Improved Performance of the P_N Scattering MOC Solvers in MPACT

Shane Stimpson and Benjamin Collins

Oak Ridge National Laboratory, 1 Bethel Valley Rd., Oak Ridge, TN 37831, stimpsonsg@ornl.gov

INTRODUCTION

As part of the Consortium for Advanced Simulation of Light Water Reactors (CASL), the Virtual Environment for Reactor Applications (VERA) is being developed to provide high-fidelity multiphysics simulations of nuclear reactor cores [1,2]. The MPACT code [3] being developed collaboratively by Oak Ridge National Laboratory and the University of Michigan is the primary deterministic neutron transport solver available within VERA.

MPACT employs the 2D/1D method to solve 3D problems using 2D method of characteristics (MOC) radial transport solvers with 1D nodal expansion method-based P_3 (NEM-P_3) axial transport solvers coupled using axial and radial transverse leakages. A 3D coarse mesh finite difference (CMFD) solver is used to provide accelerated convergence and stability [4]. Because it uses fast, isotropic MOC kernels, transport-corrected P_0 scattering is the default approach used for production level cases in MPACT. However, for benchmarking purposes, and as MPACT is being applied to more advanced reactor types, the performance of the P_N scattering MOC kernels is becoming more important.

Recent efforts to improve the performance and stability of the P_0 scattering MOC kernels has focused on applying multigroup (MG) kernels that solve all groups together in the kernel rather than solving a single-group kernel. This approach, which inherently forces a Jacobi inscatter approximation, was observed to be more stable than the one-group kernel method, which used a Gauss-Seidel inscatter source. It also improves the calculation efficiency, yielding a speedup of roughly a factor of two. For additional details on the MG kernels, please consult Ref. 5, which presents further detail and analysis with respect to isotropic scattering kernels.

Such kernels eliminate duplicate work performed in constructing long rays in the sweeper, which are repeated for every group with one-group kernels. Similarly, MG kernels have been applied to the P_N MOC kernels, along with additional improvements to remove unnecessary angular moments for 2D problems. This latter improvement is significant in that it provides speed and reduces overall memory requirements, and the percentage of angular moments removed increases with the scattering order. Table I shows the number of angular moments used between the untrimmed (default) and trimmed (new) approaches. As can be seen, eliminating the unnecessary moments can lead to more than 40% savings with P_5, and more than 20% for the more commonly used P_2 scattering. This type of approach is common among production-level transport codes.

Table I. Number of Angular Moments with Respect to Scattering Order (Trimmed vs Untrimmed)

Scattering	Untrim	Trim	Ratio
P_0	1	1	1.00
P_1	4	3	0.75
P_2	9	6	0.67
P_3	16	10	0.63
P_4	25	15	0.60
P_5	36	21	0.58

RESULTS

Both the MG and trimmed moment approaches have been tested on Progression Problems 2a (quarter lattice, Fig. 1) and 5a-2D (quarter core, Fig. 2) [5]. Both problems where analyzed using the default 51-group library generated at ORNL. The effectiveness of the MG kernels is highly dependent upon the number of groups being used. The higher the number of groups, the more effective the kernel is as more redundant work from the 1G kernel is eliminated. This holds true up to a certain point where the additional memory burden of the MG kernels can eventually yield a degradation of performance.

Table II shows the MOC run time for problem 2a. For clarity, 1G Untrimmed was the result before any of the improvements here were implemented. 1G Trimmed is only with the unnecessary angular moments removed, where ~20% run time reduction is observed for P_2. Moving into the MG results, the untrimmed results show the impact of only the MG kernels. For P_0, a roughly 2× improvement is observed, which is consistent with previous efforts, as well as a 30% reduction for P_2 scattering. Combining the trimming and MG kernels, P_2 scattering performance is improved by a factor of 1.91. These improvements make larger scale production-level cases more tractable.

It should be noted that the effect of using multigroup kernels decreases as the scattering order increases. This is primarily due to the fact that the relative importance of the eliminated duplicate work is decreasing. Using multigroup kernels

allows the work spent setting up the MOC rays to be performed only once, whereas 1G kernels repeat the process for every kernel instantiation (so for each group). When using P_0 scattering, this effort is substantially higher relative to P_5 scattering, which spends more time evaluating the transport calculation.

