INL/CON-23-72565-Revision-0

3D full core BEAVRS using OpenMOC with transport equivalence

August 2023

Guillaume Louis Giudicelli, Kord Smith, Benoit Forget

Changing the World's Energy Future

INL is a U.S. Department of Energy National Laboratory operated by Battelle Energy Alliance, LLC

INL/CON-23-72565-Revision-0

3D full core BEAVRS using OpenMOC with transport equivalence

Guillaume Louis Giudicelli, Kord Smith, Benoit Forget

August 2023

Idaho National Laboratory
Idaho Falls, Idaho 83415

http://www.inl.gov

Prepared for the
U.S. Department of Energy
Under DOE Idaho Operations Office
Contract DE-NE0008578

M&C 2023 - The International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering · Niagara Falls, Ontario, Canada · August 13 – 17, 2023

3D full-core BEAVRS using OpenMOC with transport equivalence

G. Giudicelli[1,2], B. Forget[2], K. Smith[2]
[1]Computational Frameworks, Idaho National Laboratory
1955 Fremont Ave., Idaho Falls, ID 83415

[2]Department of Nuclear Science and Engineering, Massachusetts Institute of Technology
77 Massachusetts Avenue, Cambridge, MA 02149

guillaume.giudicelli@inl.gov

ABSTRACT

Using an optimized implementation of the 3D method of characteristics (MOC) for neutron transport, along with a novel equivalence method for transport calculations that was designed to recover self-shielding errors from neglecting the angular dependence of resonant group absorption, a 3D full-core light-water reactor hybrid stochastic-deterministic eigenvalue calculation was achieved. This paper presents the optimizations developed and compares the transport solutions obtained. For the statepoint, a runtime of around 6,000 CPU hours was achieved—improving on previous works by an order of magnitude—with less than 1% error on pin fission to ^{238}U capture ratios and a few dozen pcms on the eigenvalue.

KEYWORDS: 3D MOC, equivalence, full core, high fidelity

1. INTRODUCTION

Simulation tools for the current nuclear fleet are able to achieve a few percent accuracy, thanks to decades of research by both the academic and the industrial world. Carefully crafted approximations and well-validated models ensure fast, reliable predictions of power distributions, safety parameters, isotopic depletion, and transient behavior. These predictions aid in designing improved fuel loading patterns, demonstrating safety requirements, and training operators on realistic models, and are key to nuclear power plant safety and economics. Operational requirements necessitate very fast run times, and these highly efficient methods will likely remain in use over the coming decades. Yet despite all this, the need for high-fidelity methods still remains.

Evolutionary designs that are currently being built (e.g., the Westinghouse©AP1000) can challenge existing tools. They feature increased axial heterogeneities (e.g., partial length burnable poisons and enrichment zoning) to reduce power peaking in the reactor. These new degrees of freedom in assembly design allow for longer cycle lengths, improving the economics of these new plants. However, they can can prove challenging to the current simulation suite. Homogeneous models (e.g., nodal methods) represent spatial variations by using low-order polynomials, which are insufficient for modeling sharp gradients. High-fidelity methods can be used to validate novel developments in current industrial tools so as to adapt to these new heterogeneities.

Production simulation tools for future reactors have not yet been established and validated. Qualifying these new reactors and their corresponding simulation tools will likely require years of research, as well as numerous expensive experiments. High-fidelity methods are highly flexible

in terms of the types of systems that can be modeled, and do not require experimental data, since they need not be tuned to a particular reactor. The ability to simulate these future reactors can be exploited to analyze the safety cases of new designs, and to partially replace costly experiments with numerical simulations.

However, the generality afforded by using few approximations comes at great computational cost. The gold standard in reactor physics is the Monte Carlo simulation approach, which models individual neutron random walks inside the nuclear reactor and computes the statistical averages of the grandeurs of interest. This approach is challenged in large reactors, with dominance ratio and auto-correlation effects further slowing down a convergence rate that already only follows already at best a square root law. The hybrid stochastic-deterministic scheme proposed by Boyd [1] uses Monte Carlo methods to generate group cross sections at the full-core scale. The cross-section tallies are clustered using machine-learning algorithms to increase their convergence speed. Generating cross sections at the full-core scale naturally takes into account neighborhood effects such as the presence of guide tubes near a fuel pin. Using Monte Carlo methods eliminates the need for approximate resonance models. The cross sections are then transferred to a high-fidelity deterministic solver based on the method of characteristics (MOC). This solver is able to use the same heterogeneous geometry definition as the Monte Carlo code.

The above process proved highly successful for producing high-quality power distributions, especially after the findings of Liu [2] improving the scattering model utilized in the deterministic solver. However, during the condensation process operated by the Monte Carlo tally, the angular dependence of the group cross sections is not captured. This approximation results in a well-known resonance self-shielding error worth up to -300 pcm on the reactivity and 1% on the ^{238}U absorption rates with the group structure chosen.

