

Acceleration Methods for Whole Core Reactor Simulations Using VERA

Benjamin Collins, Shane Stimpson

Oak Ridge National Laboratory, 1 Bethel Valley Rd. PO Box 2008 MS-6172 Oak Ridge, TN 37831 collinsbs@ornl.gov

INTRODUCTION

In recent years, development of VERA-CS—the core simulator for the Consortium for the Advanced Simulation of Light Water Reactors (CASL) [1]—has focused on simulating multiple cycles of operating commercial nuclear power plants. Now that the VERA-CS capabilities have advanced to the point where they are being deployed to users, attention is now directed on improving the computational performance of various components in VERA-CS. The focus of this work is the coarse mesh finite difference [2,3] (CMFD) solution in MPACT [4]. CMFD serves multiple purposes in the 2D/1D solution methodology. First, it is a natural mechanism to tie together the 2-D radial method of characteristics (MOC) transport and the 1D axial NEM-P₃ solution. Because the CMFD system solves the multigroup three-dimensional core in one system, it pulls together the global response of the system. In addition, the CMFD solution provides a framework to accelerate the convergence of the eigenvalue problem.

The CMFD methodology is based on a finite-difference approximation of the zeroth angular-moment of the neutron transport equation homogenized onto a 3D Cartesian grid:

$$\sum_{\substack{f \in n,s,e, \\ w,t,b}} J_{g,f} \square A_f + \bar{\Sigma}_{t,g,i} \bar{\phi}_{g,i} V_i = \sum_{g'} \bar{\Sigma}_{s,g' \rightarrow g,i} \bar{\phi}_{g',i} V_i + \frac{\bar{\chi}_{g,i}}{k_{eff}} \sum_{g'} \nu_{f,g',i}^- \bar{\phi}_{g',i} V_i \quad (1)$$

The CMFD methodology creates a finite-difference relationship between the surface current and flux as follows:

$$J_s = - \frac{2D^- D^+}{(D^+ + D^-)} (\bar{\phi}^+ - \bar{\phi}^-) + \hat{\mathbf{M}} (\bar{\phi}^+ + \bar{\phi}^-), \quad (2)$$

where + and - denote the cells on each side of any given surface s and Δ defines the distance between the two cells. This equation defines the relationship between the current and the cell average flux using a coarse mesh diffusion approximation; however, an additional term is

added to correct errors in the approximation. The solution marching scheme used will approximate this correction coefficient using the transport solution (high-order MOC in the radial direction and the NEM-P₃ in the axial direction) for the current iteration:

$$\hat{D}_s = \frac{J_s^{trans} + \frac{2D^+ D^-}{\Delta(D^+ + D^-)} (\bar{\phi}^+ - \bar{\phi}^-)}{(\bar{\phi}^+ + \bar{\phi}^-)}. \quad (3)$$

The iteration scheme in MPACT first assumes \hat{D}_s is zero and solves the full-core CMFD equations. Once the coarse mesh fluxes are obtained, they are projected onto the fine MOC mesh and a single MOC sweep is performed. The MOC equations solve for the current on the boundaries of the coarse mesh, which are used to estimate a new value for \hat{D}_s in the radial direction. The coarse cells are then homogenized and the 1D P₃ solution is used to compute the axial \hat{D}_s . This process is repeated until both the coarse and fine mesh solutions are converged.

Another key purpose of CMFD is to accelerate the solution to the eigenvalue problem. To understand this, it is easier to consider the matrix notation of the CMFD system. There are four major components to the matrix notation: the diffusion operator \mathbf{D} , the collision operator \mathbf{T} , the scattering operator \mathbf{S} , and the fission operator \mathbf{F} . These four terms are combined into the generalized CMFD eigenvalue problem

$$(\mathbf{D} + \mathbf{T} - \mathbf{S}) \phi = \frac{1}{k_{eff}} \mathbf{F} \phi \quad (4)$$

and is simplified to

$$= \frac{1}{k_{eff}} \mathbf{F} \phi, \quad (5)$$

where \mathbf{M} is the migration matrix.