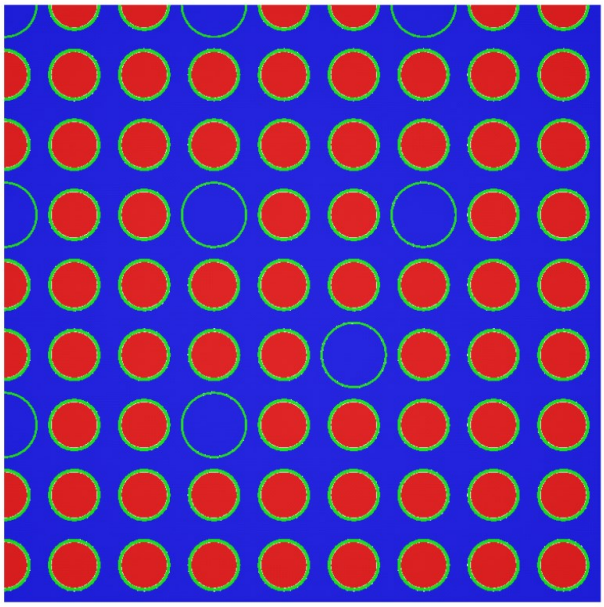

Fig. 1. Single 17 × 17 quarter lattice, Problem 2a [5].

Table II. MOC Run Time (Seconds) for Progression Problem 2a

	1G		MG		Speedup
Scat.	Untrim	Trim	Untrim	Trim	(1G-Untrim/ MG-Trim)
P_0[1]	38.35	38.55	17.56	17.50	2.19
P_1	42.65	40.84	25.43	22.07	1.93
P_2	61.69	49.04	42.29	32.34	1.91
P_3	76.48	59.87	64.32	45.38	1.69
P_4	105.02	80.30	91.95	60.36	1.74
P_5	129.67	93.35	126.01	89.84	1.44

However, these improvements come at a cost in terms of the memory burden (Table III). While removing unused moments reduces the memory burden, particularly when using an MG kernel, using MG kernels incurs some burden, as all groups of outgoing angular fluxes and temporary scalar fluxes must be stored. Overall, this leads to a 13% increase in the memory burden for this problem when using P_2 scattering. With 1G kernels, the relative importance of the trimming is smaller with respect to the total memory as the memory to store scattering source is minimized as only one group is needed at a time.

[1] P_0 here represents a P_N scattering kernel with zero angular moments, which is slightly different from the isotropic scattering kernels

Table III. Overall Memory Requirements (MB) for Progression Problem 2a

	1G		MG		Mem. Ratio
Scat.	Untrim	Trim	Untrim	Trim	(MG-Trim/ 1G-Untrim)
P_0[1]	172.1	172.1	188.1	188.2	1.09
P_1	184.2	184.4	210.0	204.2	1.11
P_2	199.9	199.9	251.9	225.0	1.13
P_3	219.8	219.9	321.3	260.6	1.19
P_4	240.6	240.5	397.8	312.8	1.30
P_5	265.6	265.6	491.2	363.8	1.37

Figure 2 shows the radial slice of the Watts Bar Nuclear Unit 1 (WBN1) quarter core layout, including the enrichments (2.1%, 2.6%, and 3.1%), as well as the burnable poison rods in each assembly. The figure also depicts a core baffle and jagged reflector region along the periphery. By default, the reflector region spans an assembly width from the outer fuel assembly, including both the baffle and moderator.

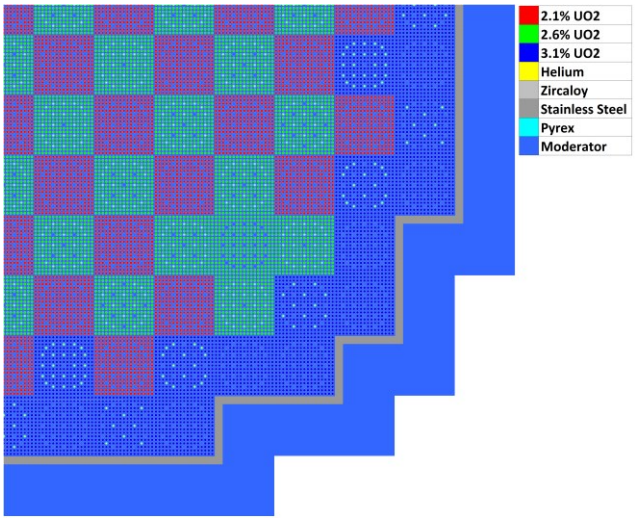

Fig. 2. WBN1 – radial slice of core layout [5].