New highly scalable 3D MOC capabilities were developed by Gunow [3], leading to the first instance of 3D full-core light-water reactor deterministic heterogeneous transport simulations. These enabled studies to be conducted on the self-shielding error in systems as large as the 3D full core, but remained impractical for all but some of the largest computing clusters in the world. But through MOC optimization for CPU-based architectures, we achieved high utilization of modern industrial computing clusters in order to demonstrate 3D MOC's potential as a high-fidelity reactor physics method that incurs reasonable computing costs—down an order of magnitude from previously published.

The present paper first describes the optimization of OpenMOC's 3D MOC solver. The code was optimized for common features of CPU-based computing clusters, such as vectorized computation capabilities and non-uniform memory architectures. An improved communication protocol was also crafted to halve the memory requirements of the short characteristics MOC solver yet still maintain scalability. Various other optimizations are briefly described herein, and the reduction in computational requirements is provided for a variety of test cases as well as several machines, when applicable. Finally, we present the transport results with a near-pellet-level resolution in a 3D full core.

2. OPTIMIZATION OF 3D MOC

Optimization of OpenMOC was a multi-faceted effort that extended far beyond the material presented herein. The information covered in the present article was selected for being of greater interest to the community, and full listings of the optimization tasks performed can be found in the relevant theses [3,4].

2.1. Single instruction, multiple data low-level parallelism

Single instruction, multiple data (SIMD) parallelism is the application of identical operations to different pieces of data. SIMD is a type of instruction-level parallelism, but in this particular case, all the instructions are the same. The amount of speedup expected from using vector instructions depends on the length of the vectors used, since it is the number of instructions being performed simultaneously. SIMD vectorization is usually performed along energy groups [5,6] in multigroup transport.

As demonstrated via profiling [4], most of the computational time in OpenMOC was spent in the main propagation/tally routine. The main tally routine computes the attenuation of angular fluxes in every energy group along a single segment, then tallies that term as a contribution to the source region's scalar flux. This tally routine is composed of four steps, each conducted for every energy group: computing the sources, computing the exponentials, computing the attenuation, and tallying the scalar flux contribution. Algorithm 1 shows how these steps can be accomplished in both a naive and an easily vectorized manner. Allocating arrays on the stack is cheap and can be done in the tally loop. Decomposing each part of the work helps the compiler vectorize. The last three steps can be performed together if the data access patterns are sufficiently simple (as with a stride of 1) and if the number of groups is known at the compile time. (The latest versions of GCC do not require this second condition.) Another critical step in improving SIMD performance was to approximate the linear source attenuation exponentials [7] via rational fractions [8,9] computed using Horner's method [10].

Algorithm 1 Main tally routine in a format that is hard for the compiler to vectorize (left), and one that is easy for the compiler to vectorize (right)

```
for each energy group do
    Compute source
    Compute exponential term
    Compute attenuation of track angular flux
    Tally track angular flux variation to scalar flux
    Attenuate track angular flux
end for
```

```
Allocate buffers for source, exponential,
and attenuation terms
for each energy group do
    Compute source, store in buffer
end for
for each energy group do
    Compute exponential term, store in buffer
end for
for each energy group do
    Compute attenuation of track angular flux,
store in buffer
    Attenuate track angular flux
    Tally track angular flux variation to scalar flux
end for
```

2.2. Single storage of angular fluxes

Another bottleneck in achieving high performance on the computing clusters available at Idaho National Laboratory (INL) at the time of this research was the large memory requirements for storing boundary angular fluxes, as quantified in [4]. Numerous compute nodes were necessary in order to have enough memory to store them.

Part of why those boundary angular fluxes take up so much storage is that, for the sake of convenience, they are stored twice. The track starting flux used at the beginning of the sweep is stored

in a different location than the track end flux. The rationale for this is very clear when considering reflective or periodic boundary conditions. Once the solver is finished with a track, it can store the track fluxes in the end array rather than the starting flux array, thus avoiding a thread safety issue if another thread is currently working on the next track and accessing its starting flux.

With spatial domain decomposition (i.e., the approach often chosen for 3D MOC distributed memory parallelism [3,11,5,12]), the boundary angular fluxes must be communicated from where they end on one domain to where they start on the next domain along the track. The "end" and the "start" angular fluxes are generally stored in separate arrays.

The communication phase progresses as follows. Each domain moves track fluxes from the "start" array into a first "send buffer" of track fluxes, allocated for each of their neighbor domains. The send buffers arrive at the receiving domains as "receive buffers." Buffers are much smaller than the total number of track fluxes and creating them incurs negligible memory costs. The receive buffers are then unpacked into the start array. For every track flux to unpack, the process checks whether the corresponding flux has already been communicated. If yes, the track flux is overwritten with the flux to unpack. If no, the track flux in the receiving array is saved in the send buffer (a process referred to as pre-filling), then the track flux from the receive buffer can safely be unpacked into the start (receiving) array without overwriting data that should be sent to another domain.