The migration matrix for this case represents the movement of neutrons through space and energy. In

This manuscript has been authored by UT-Battelle, LLC, under Contract No. DE-AC0500OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for the United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

MPACT, it is ordered such that there is a dense block matrix that represents the transfer of neutrons between energy groups caused by scatter, as seen in Fig. 1.




Figure 1: 56 Energy Group Migration Matrix

The scattering matrix as well as the nonlinear correction term creates a structural asymmetry in the migration matrix. Along with the scattering matrix, the nonlinear correction term introduces an asymmetry in the matrix. It should also be noted that the fission matrix is both asymmetric and rank deficient.

The main focus of this work is to accelerate the convergence of this eigenvalue problem by looking at advanced solution schemes to solve for the dominant eigenvalue of this linear system.

CMFD EIGENVALUE METHODS

Several methods have been proposed over the years to solve this generalized eigenvalue problem. This work compares the current methodology (power iteration with a fixed shift, previous work using the generalized Davidson method [#] and the newly implemented multilevel in energy CMFD solver

One of the most common methods for solving the CMFD eigenvalue system is to cast the generalized eigenvalue problem into a standard eigenvalue problem:

$$\mathbf{M}^{-1}\mathbf{F}\phi = k_{eff}\phi . \quad (6)$$

Since the migration operator is nonsingular, it can be inverted and moved to the left hand side of the equation. Once we have a standard eigenvalue problem, power iteration can be applied to obtain the dominant eigenvalue. Typically storing the full factorization of the

migration operator is impractical for most problems; so instead, a linear system is solved each iteration:

$$\mathbf{M}\mathbf{y} = \mathbf{z} , \quad (7)$$

where $\mathbf{z} = \mathbf{F}\mathbf{f}^n$.

Although this iteration scheme is straightforward, the convergence of power iteration is highly dependent on the dominance ratio

$$r = \frac{k_2}{k_1} , \quad (8)$$

where k_1 and k_2 are the largest and second-largest eigenvalues of system, respectively. This convergence rate can be improved by introducing a shift (μ) to the eigenvalue system

$$(\mathbf{M} - \mu\mathbf{F})^{-1}\mathbf{F}\phi = \tilde{k}\phi , \quad (9)$$

where μ is an approximation to the inverse of the dominant eigenvalue. It is straightforward to show that the eigenvalue of the original system relates to the eigenvalue of the shifted system by

$$k_{eff} = \frac{\tilde{k}}{1 + \mu\tilde{k}} , \quad (10)$$

and that the dominance ratio of the shifted system can be reduced significantly. However, shifted power iteration does not come without risk. First, an incorrect choice of shift could alter the convergence to an eigenvalue that does not reflect the largest eigenvalue of the original system. Second, as the shift becomes closer to the true solution of the original system, the linear system becomes more and more difficult to solve. The obvious extreme is if the shift is exactly the inverse of the true eigenvalue of the system; $(\mathbf{M} - \mu\mathbf{F})$ is a singular matrix, so care must be taken when trying to solve shifted systems.

Traditionally in MPACT, a fixed shift of 2/3 is used. This was determined to be an acceptable balance between performance of the solver and ensuring that the system is not overshifted, which is unstable. Another approach that has been used in many other places is to have a variable shift based on the current eigenvalue iteration and potentially previous iteration. PARCS [5] uses a variable shift based on iteration history:

$$\mu^{(n)} = \mu_p^{(n)} = \max \left\{ \frac{1}{k^{(n-1)}} - C_1 \left| \frac{1}{k^{(n-1)}} - \frac{1}{k^{(n-2)}} \right| - C_0, \mu_{\min} \right\} , \quad (11)$$

where C_0 , C_1 , and μ_{\min} are constants (0.02, 10, and 1/3, respectively). This method tries to keep enough margin

from the actual eigenvalue to ensure stability of the method.