The full WBN1 core contains 193 fuel assemblies. With quarter symmetry, this amounts to 257 quarter assembly models, including the reflector.

Similarly to the previous problem, Tables IV and V show the results for the quarter core 5a-2D problem, which was run with 16 spatial domains. Each domain is larger than the 2a case, so there are some differences in trends observed. For P_2 scattering, a 2.15× speedup is observed in the MOC runtime, and a 34% increase is seen in the overall memory requirements. This speedup decreases as scattering order increases, but a 1.6× speedup is observed at a 42% increase in memory for P_5 scattering.

Table IV. MOC Run Time (Seconds) for Progression Problem 5a-2D

	1G		MG		Speedup
Scat.	Untrim	Trim	Untrim	Trim	(1G-Untrim/ MG-Trim)
P_0[1]	12.20	12.08	5.06	5.32	2.29
P_1	17.25	14.79	8.91	7.99	2.16
P_2	23.98	17.57	14.97	11.17	2.15
P_3	29.06	22.74	24.71	16.87	1.72
P_4	40.17	30.00	34.83	21.59	1.86
P_5	48.67	35.38	51.42	30.45	1.60

Table V. Overall Memory Requirements (GB) for Progression Problem 5a-2D

	1G		MG		Mem. Ratio
Scat.	Untrim	Trim	Untrim	Trim	(MG-Trim/ 1G-Untrim)
P_0[1]	16.42	16.42	20.48	20.48	1.25
P_1	18.79	18.44	26.15	24.51	1.30
P_2	23.06	21.98	36.81	30.90	1.34
P_3	28.70	26.56	51.40	39.58	1.38
P_4	34.94	31.36	69.11	49.42	1.41
P_5	42.55	37.18	90.77	60.56	1.42

CONCLUSIONS AND FUTURE WORK

In summary, notable improvements to the P_N MOC kernels have been made by incorporating MG kernels and removing unnecessary angular moments for the 2D cases. Since MPACT employs 2D-MOC in all 2D/1D calculations, the improvements here will directly impact the runtimes of 2D/1D cases similarly. These improvements will be very useful as anisotropic scattering is considered for more advanced reactor designs. In general, transport-corrected P_0 scattering works well for LWRs, but this is not necessarily as true for more exotic designs in which such approximations may lead to higher errors in the power distribution estimates. Additionally, fully accounting for the anisotropy in 3D with 2D/1D method, would require anisotropic axial leakage terms in the 2D-MOC solvers; the current default for production-level cases is to use isotropic leakage terms.

Though these improvements yielded a net increase in the memory burden (~34% increase for P_2 scattering), the memory requirement is still manageable when spread over a moderate number of cores through radial decomposition. MPACT typically runs with 16 radial domains, which would yield roughly 2 GB per core. As future work applies this approach to 3D quarter core problems, the authors expect this memory burden to still be practical, particularly on machines with 4+GB of memory per core, which is fairly typical.

ACKNOWLEDGMENTS

This work was supported by the Consortium for Advanced Simulation of Light Water Reactors (www.casl.gov), an energy innovation hub (http://www.energy.gov/hubs) for modeling and simulation of nuclear reactors under U.S. Department of Energy (DOE) Contract No. DE-AC05-00OR22725.

REFERENCES

1. "Consortium for Advanced Simulation of Light Water Reactors (CASL)." (2015) http://www.casl.gov/.
2. J. TURNER et al., "The Virtual Environment for Reactor Applications (VERA): Design and architecture," *Journal of Computational Physics*, **326**, 544 (2016).
3. MPACT TEAM, *MPACT Theory Manual,* Version 2.1.0, Technical Report CASL-U-2015-0078-001, Oak Ridge National Laboratory and University of Michigan (2017).
4. B. COLLINS et al., "Stability and Accuracy of 3D Neutron Transport Simulations Using the 2D/1D Method in MPACT," *Journal of Computational Physics*, **326**, 612 (2016).
5. S. STIMPSON, B. COLLINS, and B. KOCHUNAS, "Improvement of Transport-Corrected Scattering Stability and Performance Using a Jacobi Inscatter Algorithm for 2D-MOC," *Annals of Nuclear Energy*, **105**, 1–10 (2017).
6. A. GODFREY et al. *VERA Benchmarking Results for Watts Bar Nuclear Plant Unit 1 Cycles 1–12*. Technical Report CASL-U-2015-0206-000, Oak Ridge National Laboratory. http://www.casl.gov/docs/CASL-U-2015-0206-000.pdf (2015).