When packing the next round of send buffers, the starting index into the send buffer represented the end of the pre-filling. The process loops through a list of tracks ending in each neighbor. The process checks whether each track flux has not already been put into a send buffer by pre-filling. This is done to avoid communicating it twice and to facilitate the determination of the end of the track communication phase.

Modifications of the algorithm are restricted to cases involving vacuum boundary conditions. This simplifies the treatment of domains on the border of the global system, since the boundary fluxes are always 0 there. During the transport sweep, outside of the communication phase, the track fluxes remain in the start array. Thus, there is no conflict between threads, since each thread deals with different track fluxes found at different locations in the array.

2.3. Using mixed-precision floating-point arithmetic

MOC has long made use of single precision for various quantities [13,14] such as boundary angular fluxes [15]. These constitute the majority of the memory cost for light-water reactor MOC applications. Double precision was reserved for accumulated variables such as the fission sums defining k_{eff}. The round-off error when summing a large number of floats could become an issue when trying to converge those quantities to 1 pcm.

For scalar fluxes, sources, and most other variables, the numerical precision is set at compile time. The FP_PRECISION flag can be set to either single or double—depending on the precision wanted in the MOC solver—and the CMFD_PRECISION flag governs the numerical precision used in the CMFD solver.

Numerical convergence issues were observed for all 3D linear source cases when using single precision in the MOC solver. By using a deprecated SF_PRECISION flag, scalar fluxes could be set independently to single precision while keeping other variables set to double precision. This eliminated the issue. By looking at the accumulated sums of centroids and linear source moments over a symmetric system for which they should analytically have summed to 0, the issue was traced back to ray-tracing precision. Segment lengths were stored in FP_PRECISION. This was

changed to double precision. Ray tracing is now performed consistently in double precision, while sweeping can now be performed in single precision without causing any convergence issues.

3. RESULTS

The system of interest is the 3D cycle 1 fresh core, as defined in the Benchmark for Evaluation and Validation of Reactor Simulations (BEAVRS) specification [16]. We first present the computational expense of a 3D full-core MOC simulation after implementing the optimizations described in Section 2, then examine the accuracy of the full-scale simulation by comparing it against a reference Monte Carlo simulation. The Monte Carlo reference solution was computed using OpenMC [17], with 1,000 inactive and 4,200 active batches of 100 million particles each.

Based on the 2D and 3D core convergence studies [3,4], a discretization that should achieve 1% convergence on the pellet reaction rates and an error of less than 10 pcm on the eigenvalue would resemble the one shown in Table 1. By tolerating a higher error on the eigenvalue, the more commonly used 64 azimuthal angles may be preferred. The pin reaction rates are much more converged than the pellet reaction rates. If the objective was to obtain accurate pin-wise reaction rates instead of pellet-wise rates, this would enable a significant reduction in the computational cost. All studies were performed using the CASMO 70-group structure [14]. Neutron scattering anisotropy was modeled using P_0 transport-corrected cross sections using the flux-limited in-scatter approximation [18] for the fuel, gap, clad, and burnable poisons, along with the migration area method for the water [19].

Table 1: Minimal discretization parameters for ensuring sufficient convergence of the K-effective and pellet reaction rates, using a 70-group energy structure for an unrodded 3D core case.

Source	Linear
Azimuthal angles (-)	128 (recommended 64)
Track spacing (cm)	0.1
Polar angles (-)	6
Axial spacing (cm)	1.0
Fuel rings	1 (no Gd)
Fuel sectors	2
Water rings	1
Water sectors	8
Reflector & baffle source region discretization (per pin cell mesh)	1
Axial reflector height -top	15 cm
-bottom	15 cm

3.1. Computational expense

Table 2 summarizes the current computational expense for simulating a 3D core on 50 nodes of the previous-generation INL computing clusters. At 6,000 CPU hours, the total computational cost is

reasonable for modern industrial computing clusters. All major parts of the solver contribute significantly to the computing time, indicating that further optimization efforts will need to consider OpenMOC as a whole. The discretization followed the recommendations of Table 1.

Table 2: Run-time decomposition of a 3D full-core run on 50 nodes of Lemhi.

	3D core	
	CPU hr	%
Initial track generation	1355	12
Transport sweep (excl. comm.)	3440	30
Track fluxes MPI comm.	2767	24
Idle time	1238	11
CMFD run time	1786	16
Time per integration (sweep, ns)	5.66	
Number of transport sweeps	12	
Total CPU time (CPU hr)	6277	
Memory used per node (GB)	46	

3.2. Verification

We compared the fission distribution, ^{238}U capture distribution, and fission-to-^{238}U-absorption ratios for the 3D core at the pin cell, axial, and pellet levels. The Monte Carlo reference seems sufficiently accurate at the pin cell and axial levels, with errors of only a few tenths of a percent. However, as shown in Fig. 1, even though these 4,200 active cycles took nearly 100,000 CPU hours to complete, the maximum pellet fission rate fractional error remained close to 10%.

