Preliminary Accuracy and Performance Assessment of Depletion Analysis in the GPU-Optimized Monte Carlo Code GREAPMC Using MCS

Ian Nam[a], Murat Serdar Aygul[a], Kyeongwon Kim[a], Deokjung Lee[a*]
[a]*Department of Nuclear Engineering, Ulsan National Institute of Science and Technology, 50 UNIST-gil, Ulsan, 44919, Republic of Korea*
**Corresponding author: deokjung@unist.ac.kr*

Keywords : GREAPMC, GPU, Monte Carlo, depletion analysis, MCS

1. Introduction

Recent advances in graphics processing units (GPUs) have promoted the development of GPU-accelerated continuous-energy Monte Carlo codes for reactor analysis. Representative examples include PRAGMA developed at Seoul National University [1], Shift developed at Oak Ridge National Laboratory [2], and GPU-enabled OpenMC [3]. These codes have achieved substantial computational advantages over conventional CPU-based Monte Carlo calculations, with a single modern GPU providing performance comparable to hundreds of CPU cores for selected reactor problems. Such advances have improved the feasibility of high-fidelity Monte Carlo calculations for practical reactor analysis. GPU acceleration is particularly beneficial for depletion analysis because it requires repeated neutron transport calculations, reaction-rate tallying, and nuclide density updates over multiple burnup steps.

The Computational Reactor Physics & Experiment (CORE) Laboratory at Ulsan National Institute of Science and Technology (UNIST) is developing the GPU-optimized REActor Physics Monte Carlo code (GREAPMC). GREAPMC is a CUDA C++ continuous-energy Monte Carlo code based on history-based neutron tracking. Its continuous-energy transport performance was evaluated previously [4], and its constructive solid geometry modeling capability was verified in a subsequent study [5]. The version evaluated in the present work extends these capabilities with depletion analysis, an internal critical boron concentration (CBC) search, and multi-GPU domain decomposition. These newly implemented capabilities are evaluated here before application to practical full-core depletion calculations.

In this study, the newly implemented depletion framework of GREAPMC, including its CBC-search capability, is verified using a 3×3 PWR model derived from VERA Benchmark Problem 4 [6]. The GREAPMC depletion results up to 16.5 MWd/kgU are compared with reference solutions obtained using MCS [7], a continuous-energy Monte Carlo code developed by the UNIST CORE Laboratory. The comparison includes CBC, peaking factor (F_q), axial shape index (ASI), radial power distributions (RPD), and axial power distributions (APD). The preliminary computational performance of GREAPMC is also evaluated relative to MCS.

2. Methodologies

2.1. GREAPMC

Previous GREAPMC studies reported its continuous-energy transport capability [4] and constructive solid geometry modeling capability [5]. The version evaluated in this study additionally includes depletion, an internal CBC-search module, and multi-GPU domain decomposition. GREAPMC currently allocates a fixed transport array for 422 nuclides on each GPU and uses a 1,441-nuclide depletion set. These values correspond to the implementation used in the present calculations. Future work will reduce the fixed transport list using a retention criterion derived from representative depletion inventories, although this reduction is not applied in the present study. The GREAPMC calculation reported in this study uses assembly-wise domain decomposition. Accordingly, the reported GREAPMC runtime includes communication overhead associated with the domain-decomposed calculation, in addition to the computational and memory costs associated with the fixed transport-nuclide allocation.

2.2. Modeling

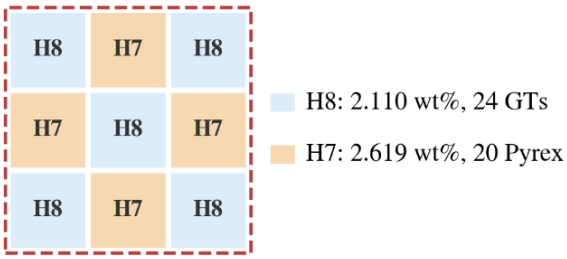


Fig. 1. The 3×3 H7 and H8-assembly loading pattern with reflective radial boundary conditions.

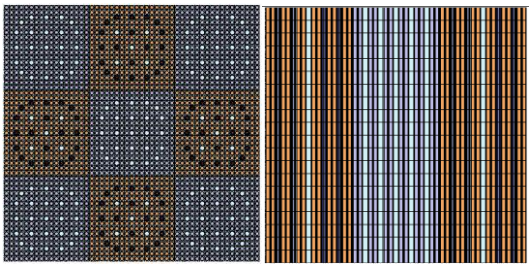


Fig. 2. Pin-level radial (left) and axial (right) geometry of the PWR model.

