Implementation of a Red-Black SOR CMFD Solver in MPACT

Shane Stimpson and Benjamin Collins

Oak Ridge National Laboratory
1 Bethel Valley Rd., Oak Ridge, TN 37831
stimpsonsg@ornl.gov
collinsbs@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 being developed collaboratively by Oak Ridge National Laboratory and the University of Michigan includes the primary deterministic neutron transport solver available within VERA.

MPACT employs the 2D/1D method to solve three-dimensional problems, using 2D method of characteristics (MOC) radial transport solvers with 1D NEM-SP_3 axial transport solvers, which are coupled using axial and radial transverse leakages [3]. A 3D coarse mesh finite difference (CMFD) solver [4] is used to provide accelerated convergence that is more stable.

The Portable, Extensible Toolkit for Scientific Computation (PETSc) [5] generalized minimal residual (GMRES) iterative solver with a block Jacobi preconditioner has been used to solve the CMFD linear systems. This work focused on the addition of a red-black successive over-relaxation (RBSOR) solver to reduce overall runtime by reducing time spent solving the CMFD linear systems.

THEORETICAL BACKGROUND

Coarse Mesh Finite Difference (CMFD)

CMFD is one of the mostly widely used acceleration techniques for neutron transport solvers [4]. With this approach, the fine mesh used for the transport solvers, which can include sub-pin level resolution, is homogenized to produce coarse mesh-averaged cross sections and fluxes. The size of the coarse mesh can vary based on the method, but MPACT uses pin-wise coarse mesh cells. From the homogenized parameters, coupling coefficients are formed, one being the traditional finite difference coefficients and the other a nonlinear current coupling coefficient which uses the higher fidelity transport-tallied currents on the coarse mesh surfaces. The details of these coefficients will be omitted here, but they are well documented and can be found in several references, such as the MPACT theory manual [6].

These coupling coefficients are then used to construct a linear system relating the multigroup coarse cells fluxes. This system can be solved by any number of means, though the parallel capabilities within the PETSc solver framework have become a popular and efficient means of solving systems like these.

Once the system is solved and the new coarse mesh fluxes have been obtained, the fine mesh fluxes are updated by multiplying them by a ratio of the new and old coarse mesh fluxes.

Red-Black Gauss-Seidel

As previously mentioned, the CMFD linear systems can be solved through a variety of techniques. One of the most basic iterative solvers is the Gauss-Seidel scheme, which loops over the coarse mesh cells, solving for the flux in each cell using the most up-to-date neighboring cell data available.

There are several different techniques for determining the order in which the solution is obtained; one popular approach is the red-black scheme [7], which tags each coarse cell as red or black to produce a checkerboard pattern (as in Fig. 1 for a simple 7×7 problem).

Fig. 1. Red-black indexing.

It should be noted that not all numerical index patterns will be lexicographic or naturally ordered as is observed in Fig. 1. For example, the coarse cells in MPACT are naturally ordered within each quarter assembly, so it does not hold that core-wide indexing is lexicographic.

Figure 2 is a flow chart of the solver's iteration strategy, which involves looping over and solving the fluxes for the red indexes, passing data as necessary, looping over the black indexes, and again passing data as necessary.

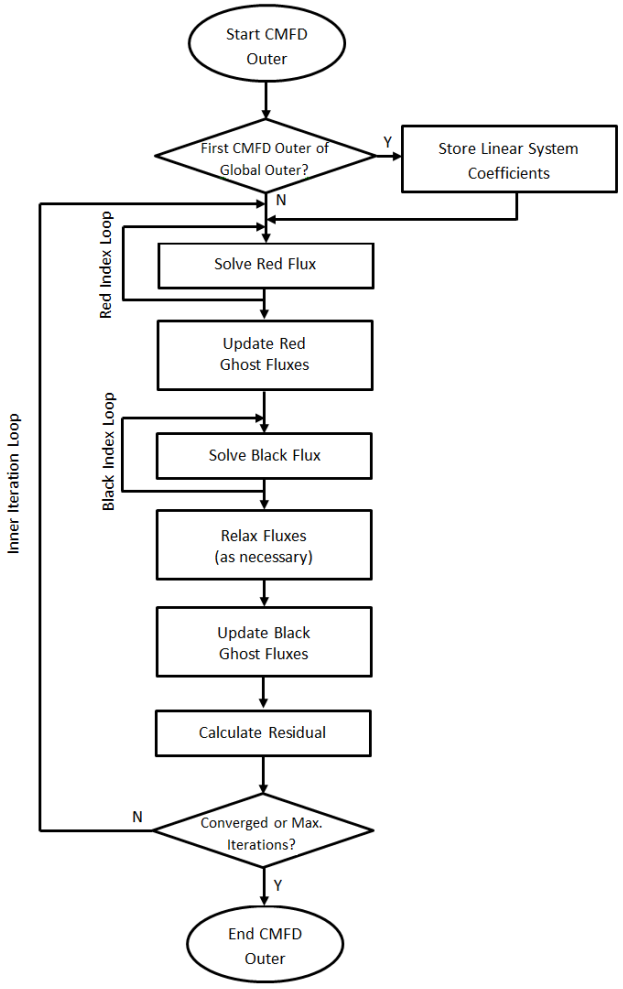

Fig. 2. Red-black SOR CMFD iteration strategy.