Table 3 summarizes the accuracy of 3D full-core simulations with various multi-scale equivalence schemes using the local neighbor symmetry (LNS) [1] group cross-section discretizations. The base case is without equivalence. Equivalence factors are generated from pin-cell calculations for the Generalized Equivalence Theory (GET) ([20] extended to transport [21]) and Super Homogenization (SPH) [22] cases. Reference [23] shows that fuel pin neighborhood effects only have a minor impact on equivalence parameters. The number of fuel sectors which raised to four for consistency with the pin-cell calculations; the full test case specification is available in [4].

We first observed that the reactivity biases were very close to what was reported for the 2D core (i.e., near -200 pcm). Equivalence recovers most of this bias, which is also consistent with the 2D core results [4]. On the other hand, the pin reaction rates are not as accurate as for the 2D core. The maximum fractional error near the reflector is close to 1.5% for the homogeneous case, and close to 2% for the LNS case—as opposed to 1.3% in 2D [4]. The RMS of the fractional error also rose from 0.3% [4] to 0.45% and 0.66%, partly due to insufficient convergence of the Monte Carlo simulation.

Examining reaction rates via axial slices of the core shows that the uppermost and lowermost planes display a large fractional error ($\approx$5%) for ^{238}U absorption rates, and a much smaller error for fission rates. This error level is consistent with that observed when studying the 3D pin cell system, though the fission rate errors are actually much lower, which was unexpected.

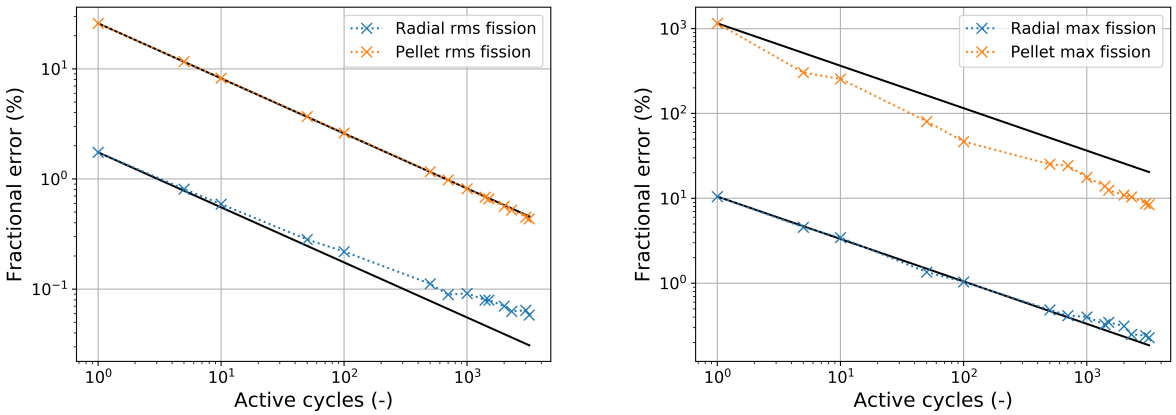

(a) RMS of fission rate fractional errors

(b) Maximum fractional error on fission rates

Figure 1: Convergence of the Monte Carlo reference fission rate distribution with the number of active batches

Finally, the pellet error rates are very high, mostly because they are not converged in the Monte Carlo simulation. The RMS error is close to 2%, whereas the maximum error is close to 30%. Equivalence does little to improve the reaction rates. Very high fractional errors are expected, as axial extreme regions are the most difficult to predict for MOC (given the proximity of the reflector) and have terrible statistics in Monte Carlo. The absolute pellet errors are closer to 5%, but could be further reduced by running more neutron generations.

The fuel pin reaction rate fractional errors are shown in Figures 2 and 3 for fission and ^{238}U absorption, and in Figure 4 for the ^{238}U-absorption-to-fission ratio. Results are presented for a quarter core, as the MOC solution is symmetric to numerical precision and the Monte Carlo solution is octant-symmetrized then re-expanded onto the quarter core.

The pin reaction rates show the same error pattern as in the 2D core cases, with a clear in-out tilt. On top of this in-out tilt, there is a 1% for the ^{238}U absorption rates, which is recovered by introducing equivalence. The error also peaks near the reflector, due to both the transport correction approximation and the inadequacy of recursive LNS (for the second pin row).

Taking the ^{238}U-absorption-to-fission ratio removes the in-out tilt, leaving only a checkerboard error pattern and the peaks near the reflector. This inter-assembly checkerboard error pattern was investigated in [4] and is likely caused by the transport correction approximation.

Figure 5 shows, for individual pellets on each radial and axial slice of the core, the ^{238}U-absorption-to-fission ratio fractional errors. The numerous spikes in error are due to statistical uncertainty and are shared across cases. This is especially evident near the radial, top, and bottom reflector, as the statistics on the pellet reaction rates there are worse. These end effects create a large maximum error that inflates the RMS of the error. Overall, the absolute error, which neglects these end effects, is smaller than the fractional error, and as Table 3 shows, is closer to 5%.