Thus far, all of the methods discussed here have been fixed-point methods; that is, the next estimate of the solution depends only on the estimate immediately preceding it. A set of alternatives to fixed-point iterations are subspace eigenvalue solvers in which information from several previous vectors is used to generate the next approximate solution. One such method that has been used successfully for the solution of both the discrete ordinates and SP_N forms of the transport equation is the generalized Davidson method [7,8]. The advantage of this method relative to other approaches (such as Arnoldi's method) is that it can be applied directly to a generalized eigenvalue problem without requiring conversion to a standard eigenvalue problem. It accomplishes this by using a Rayleigh-Ritz procedure, which solves the projected eigenvalue problem

$$\mathbf{V}^T \mathbf{M} \mathbf{V} \mathbf{y} = \lambda \mathbf{V}^T \mathbf{F} \mathbf{V} \mathbf{y} , \quad (13)$$

where \mathbf{V} contains a set of (typically orthogonal) basis vectors for the current subspace. For an appropriate selection of the subspace, the eigenvalues of the projection problem will closely approximate the original system and the vectors of $\mathbf{V} \mathbf{y}$ will approximate the eigenvector. New vectors are added to the subspace by applying a preconditioner to the current eigenvalue residual:

$$\mathbf{P} \mathbf{v} = -\mathbf{r}^{(n)} . \quad (14)$$

The Davidson method is extremely attractive because unlike the methods discussed previously, a linear system solve is not required. Instead, only a preconditioner application that approximates the solution of a linear system is needed. The choice of preconditioner is important for the Davidson solver to quickly converge. To avoid singularities in the preconditioner system, the preconditioner is chosen based on the migration operator \mathbf{M} rather than the shifted operator $(\mathbf{M} - \lambda^{(n)} \mathbf{F})$.

The final method considered is to recursively apply CMFD to a coarser and coarser energy grid. Several methods have looked at the effect of using 2 different energy grids [#, #, #, #] and more recently multiple energy grids have been considered [#, #, #, #]. The method considered here is a modified V-cycle iteration, Figure 2, with only energy restriction and projection operations. Each energy grid restriction is divided uniformly until two groups are obtained. The two-group grid is divided along the upscatter cutoff to ensure no upscattering is present in the coarsest grid. The modified V-cycle means

that the eigenvalue calculation only occurs on the coarsest grid. Every level in the iteration scheme solves a fixed source problem as a pre-sweep step except the coarsest, which performs a shifted power iteration to compute the eigenvalue. Then each level is solved again using the updated fission source as a post-sweep step. In this scheme the pre-sweep or post-sweep could be disabled to improve efficiency but that is not optimized in this study.

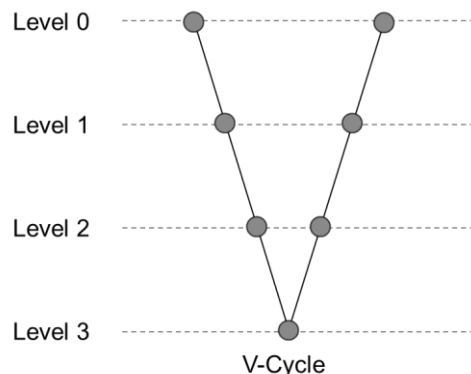


Figure 2: Standard V-Cycle Iteration Scheme

REACTOR MODELS AND RESULTS

The first case considered is the VERA Progression Problem [14] 5-2D. A 2D slice of a quarter-core pressurized-water reactor (PWR) is shown in Fig. 5.