As shown in Fig. 1, the model is derived from VERA Benchmark Problem 4 and consists of 9 Westinghouse-type fuel assemblies arranged in a 3×3 configuration, with each assembly having a 17×17 lattice [6]. The H8 and H7 assemblies contain fuel enriched to 2.110 and 2.619 wt% U-235. Reflective boundary conditions are applied, and control rods are not included in the model.

Fig. 2 shows the radial and axial geometry of the pin-level model. The active fuel length is 365.76 cm and is divided into 20 equal axial meshes.

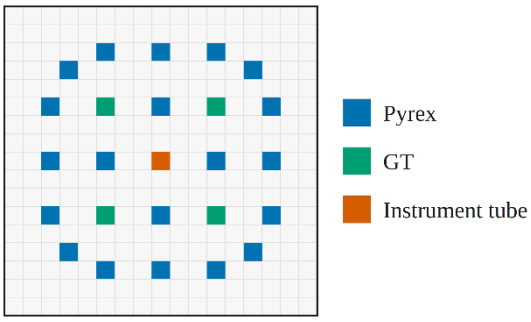


Fig. 3. The 17×17 lattice of the H7-assembly.

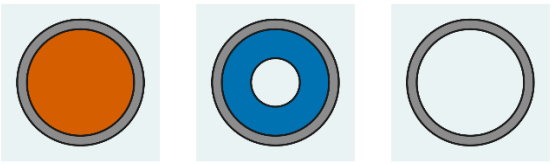


Fig. 4. Radial pin-cell models for (a) the fuel rod, (b) the Pyrex rod, and (c) the guide tube.

Fig. 3 and 4 show the H7-assembly lattice and radial pin-cell models. Each assembly has a 17×17 lattice comprising 264 fuel rods, 24 guide tubes, and one instrumentation tube. In the H7-assembly, 20 guide tubes contain Pyrex rods, and the remaining 4 are filled with moderator. All guide tubes in the H8-assembly are filled with moderator. The fuel pellet radius is 0.4096 cm, and the cladding inner and outer radii are 0.418 and 0.475 cm. The fuel rod pitch is 1.26 cm. The guide tube inner and outer radii are 0.561 and 0.602 cm, while those of the instrument tube are 0.559 and 0.605 cm. The assembly pitch is 21.50 cm, with a half-gap of 0.04 cm between assemblies. The Pyrex rods contain 12.5 wt% B_2O_3 glass and have inner and outer radii of 0.241 and 0.427 cm.

2.3. Calculation Conditions

GREAPMC and MCS use ENDF/B-VII.1 nuclear data, equilibrium xenon, SERPENT kappa data, a thermal power of 100 MW, the predictor depletion scheme, and the same burnup steps and number of neutron histories per cycle. Both calculations are performed at prescribed temperatures without thermal-hydraulic feedback. Fuel and Pyrex are set to 900 K, while moderator and structural materials are set to 600 K at a moderator density of 0.743 g/cm³. The depletion calculation contains 38 burnup steps from 0 to 16.5 MWd/kgU and uses 100 inactive and 100 active cycles. Beginning of cycle (BOC), middle of cycle (MOC), and end of cycle (EOC) are defined as 0.0, 8.0, and 16.5 MWd/kgU, respectively. Table I summarizes the common conditions and Table II summarizes the hardware and parallel-resource allocations.

Table I: Calculation conditions.

Parameter	Setting
Temperature	Fuel/Pyrex: 900 K moderator/structure: 600 K
Moderator density	0.743 g/cm³
Power / Library	100 MWth / ENDF/B-VII.1
Depletion	Predictor / equilibrium Xe 38 burnup steps to 16.5 MWd/kgU
Monte Carlo cycles	100 inactive + 100 active
Histories per cycle	1.00E+07 (10M)
Total histories	2.00E+09

Table II: Execution hardware and parallel allocation.

Case	Resources
GREAPMC (GPU)	7 RTX A5000 + 2 RTX 3090
MCS (CPU)	160 Intel Xeon Gold 6242/6242R

GREAPMC was executed on nine GPUs, comprising six RTX A5000 and two RTX 3090 GPUs on the first node and one RTX A5000 GPU on the second node. The 3×3 PWR model was partitioned into nine assembly-wise domains, with one assembly assigned to each GPU. The executable was built with CUDA 12.4. MCS was executed on Intel Xeon Gold 6242/6242R CPUs with one thread per MPI process.