The radial tilt is clearly seen in the radial distribution of the per-pin RMS of the pellet fission and absorption reaction rates. The error also peaks on the corner of the assemblies near the reflector.

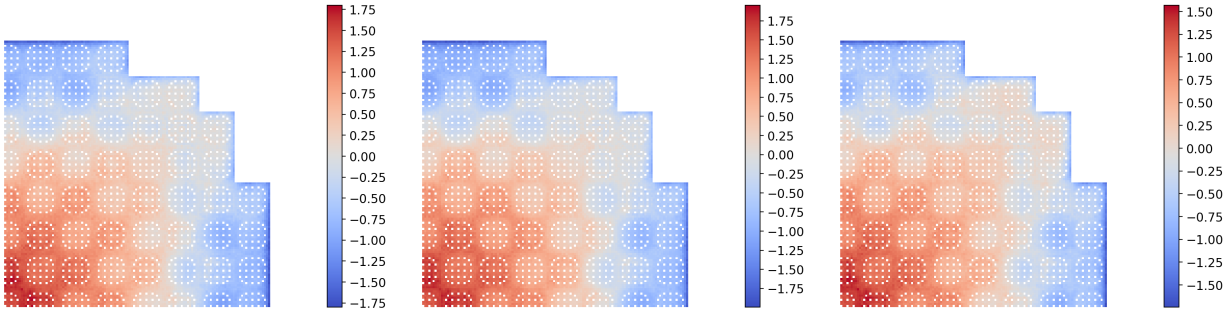

(a) Without equivalence **(b) GET** **(c) SPH**

Figure 2: Radial pin-wise fission rate fractional errors, both without equivalence and with equivalence factors generated from pin-cell calculations

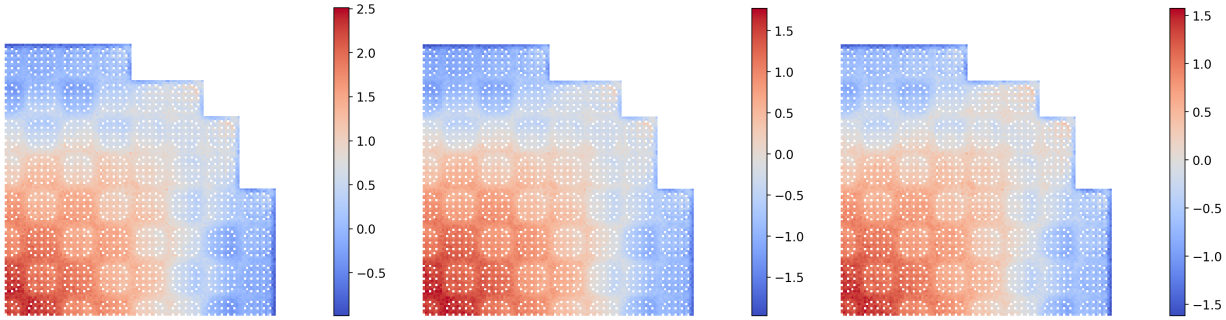

(a) Without equivalence **(b) GET** **(c) SPH**

Figure 3: Radial pin-wise ^{238}U absorption rates fractional errors, both without equivalence and with equivalence factors generated from pin-cell calculations

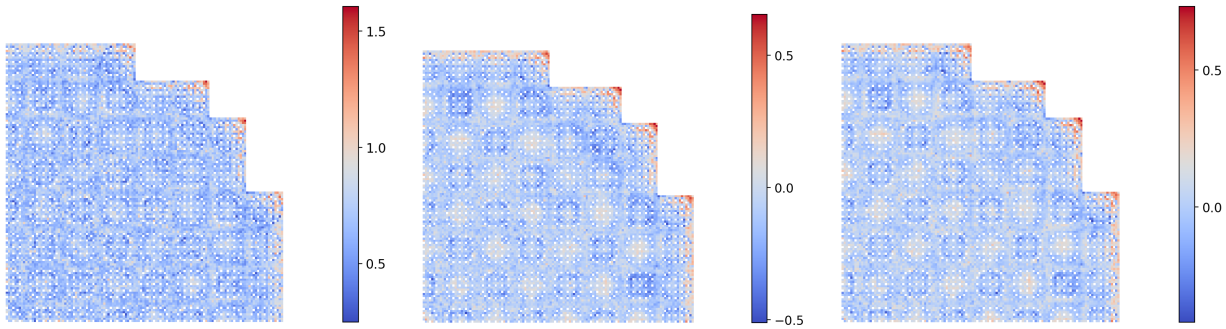

(a) Without equivalence **(b) GET** **(c) SPH**

Figure 4: Radial pin-wise ^{238}U-absorption-to-fission-ratio fractional errors, both without equivalence and with equivalence factors generated from pin-cell calculations

Table 3: Study of the 3D core, using equivalence factors generated via multi-scale equivalence procedures