Because the inner iteration convergence for Gauss-Seidel solvers tends to be slower than that for many other iterative solvers, each CMFD eigenvalue update iteration is restricted to performing a user-specified maximum number of iterations. In the results shown later, a maximum of 50–100 inner iterations seems sufficient for 2D cases, but 3D cases with feedback perform best with a bound of 100–150 iterations. Otherwise, the differences in the total number of outer iterations are noticeable.

To increase the applicability of the solver, both message passing interface (MPI) and open multi-processing (OpenMP) are currently available to provide parallelization. When MPACT is executed, the user provides inputs specifying the number of processors used for spatial decomposition (MPI), MOC angle decomposition (MPI), and MOC ray decomposition (OpenMP). The current solver takes advantage of the processors used for spatial decomposition, as well as the OpenMP threads, which are repurposed from the application to MOC to provide additional spatial decomposition during the sweeps over the red and black indexes. One advantage of this approach is that PETSc is restricted to using MPI [5], though a hybrid MPI/OpenMP approach is being considered by the PETSc development team. While future work should consider repurposing the processors dedicated to angle decomposition for MOC, this work only evaluates cases with spatial decomposition.

Successive Over-Relaxation (SOR)

Gauss-Seidel is a special case of the successive over-relaxation (SOR) solver, where the relaxation factor equals unity. However, SOR typically applies a relaxation factor larger than unity to accelerate the convergence of the system, as in Eq. 1 [7,8]:

$$\phi_{n,g}^{i} = \omega\phi_{n,g}^{i} + (1-\omega)\phi_{n,g}^{i-1}, \qquad (1)$$

where i denotes the iteration index, n denotes the cell index, and g denotes the group index.

The determination of optimal relaxation factors has been studied extensively. This implementation uses adaptive relaxation factors based on the Cyclic Chebyshev Semi-Iterative (CCSI) method [8,9], where the red and black fluxes use different relaxation factors (Eqs. 2–3), which eventually converge to the same value. Equation 2 shows the initial relaxation factors, and Eq. 3 shows the relaxation factor for subsequent iterations:

$$\omega_R = 1, \qquad \omega_B = \frac{1}{1-\frac{1}{2}\rho_J^2}, \qquad (2)$$

$$\omega_R^{i+1} = \frac{1}{1-\frac{1}{4}\rho_J^2\omega_B^i}, i \geq 1, \qquad \omega_B^{i+1} = \frac{1}{1-\frac{1}{4}\rho_J^2\omega_R^{i+1}}, i \geq 1. \qquad (3)$$

These equations make use of the Jacobi spectral radius, which, in this work, is estimated during the first 10 inner iterations, applying no relaxation, to obtain the Gauss-Seidel spectral radius. The Gauss-Seidel spectral radius is then converted to the Jacobi spectral radius.

It is worth noting that these relaxation factors only seem to be valid when rotational symmetry is not applied to any boundaries. With rotational symmetry, some red-red and black-black neighbors will occur along those boundaries,

which will likely invalidate the underlying theory of the derivation. The results presented in this work exclusively consider mirror symmetry, though roughly similar time reductions are observed with rotational symmetry where the relaxation is disabled.

Global Iteration Strategy

What has been presented to this point pertains to the inner iteration strategy with RBSOR. To clarify some of the results, information on the global iteration strategy is presented. An "outer iteration" is considered to be a single CMFD converged solve plus a single MOC sweep. A CMFD converged solve involves several eigenvalue iterations (maximum of 20) with the goal of reducing the residual (2-norm of Ax-b) by two orders of magnitude. Each eigenvalue iteration involves inner iterations as previously specified, where GMRES or RBSOR are used to solve the linear system. An MOC sweep can also consist of a few inner iterations to resolve self- and up-scatter distributions. Outer iterations are repeated until global convergence of the eigenvalue and fission source is achieved.

RESULTS

To demonstrate this new capability and compare it to the preexisting PETSc GMRES solver, a zero-power 2D quarter core problem known as VERA Problem 5-2D [10] is evaluated. This case was run on Titan [11] with 73 spatial decomposition cores.

2D Quarter Core (VERA Problem 5a)

Figure 3 shows the assembly layout of the quarter core, including the enrichments (2.1%, 2.6%, and 3.1%), as well as the number of Pyrex burnable poison rods in each assembly. The problem also includes a core baffle and jagged reflector region.