3. Results

3.1. Critical Boron Concentration

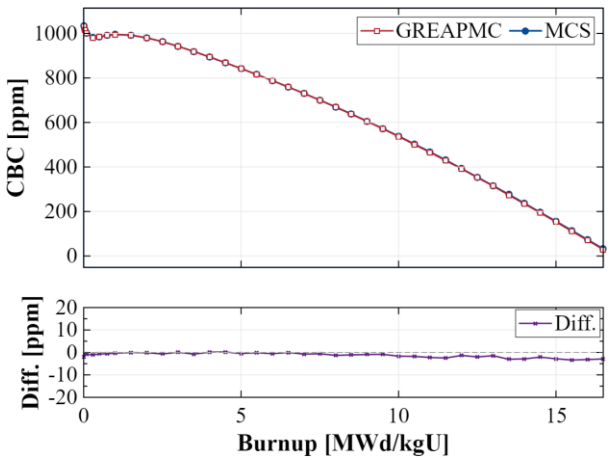


Fig. 5. CBC comparison between GREAPMC and MCS.

The GREAPMC and MCS CBC values are 1031.04 and 1033.10 ppm at BOC, 668.45 and 669.77 ppm at MOC, and 29.33 and 32.25 ppm at EOC. The absolute difference is 2.06 ppm at BOC, 1.32 ppm at MOC, and 2.92 ppm at EOC. Over the full burnup range, the maximum and mean absolute differences are 3.366 and 1.267 ppm, respectively. These results indicate that the CBC predictions of GREAPMC are in close agreement with those of MCS.

3.2. Axial Shape Index (ASI)

)

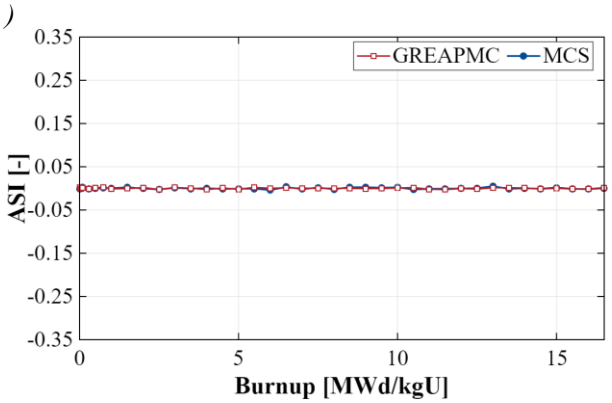


Fig. 6. ASI comparison between GREAPMC and MCS.

Fig. 6 compares the ASI values calculated by GREAPMC and MCS. The absolute differences are 0.002893 at BOC, 0.002120 at MOC, and 0.001307 at EOC. The corresponding relative differences are 0.289%, 0.317%, and 0.196%, respectively. The mean absolute difference is 0.001783, and the maximum absolute difference is 0.004408 at 11.0 MWd/kgU.

3.3. Peaking Factor (F_q)

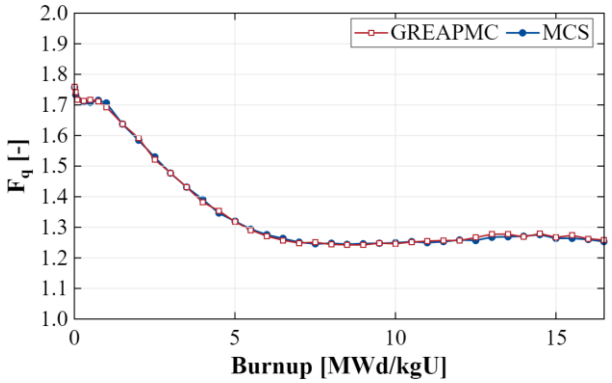


Fig. 7. F_q comparison between GREAPMC and MCS.

Fig. 7 compares the F_q values calculated by GREAPMC and MCS. At BOC, the GREAPMC and MCS values are 1.75965 and 1.75746, respectively. The corresponding values are 1.24382 and 1.24762 at MOC and 1.25833 and 1.25317 at EOC. The minimum and maximum absolute differences over the full burnup range are 0.000269 and 0.013217, respectively. The absolute relative differences range from 0.019% to 0.775%, with a mean absolute difference of 0.004629 and a mean absolute relative difference of 0.336%. At the representative BOC, MOC, and EOC, the differences are within combined 1σ statistical uncertainties of the two calculations. The maximum difference corresponds to approximately 1.56 combined standard deviations. These results indicate good statistical agreement throughout the depletion calculation between GREAPMC and MCS.