Case	Cross-section discr	Equiv. method	$\Delta\rho$ (pcm)	Frac. error on pin: Fission rates RMS (%)	Fission rates max (%)	^{238}U absorption RMS (%)	^{238}U absorption max (%)	$\frac{^{238}\text{U abs.}}{\text{fission}}$ RMS (%)	$\frac{^{238}\text{U abs.}}{\text{fission}}$ max (%)
Unrodded	LNS	None	-197	0.66	1.79	1.02	2.5	0.7	1.43
-	-	GET*	-29	0.74	1.99	0.73	1.96	0.14	0.65
-	-	SPH	4	0.58	1.74	0.57	1.61	0.11	0.73
				Frac. error on axial slice: Fission rates RMS (%)	Fission rates max (%)	^{238}U absorption RMS (%)	^{238}U absorption max (%)	$\frac{^{238}\text{U abs.}}{\text{fission}}$ RMS (%)	$\frac{^{238}\text{U abs.}}{\text{fission}}$ max (%)
Unrodded	LNS	None	0.19	0.47	0.83	5.35	0.86	5.36	
-	-	GET	0.17	0.35	0.51	5.86	0.51	6.08	
-	-	SPH	0.17	0.36	0.53	6.13	0.51	5.9	
				Frac. error on radial pellet: Fission rates RMS (%)	Fission rates max (%)	^{238}U absorption RMS (%)	^{238}U absorption max (%)	$\frac{^{238}\text{U abs.}}{\text{fission}}$ RMS (%)	$\frac{^{238}\text{U abs.}}{\text{fission}}$ max (%)
Unrodded	LNS	None	0.82	8.11	1.41	22.91	1.25	20.26	
-	-	GET	0.88	8.12	1.23	23.62	1.05	20.94	
-	-	SPH	0.75	7.75	1.16	23.48	1.05	20.74	

Taking the ^{238}U-absorption-to-fission ratio removes the tilt but leaves the error near the reflector. As mentioned previously when covering the 2D core, the LNS discretization is not totally satisfactory there, so this may only be a cross-section discretization effect. The inter-assembly tilt can also be observed both with and without equivalence for the pellet fractional errors.

The axial distribution of pellet fractional errors shows spikes near every spacer. This is the same effect as shown for the 3D assembly case; the cross-section discretization should be increased near the spacers to account for reduced moderation. This effect is less prominent on the 3D core, as the statistical errors in the reference reaction rates are dominating.

Introducing equivalence, whether GET or SPH, reduces the RMS of the ^{238}U absorption rate fractional error, bringing it close to the fission rate fractional error. The agreement is not as close as for the 2D core case, as the pellet ^{238}U absorption rates present much larger errors near the axial reflectors, driving up the pin's axially integrated RMS. A similar effect would be observed axially for the RMS if the three distributions (Flux-Volume-Weighting, GET, SPH) were plotted together.

4. CONCLUSION

A 3D full-core high-fidelity deterministic simulation of a single statepoint of BEAVRS was achieved using the MOC for neutron transport. This calculation obtained converged and accurate results in less than 10,000 CPU hours on 50 nodes of the INL computing cluster Lemhi. It heavily leveraged three levels of parallelism in its implementation of the MOC [9], and is 1–2 orders of magnitude faster than previous methods [24,25]. It made use of a novel equivalence technique [21,26,4], Generalized Equivalence Theory for transport, designed to reduce the resonance self-shielding er-