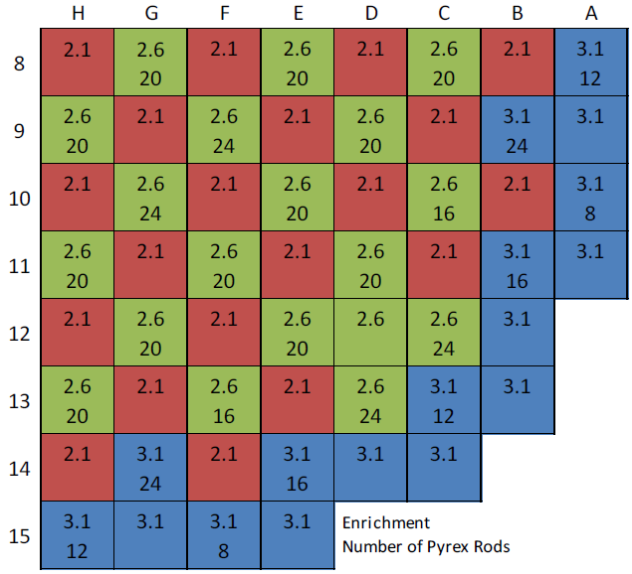

Fig. 3. Quarter core assembly layout [10].

Tables I and II show the results for PETSc GMRES and the RBSOR solvers, varying the maximum number of inner iterations allowed in each eigenvalue update. With only 25 inner iterations, the converged eigenvalue was ~150 pcm from the asymptotic value, despite the 1×10^{-6} convergence criteria on the eigenvalue. Even with 50 inner iterations, the eigenvalue is 2 pcm different. Therefore, at least 75 inner iterations are necessary to guarantee true convergence. Not surprisingly, as the number of inner iteration increases, the total number of outer iterations does not decrease, and more time is spent solving CMFD, to the detriment of the overall runtime.

Table I. PETSc GMRES Results for 2D Quarter Core

Inner Iterations	Eigenvalue	Outer Iterations	Total Time (min)	CMFD Solve Time (min)
25	1.00159	11	5.35	2.49
50	1.00310	11	7.43	4.62
75	1.00312	11	8.64	5.83
100	1.00312	11	9.99	7.20
150	1.00312	11	13.06	10.28
200	1.00312	11	13.84	11.05

The RBSOR solver does bring some noteworthy reductions in the run time, with 50 inner iterations yielding the best performance, reducing the total run time by roughly 39% (5.26 min vs 8.64 min) and the CMFD solve time by slightly more than a factor of two (2.86 vs. 5.83 min).

Table II. RBSOR Results for 2D Quarter Core

Inner Iterations	Eigenvalue	Outer Iterations	Total Time (min)	CMFD Solve Time (min)
25	1.00312	15	7.00	3.77
50	1.00312	11	5.62	2.86
75	1.00312	11	6.06	3.35
100	1.00312	11	6.48	3.78
150	1.00312	11	6.90	4.20
200	1.00312	11	7.34	4.64

CONCLUSIONS

Though the red-black SOR scheme is one of the most basic iterative methods, it performs well compared to more modern approaches like GMRES, at least within the context of these applications. For the 2D quarter core modelled, RBSOR does provide some notable reduction in the overall run time (~39%) by reducing the CMFD solve time by roughly a factor of two.

Additional developments such as incorporating advanced eigensolver techniques—as compared to the shifted power iteration used here—are already underway [12], and the RBSOR solver will provide for a more thorough comparison.

ACKNOWLEDGMENTS

This research 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 US Department of Energy Contract No. DE-AC05-00OR22725.

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

REFERENCES

1. "Consortium for Advanced Simulation of Light Water Reactors (CASL)." (2015) http://www.casl.gov/.
2. S. PALMTAG, "Demonstration of Neutronics Coupled to Thermal-Hydraulics for a Full-Core Problem using VERA," CASL-U-2013-0196-000 (2013).
3. B. S. COLLINS et al, "Three-Dimensional Nuclear Reactor Core Simulations of the Boltzmann Transport Equations with the 2D/1D Method Using MPACT," *Journal of Computational Physics* (forthcoming).
4. K. S. SMITH, "Nodal Method Storage Reduction by Nonlinear Iteration," *Trans. Am. Nucl. Soc.*, **44**, 265 (1983).
5. S. BALAY et al., "PETSc User Manual," ANL-95/11 – Revision 3.5, Argonne National Laboratory, http://www.mcs.anl.gov/petsc (2014).
6. MPACT Team, *MPACT Theory Manual, Version 1.0*, University of Michigan, Ann Arbor, MI (2013).
7. J. W. DEMMEL. *Applied Numerical Linear Algebra*, SIAM, 1997.
8. R. S. VARGA, *Matrix Iterative Analysis*, Prentice-Hall, Englewood Cliffs, N.J. (1962).
9. J. GEHIN, "A Quasi-Static Polynomial Nodal Method for Reactor Analysis," Ph. D. Dissertation, Massachusetts Institute of Technology, September (1992).
10. A. T. GODFREY, "VERA Core Physics Benchmark Progression Problem Specifications," CASL-U-2012-0131-004 (2014).
11. OAK RIDGE LEADERSHIP COMPUTING FACILITY, "Introducing Titan – The World's #1 Open Science Supercomputer," http://www.olcf.ornl.gov/titan/ (2013).
12. B. COLLINS et al. "Performance Improvements for Coarse Mesh Finite Difference Acceleration," CASL-X-2016-1106-000, Oak Ridge National Laboratory (2016).