3.4. Radial Power Distribution

BOC 0.00 MWd/kgU

1.083	0.918	1.083	
0.02	0.02	0.02	
1.082	0.918	1.082	
0.02	0.01	0.02	
0.10	0.01	0.08	
0.918	0.998	0.918	Rel. Diff. [%]
0.01	0.01	0.01	RMS 0.06
0.917	0.998	0.918	Min -0.11
0.02	0.01	0.01	Max 0.10
0.02	0.01	-0.02	
1.083	0.917	1.082	GREAPMC
0.016	0.015	0.018	GREAPMC Rel. SD [%]
1.083	0.918	1.083	MCS
0.02	0.02	0.02	MCS Rel. SD [%]
-0.03	-0.11	-0.08	Rel. Error [%]

MOC 8.00 MWd/kgU

0.996	1.006	0.996	
0.01	0.01	0.02	
0.996	1.004	0.997	
0.01	0.01	0.02	
0.01	0.21	-0.11	
1.005	0.995	1.006	Rel. Diff. [%]
0.01	0.01	0.01	RMS 0.13
1.004	0.996	1.005	Min -0.23
0.01	0.01	0.01	Max 0.21
0.16	-0.05	0.05	
0.995	1.005	0.995	GREAPMC
0.01	0.01	0.02	GREAPMC Rel. SD [%]
0.996	1.004	0.997	MCS
0.01	0.02	0.02	MCS Rel. SD [%]
-0.12	0.07	-0.23	Rel. Error [%]

EOC 16.50 MWd/kgU

0.974	1.030	0.973	
0.01	0.01	0.01	
0.975	1.029	0.975	
0.02	0.02	0.02	
-0.17	0.13	-0.16	
1.030	0.988	1.029	Rel. Diff. [%]
0.01	0.01	0.01	RMS 0.14
1.028	0.988	1.028	Min -0.17
0.01	0.01	0.01	Max 0.21
0.21	-0.02	0.10	
0.974	1.030	0.973	GREAPMC
0.01	0.01	0.01	GREAPMC Rel. SD [%]
0.975	1.028	0.974	MCS
0.02	0.02	0.02	MCS Rel. SD [%]
-0.09	0.13	-0.16	Rel. Error [%]

Fig. 8. Assembly-wise radial power distribution at BOC (top), MOC (middle) and EOC (bottom).

Fig. 8 shows the radial assembly power distributions at BOC, MOC, and EOC. At BOC, the normalized power is approximately 1.08 in the corner H8-assemblies and 0.92 in the H7-assemblies containing Pyrex rods. At MOC, the power distribution becomes nearly uniform, with normalized values ranging from 0.995 to 1.006. At EOC, the BOC distribution is reversed, with normalized power ranging from 1.028 to 1.030 in the H7-assemblies and from 0.973 to 0.975 in the corner H8-assemblies. The maximum absolute relative differences between GREAPMC and MCS are 0.113% at BOC, 0.231% at MOC, and 0.215% at EOC, with RMS errors of 0.065%, 0.134%, and 0.142%, respectively. The absolute relative differences remain below 0.25% for all assemblies throughout the depletion calculation.