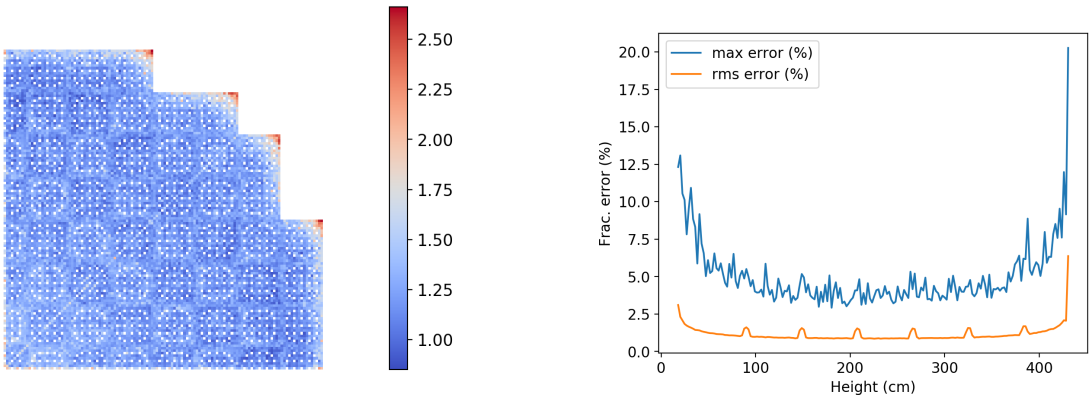

(a) No equivalence - radial

(b) No equivalence - axial

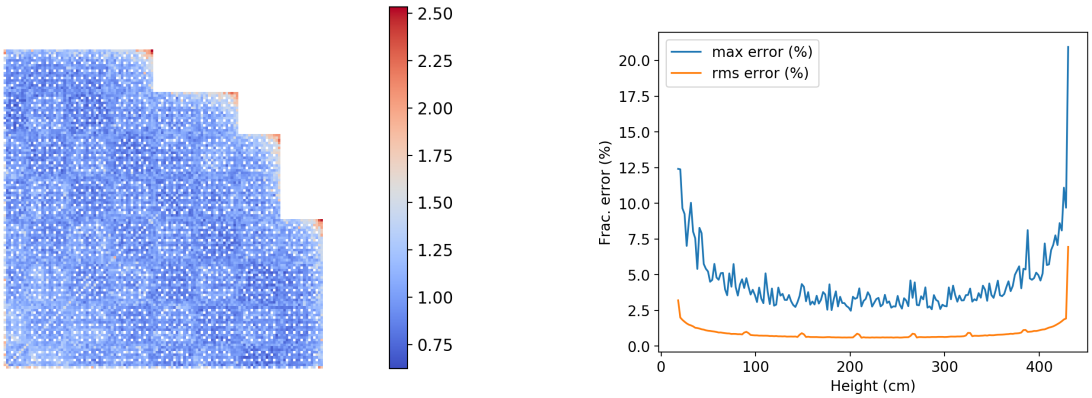

(c) GET - radial

(d) GET - axial

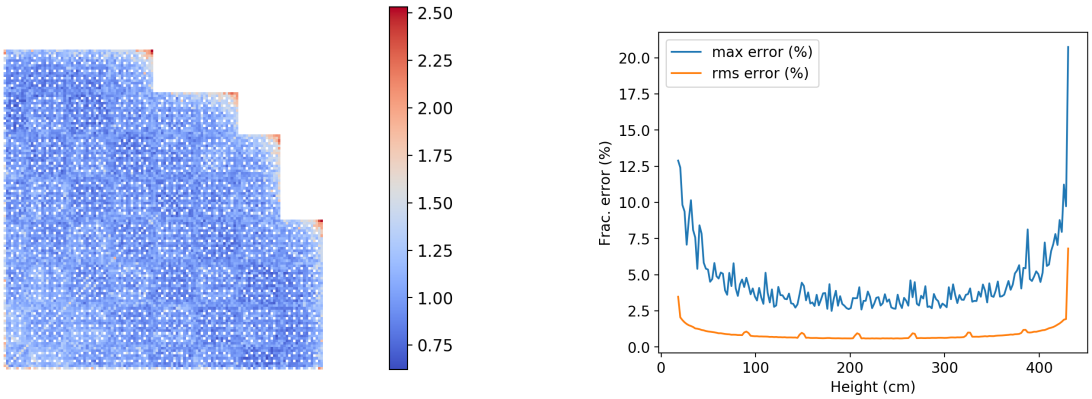

(e) SPH - radial

(f) SPH - axial

Figure 5: RMS and maximum pellet fractional errors on the pellet ^{238}U-absorption-to-fission ratio per radial (left) and axial (right) slices of the core

ror caused by neglecting the angular dependence of resonance group cross sections. It generally achieves 1% accuracy on fission-to-^{238}U-absorption ratios, and future efforts should aim to further improve this accuracy by concentrating on the scattering model. Several leads to further reduce computational costs can be found in [4], and include more balanced angular quadratures and better partitioning.

5. Acknowledgements

This research made use of the resources of the High Performance Computing Center at Idaho National Laboratory, which is supported by the Office of Nuclear Energy of the U.S. Department of Energy and the Nuclear Science User Facilities under contract no. DE-AC07-05ID14517 for all large-scale simulations presented. The research work presented in this paper was supported by U.S. Department of Energy Nuclear Energy University Program contract no. DE-NE0008578.

REFERENCES

[1] W. R. D. Boyd. *Reactor Agnostic Multi-Group Cross Section Generation for Fine-Mesh Deterministic Neutron Transport Simulations*. Ph.D. thesis, Massachusetts Institute of Technology (2017).

[2] Z. Liu. *Cumulative Migration Method for Computing Multi-Group Transport Cross Sections and Diffusion Coefficients with Monte Carlo Calculations*. Ph.D. thesis, Massachusetts Institute of Technology (2020).

[3] G. Gunow. *Full Core 3D Neutron Transport Simulation Using the Method of Characteristics with Linear Sources*. Ph.D. thesis, Massachusetts Institute of Technology (2018).

[4] G. Giudicelli. *A novel equivalence method for a high fidelity hybrid stochastic-deterministic reactor physics computations*. Ph.D. thesis, Massachusetts Institute of Technology (2020).

[5] J. R. Tramm. *Development of The Random Ray Method of Neutral Particle Transport for High Fidelity Nuclear Reactor Simulation*. Ph.D. thesis, Massachusetts Institute of Technology (2018).

