. When using weight windows, weight is only checked at surface crossings and collisions, so streaming particles are not affected. Therefore, most of the speedup comes from the radial biasing of the source, which preferentially samples particles born in the region aligned with the central streaming path. This was confirmed by MC simulations performed with CADIS and FW-CADIS weight windows, but without source biasing; virtually no improvement in FOM was achieved.

The outermost neutron tally (134) decreased in FOM with both CADIS and FW-CADIS weight windows. This is due to the inability to accurately represent the cylindrical geometry on Cartesian mesh, which has been shown to impact the efficiency of simulations using CADIS/FW-CADIS [16]. Tally 134 is aligned with the 2 cm annular streaming path between the steel cylinder and the steel plate. The contribution of neutrons from this streaming annulus is significant—as suggested by the fact that the tally 134 dose rate is much higher than the tally 124 dose rate. However, the Cartesian mesh used for the deterministic flux estimates does not fully resolve this streaming annulus, as seen in Figure 9. This results in inaccurate deterministic flux estimates in this region, resulting in weight windows that do not improve the efficiency of the simulation. The large portion of particles that score via streaming will also not play the weight window game, which means their weights will be much higher than particles that score after diffusing through the steel plate. In other words, the

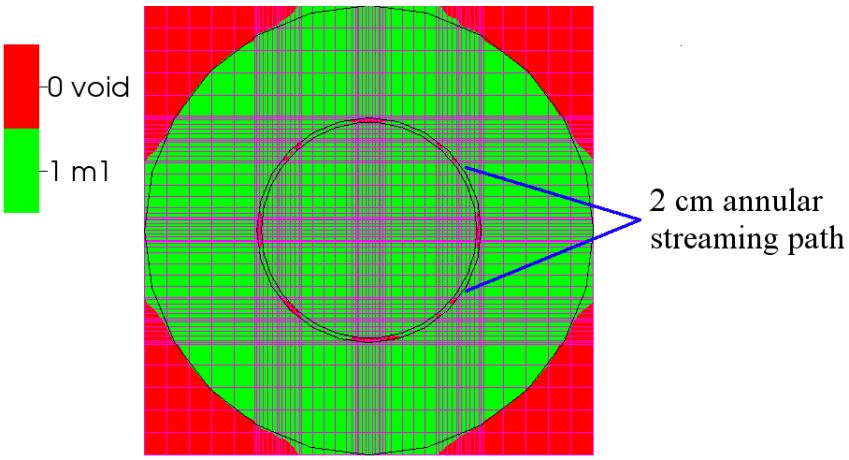

Figure 9. Slice through the material map halfway through the steel plate. This slice shows that the annular region is not well represented as void on the Cartesian mesh.

distribution of tally score weights will be bimodal, resulting in large variances.

For photon transport, speedup is achieved for all tallies using both CADIS and FW-CADIS weight windows. The photon contribution is likely dominated by photon production in the steel plate (as supported by the adjoint flux distribution in Figure 7). This steel plate provides a site for photon interaction that yields particle splitting, resulting in more tally scores, thereby accelerating the convergence of the tallies.

The results in Tables II and III show good statistical agreement to analog results. Out of the 16 dose rates presented, 10 (62.5%) were within $1S_{\text{diff}}$ of the analog dose rate, 14 (87.5%) were within $2S_{\text{diff}}$, and 16 (100%) were within $3S_{\text{diff}}$, which is close to the expected Gaussian distribution (68.3%, 95.4%, 99.7%).

Table I. Analog MC results

tally	dose rate (rem/h)	relative error
104	4706	0.0232
114	3082	0.0161
124	1628	0.0165
134	3078	0.0114
204	19.48	0.0582
214	17.84	0.0368
224	12.09	0.0351
234	8.24	0.0353

Table II. Results using CADIS

tally	dose rate (rem/h)	relative error	number of S_{diff} from analog result	speedup
104	4647	0.0105	0.49	7.10
114	3076	0.0081	0.11	5.58
124	1621	0.0105	0.22	3.53
134	3153	0.0434	0.53	0.10
204	19.97	0.0090	0.43	59.55
214	16.35	0.0090	2.21	23.74
224	11.42	0.0100	1.52	17.46
234	7.57	0.0140	2.16	9.05

Table III. Results using FW-CADIS

tally	dose rate (rem/h)	relative error	number of S_{diff} from analog result	speedup
104	4790	0.0142	0.65	3.85
114	3131	0.0107	0.82	3.19
124	1639	0.0128	0.32	2.35
134	3133	0.0579	0.30	0.05
204	20.45	0.0115	0.84	36.19
214	16.74	0.0114	1.61	14.61
224	11.64	0.0127	1.00	10.83
234	7.82	0.0250	1.20	2.82

5 CONCLUSIONS

The ability to generate variance reduction parameters directly from CAD geometries represents a significant new capability that will be applicable to ITER analysis. The weight windows produced using DAG-MCNP5 geometry agree within a reasonable degree with the weight windows generated from MCNP5 CSG geometry. Discrepancies are the result of fundamental differences in the geometry representations and are not expected to affect performance. These discrepancies can be eliminated by increasing the number of rays fired for discretization.