3.5. Axial Power Distribution

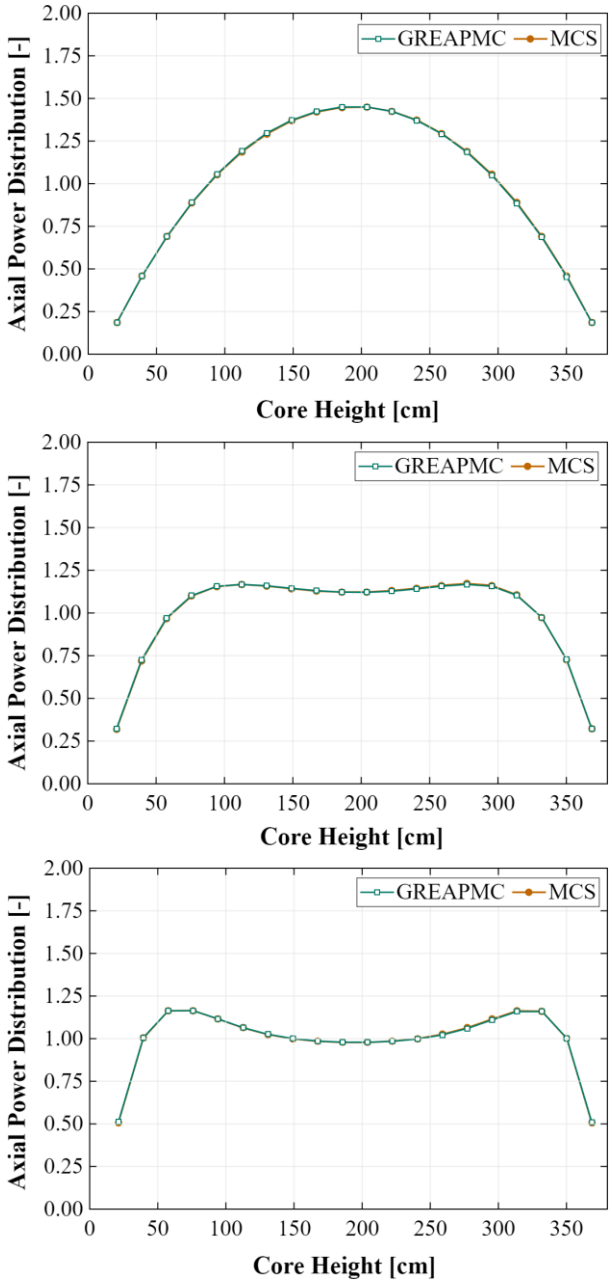


Fig. 9. Axial power distributions at BOC (top), MOC (middle), and EOC (bottom).

Fig. 9 shows the axial power distributions at BOC, MOC, and EOC, each normalized to an axial mean of 1.0. At BOC, the maximum normalized powers are 1.4500 and 1.4493 for GREAPMC and MCS, respectively. At MOC, peak values are 1.1673 and 1.1721. At EOC, maximum powers are 1.1640 and 1.1644. The BOC, MOC, and EOC RMS errors are 0.175%, 0.176%, and 0.112%, respectively, while the maximum local differences are 0.895%, 1.227%, and 0.889%, respectively. Table III summarizes the differences between GREAPMC and MCS at representative burnup steps.

Table III: Summary of differences between GREAPMC and MCS (GREAPMC − MCS).

Values	BOC	MOC	EOC
CBC (ppm)	-2.060	-1.319	-2.921
F_q (−)	0.00219	-0.0038	0.00516
ASI (−)	0.00289	0.00212	0.00131
RPD RMS (%)	0.065	0.134	0.142
RPD max. abs. (%)	0.113	0.231	0.215
APD RMS (%)	0.175	0.176	0.112
APD max. abs. (%)	0.895	1.227	0.889

At the representative BOC, MOC, and EOC, the CBC difference is at most 2.921 ppm, the RPD and APD RMS errors remain below 0.142% and 0.176%, respectively, and the Fq and ASI differences remain small, indicating overall consistency between GREAPMC and MCS throughout the depletion calculation.

3.6. Computational Performance

Table IV and V summarize the preliminary runtime performance of GREAPMC and MCS for the VERA benchmark. In the following discussion, the 1M and 10M cases refer to calculations using 1.0×10^6 and 1.0×10^7 neutron histories per cycle, respectively.

Table IV: Preliminary runtime performance for 1M case.

Case	GREAPMC	MCS
Processor Type	GPU	CPU
No. of Processors	9	72
Multi-Processor Computation Time	8.99 hours	126.93 hours
Single-Processor Equivalent Time	81 hours	9,139 hours
Speedup (1 GPU vs. 1 CPU)	112.95 faster	1.0 baseline

For the 1M case, GREAPMC required 8.99 h using 9 GPUs, whereas MCS required 126.93 h using 72 CPU cores. This corresponds to the complete 9-GPU configuration providing throughput equivalent to 1,016.57 CPU cores, or approximately 1,017 cores. Dividing this value by nine gives an average throughput equivalent to 112.95 CPU cores per GPU. Therefore, 1 GPU provides, on average, approximately 112.95 times the throughput of 1 MCS CPU core.

Table V: Preliminary runtime performance for 10M case.