[6] J. R. Tramm, G. Gunow, T. He, K. S. Smith, B. Forget, and A. R. Siegel. "A task-based parallelism and vectorized approach to 3D Method of Characteristics (MOC) reactor simulation for high performance computing architectures." *Computer Physics Communications*, **volume 202**, pp. 141 – 150 (2016). URL http://www.sciencedirect.com/science/article/pii/S0010465516000266.

[7] R. Ferrer and J. Rhodes. "A Linear Source Approximation Scheme for the Method of Characteristics." *Nuclear Science and Engineering*, **volume 77**, pp. 119–136 (1981).

[8] E. Y. Remez. "Sur la détermination des polynômes d'approximation de degré donné." *Comm Soc Math Kharkov*, **volume 10**, p. 196 (1934).

[9] G. Giudicelli, W. Wu, C. Josey, K. Smith, and B. Forget. "Adding a third level of parallelism to OpenMOC, an open source deterministic neutron transport solver." In *Mathematics & Computation 2019* (2019).

[10] W. G. Horner and D. Gilbert. "XXI. A new method of solving numerical equations of all orders, by continuous approximation." *Philosophical Transactions of the Royal Society of London*, **volume 109**, pp. 308–335 (1819). URL https://royalsocietypublishing.org/doi/abs/10.1098/rstl.1819.0023.

[11] D. R. Gaston. *Parallel, Asynchronous Ray-Tracing for Scalable Method of Characteristics Neutron Transport on Unstructured Mesh*. Ph.D. thesis, Massachusetts Institute of Technology (2019).

[12] B. Kochunas and T. Downar. "Performance Model Development for the 3-D Method of Characteristics." In *International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering*. Nashville, TN, USA (2015).

[13] S. Choi, Q. Shen, D. Jabaay, Y. Liu, B. Kochunas, and T. Downar. "MPACT Efficiency and Robustness Enhancements for Full-Core BWR Modeling NURAM-2021-002-00." Technical report, University of Michigan (2021).

[14] J. Rhodes and al. "CASMO-4 User's Manual Rev. 4, SSP-01/400 Rev 4." *Studsvik Scandpower* (2004).

[15] W. R. D. Boyd III. *Massively Parallel Algorithms for Method of Characteristics Neutral Particle Transport on Shared Memory Computer Architectures*. Master's thesis, Massachusetts Institute of Technology (2014).

[16] N. Horelik, B. Herman, B. Forget, and K. Smith. "Benchmark for Evaluation and Validation of Reactor Simulations (BEAVRS), v1.0.1." In *Int. Conf. Math. and Comp. Methods Applied to Nuc. Sci. & Eng.* Sun Valley, Idaho, USA (2013).

[17] P. K. Romano and B. Forget. "The OpenMC Monte Carlo Particle Transport Code." *Annals of Nuclear Energy*, **volume 51**, pp. 274–281 (2013).

[18] G. C. Pomraning. "Flux-Limited Diffusion Theory with Anisotropic Scattering." *Nuclear Science and Engineering*, **volume 86**, pp. 335–343 (1984).

[19] Z. Liu, K. Smith, and B. Forget. "Calculation of multi-group migration areas in deterministic transport simulations." *Annals of Nuclear Energy*, p. 107110 (2019). URL http://www.sciencedirect.com/science/article/pii/S0306454919306206.

[20] K. Smith. "Assembly homogenization techniques for light water reactor analysis." *Progress in Nuclear Energy*, **volume 17**(3), pp. 303 – 335 (1986). URL http://www.sciencedirect.com/science/article/pii/0149197086900351.

[21] G. Giudicelli, K. Smith, and B. Forget. "Generalized equivalence methods for 3D multigroup neutron transport." *Annals of Nuclear Energy*, **volume 112**, pp. 9–16 (2018).

[22] A. Hébert. *Développement de la méthode SPH: Homogénéisation de cellules dans un réseau non uniforme et calcul des paramètres de réflecteur*. Ph.D. thesis, CEA (1981).

[23] H. Park and H. G. joo. "Practical resolution of angle dependency of multigroup resonance cross sections using parametrized spectral superhomogenization factors." *Nuclear Engineering and Technology*, **volume 49**, pp. 1287–1300 (2017).

[24] D. Gaston, B. Forget, K. Smith, and R. Martineau. "Verification of MOCkingbird, an Unstructured-Mesh, Method of Characteristics Implementation Using the MOOSE Multiphysics Framework." In *Int'l Conf. on Mathematics & Computational Methods Applied to Nuclear Science & Engineering*. Jeju, Korea (2017).

[25] G. Gunow, S. Shaner, W. Boyd, B. Forget, and K. Smith. "Accuracy and Performance of 3D MOC for Full-core PWR Problems." *Proceedings of M&C2017, Jeju Korea* (2017).

[26] G. Giudicelli, K. Smith, and B. Forget. "Generalized Equivalence Theory used with Spatially Linear Sources in the Method of Characteristics for Neutron Transport." *Nuclear Science and Engineering*, **volume 0**(0), pp. 1–12 (2020). URL https://doi.org/10.1080/00295639.2020.1765606.