Both the CADIS and FW-CADIS methods produced weight windows and source biasing parameters that resulted in significant speedup for all but the outermost neutron tally. This is due to the fact that the overlaid mesh did not provide a high-fidelity representation of the annular streaming path and that particles contribute to this tally by both streaming and scattering. A finer deterministic mesh could be used to more completely characterize this behavior.

6 ACKNOWLEDGMENTS

This work was funded in part by the ORNL Nuclear Engineering Science Laboratory Synthesis (NESLS) program. Funding was also provided by the U.S. Nuclear Regulatory Commission

Fellowship program. The authors wish to thank Douglas Peplow and Seth Johnson for facilitating this work and providing additional guidance. Thanks go to Paul Romano and Tom Sutton, who provided the LaTeX template for this paper.

7 REFERENCES

[1] J. C. Wagner and A. Haghighat, "Automated Variance Reduction of Monte Carlo Shielding Calculations Using Discrete Ordinates Adjoint Function," *Nuclear Science and Engineering*, **128** (1998).

[2] J. C. Wagner, D. E. Peplow, and S. W. Mosher, "FW-CADIS Method for Global and Regional Variance Reduction of Monte Carlo Radiation Transport Calculations," *Nuclear Science and Engineering*, **176** (2014).

[3] A. M. Ibrahim et al., "ITER Neutronics Modeling Using Hybrid Monte Carlo/SN and CAD-based Monte Carlo Methods," *Nuclear Technology*, **175**, pp. 251–258 (2011).

[4] S. W. Mosher et al., *ADVANTG–An Automated Variance Reduction Parameter Generator*, ORNL/TM-2013/416, Oak Ridge National Laboratory (2013).

[5] X-5 Monte Carlo Team, "MCNP – A General Monte Carlo N-Particle Transport Code, Version 5," (2004).

[6] 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**, pp. 171–200 (2010).

[7] U. Fischer et al., "Use of CAD Generated Geometry Data in Monte Carlo Transport Calculations for ITER," *Fusion Science and Technology*, **56**, *2*, pp. 702–709 (2009).

[8] T. J. Tautges, P. P. H. Wilson, J. Kraftcheck, B. F. Smith, and D. L. Henderson, "Acceleration Techniques for Direct Use of CAD-Based Geometries in Monte Carlo Radiation Transport," *International Conference on Mathematics, Computational Methods & Reactor Physics*, Saratoga Springs, NY, (2009).

[9] T. J. Tautges, R. Meyers, K. Merkley, C. Stimpson, and C. Ernst, *MOAB: A Mesh-Oriented Database*, SAND2004-1592, Sandia National Laboratories (2004).

[10] M. E. Sawan et al., "Application of CAD-Neutronics Coupling to Geometrically Complex Fusion Systems," *Proc. 23rd Symp. Fusion Engineering*, San Diego, California, (2009).

[11] T. J. Tautges, "The Common Geometry Module (CGM): a Generic, Extensible Geometry Interface," *Engineering with Computers*, **17**, pp. 299–314 (2001).

[12] M. Loughlin, "Conclusion of the Shutdown Dose Rate Benchmark Study," 6th ITER Neutronics Meeting, Hefei, China, (2011).

[13] G. D. Sjaardema et al., *CUBIT mesh generation environment Volume 1: Users manual*, Sandia National Laboratories (1994).

[14] "Neutron and Gamma-Ray Flux-to-Dose-Rate Factors," American Nuclear Society, ANSI Standard ANSI/ANS-6.1.1-1977, 1977.

[15] D. Wiarda, M. E. Dunn, D. E. Peplow, T. M. Miller, and H. Akkurt, *Development and Testing of ENDF/B-VI.8 and ENDF/B-VII.0 Coupled Neutron-Gamma Libraries for SCALE 6*, ORNL/TM-2008/047, NUREG/CR-6990, Oak Ridge National Laboratory (2008).

[16] A. Ibrahim, *Automatic Mesh Adaptivity for Hybrid Monte Carlo/Deterministic Neutronics Modeling of Difficult Shielding Problems*, PhD nuclear engineering and engineering physics, University of Wisconsin-Madison, Madison, WI, United States, 2012.