Case	GREAPMC	MCS
Processor Type	GPU	CPU
No. of Processors	9	160
Multi-Processor Computation Time	34.64 hours	628.6 hours
Single-Processor Equivalent Time	312 hours	100,576 hours
Speed up (1 GPU vs. 1 CPU)	322.6 faster	1.0 baseline

For the 10M case, GREAPMC required 34.64 h using 9 GPUs, whereas MCS required 628.6 h using 160 CPU processes. This corresponds to an overall system-level speedup of 18.15 times. The complete 9-GPU GREAPMC calculation corresponds to 312 GPU-hours, whereas the 160-process MCS calculation corresponds to 100,576 CPU-hours. Based on this resource-normalized comparison, 1 GPU provides approximately 322.6 times the throughput of 1 MCS CPU process. The reported GPU performance values are averaged over the 7 RTX A5000 and 2 RTX 3090 GPUs and are not architecture, power, or cost-normalized. The comparison uses the same histories per cycle, cycle counts, burnup mesh, and CBC-search calculation.

For comparison with the 1M case, increasing the number of neutron histories by a factor of ten increased the GREAPMC runtime from 8.99 h to 34.64 h, corresponding to a 3.85-fold increase in runtime. Consequently, the average particle throughput increased by a factor of 2.60. This sublinear increase in runtime suggests that the GPU resources were not fully utilized. The first 2–3 MCS burnup steps are faster than later steps as its active isotope list grows, whereas GREAPMC retains the fixed 422-nuclide allocation at every step.

4. Conclusion

The preliminary GREAPMC depletion calculation was assessed in comparison with MCS for the 3×3 PWR model up to 16.5 MWd/kgU. The maximum absolute CBC difference over all burnup steps was 3.366 ppm, while the RPD and APD RMS errors remained below 0.142% and 0.176%. The F_q and ASI results also agreed closely with the MCS results throughout the depletion calculation. These results demonstrate that GREAPMC closely predicts the reactivity and spatial power behavior calculated by MCS for the present model.

For the 10M-particle case, the 9-GPU GREAPMC calculation required 34.64 h, whereas the 160-process MCS calculation required 628.6 h. This corresponds to an overall system-level speedup of 18.15 times. On a resource-normalized basis, the complete GPU configuration corresponds to 312 GPU-hours, while the CPU calculation corresponds to 100,576 CPU-hours, resulting in an average throughput advantage of approximately 322.6 times for 1 GPU over 1 MCS CPU process. Because the present calculation uses a fixed transport-nuclide list and a non-optimized depletion layer, these values do not represent the final optimized performance.

Future work will reduce the transport-nuclide list, apply the optimized depletion layer, integrate thermal-hydraulic feedback, and extend the verification to full-core reactor depletion analyses. These developments will support the application of GREAPMC to efficient high-fidelity full-core depletion analysis.

REFERENCES

[1] N. Choi, K. M. Kim, H. G. Joo, "Initial Development of PRAGMA – A GPU-Based Continuous Energy Monte Carlo Code for Practical Applications," Transactions of the Korean Nuclear Society Autumn Meeting, Goyang, Korea, Oct. 24-25, 2019.

[2] S. P. Hamilton and T. M. Evans, "Continuous-energy Monte Carlo neutron transport on GPUs in the Shift code," Annals of Nuclear Energy, Vol. 128, pp. 236-247, 2019.

[3] J. R. Tramm, P. K. Romano, J. Doerfert, A. Lund, P. Shriwise, A. Siegel, G. Ridley, A. Pastrello, "Toward Portable GPU Acceleration of the OpenMC Monte Carlo Particle Transport Code," International Conference on Physics of Reactors (PHYSOR 2022), Pittsburgh, PA, USA, 2022.

[4] M. R. Ali, M. S. Aygül, and D. Lee, "Performance Evaluation of GREAPMC for Continuous Energy Reactor Physics Problems," in Proceedings of the International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2025), Denver, Colorado, USA, pp. 1811–1820, Apr. 27–30, 2025, doi: 10.13182/MC25-47517.

[5] H. Sohn, M. S. Aygül, and D. Lee, "Verification of Constructive Solid Geometry Modeling in GPU-Optimized Reactor Physics Monte Carlo Code, GREAPMC," Transactions of the KNS Spring Meeting, Jeju, Korea, May 7–8, 2026.

[6] A. T. Godfrey, "VERA Core Physics Benchmark Progression Problem Specifications," Rev. 4, CASL-U-2012-0131-004, ORNL, 2014.

[7] T. D. C. Nguyen, H. Lee, S. Choi, and D. Lee, "Validation of UNIST Monte Carlo Code MCS Using VERA Progression Problems," Nuclear Engineering and Technology, Vol. 52, No. 5, pp. 878–888, 2020.