	H	G	F	E	D	C	B	A
8	2.1 20	2.6 20	2.1 20	2.6 20	2.1 20	2.6 20	2.1 20	3.1 12
9	2.6 20	2.1 24	2.6 24	2.1 20	2.6 20	2.1 24	3.1 24	3.1 8
10	2.1 20	2.6 24	2.1 20	2.6 20	2.1 20	2.6 16	2.1 16	3.1 16
11	2.6 20	2.1 20	2.6 20	2.1 20	2.6 20	2.1 16	3.1 16	3.1 16
12	2.1 20	2.6 20	2.1 20	2.6 20	2.6 24	2.6 24	3.1 24	3.1 24
13	2.6 20	2.1 16	2.6 16	2.1 24	2.6 24	3.1 12	3.1 12	3.1 12
14	2.1 24	3.1 24	2.1 16	3.1 16	3.1 16	3.1 16	3.1 16	3.1 16
15	3.1 12	3.1 8	3.1 8	3.1 8	Enrichment Number of Pyrex Rods			

Figure 3: Core Layout for Problem 5 Cases

This case is run with a 16-processor partition. Table I shows the number of iterations, the total run time for MPACT, and the CMFD solve time. Five cases are shown: (1) a shifted power iteration with a constant shift of 1.5 using the GMRES solver from PETSc for a linear

solver with a block ILU preconditioner; (2) the generalized Davidson solver from Trilinos, preconditioned with algebraic multigrid (AMG) from ML; and (3-5) the multilevel cases with 2, 3, and 4 energy grids respectively.

Table 1: 2D Core Runtime

	Total Time	CMFD Time
Power Iteration	09:13.0	04:23.1
GD	05:21.6	00:39.9
ML-2	05:21.5	00:43.9
ML-3	05:26.4	00:49.8
ML-4	05:47.2	01:01.2

The second case considered is the full-height, quarter-core VERA Progression Problem 5a. This case is run on 464 processors with an 8 radial partition. Table II shows the problem 5a performance. Again, the five solution methods are evaluated at the ~500-core solution. As with the 2D cases, significant decreases in CMFD and total run times are observed in all cases.

Table 2: 3D Core Runtime

	Total Time	CMFD Time
Power Iteration	54:13.0	30:05.0
GD	20:57.4	04:54.7
ML-2	22:37.7	05:49.5
ML-3	21:17.9	05:40.6
ML-4	21:14.7	05:49.7

While it is important to understand the behavior of CMFD for beginning of life cases without feedback, it is crucial to access the performance with thermal-hydraulic feedback and isotopic depletion. Progression problem 9 is an ideal problem to demonstrate the effect of reactor feedback on the acceleration method chosen. Additionally, cycle 2 of Watts Bar Unit 1 is also simulated to determine if the performance changes with a reload core.

As can be observed in Figure 4 and 5, the performance at zero power is significantly improved over power iteration. As the cycle depletes the runtime improvements are decreased. Cycle 1 shows an approximate 20% improvement over power iteration over the course of the cycle and cycle 2 shows closer to 10% improvement. This suggests that while CMFD can get a significant speedup but cases with TH feedback but the benefits is significantly less.



Figure 4: Time per Statepoint Watts Bar Cycle 1

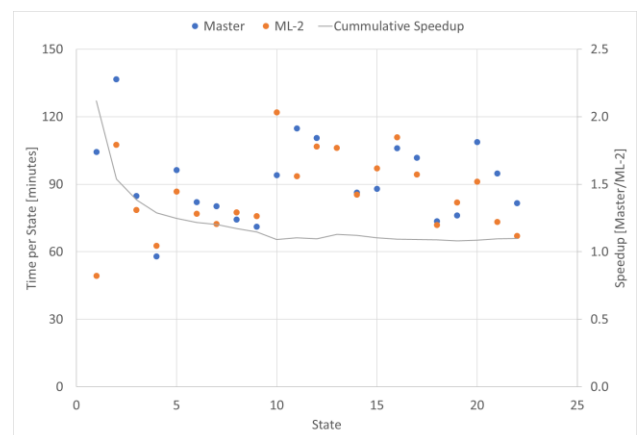


Figure 5: Time per Statepoint Watts Bar Cycle 1