CASL-U-2015-0157-000

Stabilization Methods for CMFD Acceleration

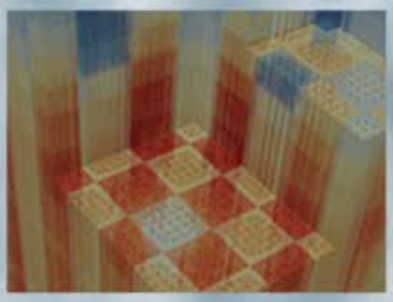

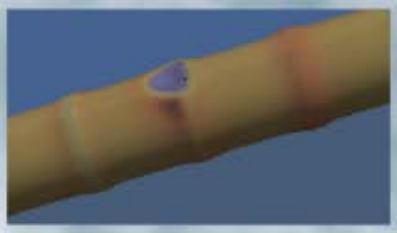

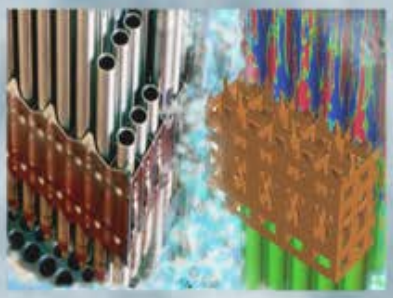

M. Jarrett, B. Kelley, B. Kochunas, T. Downar, E. Larsen
University of Michigan

April 19, 2015

U.S. DEPARTMENT OF ENERGY | Nuclear Energy

CASL-U-2015-0157-000

ANS MC2015 - Joint International Conference on Mathematics and Computation (M&C), Supercomputing in Nuclear Applications (SNA) and the Monte Carlo (MC) Method • Nashville, TN • April 19-23, 2015, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2015)

STABILIZATION METHODS FOR CMFD ACCELERATION

M. Jarrett, B. Kelley, B. Kochunas, T. Downar, E. Larsen
Department of Nuclear Engineering
University of Michigan
Ann Arbor, MI, 48109-2104 USA
jarremic@umich.edu

ABSTRACT

The Coarse Mesh Finite Difference (CMFD) method is one of the most widely used methods for accelerating the convergence of numerical transport solutions. However, in some situations, iterative methods using CMFD can become unstable and fail to converge. In this paper we evaluate several different modifications of the CMFD scheme that are known to stabilize the iterative method. We perform Fourier analysis on a linearized version of each scheme applied to an idealized (monoenergetic 1D infinite homogeneous medium planar S_N) problem to characterize the stability and rate of convergence. We also compare the effectiveness of the methods numerically by applying each to a 2D benchmark problem and a 2D/1D solution of a standard 3D benchmark problem using the MPACT code. We show that several methods are capable of stabilizing a 2D MOC solution, and examine the advantages and disadvantages of each. The numerical results show that there is potential for significant reductions in MPACT run time using improved CMFD stabilization methods.

Key Words: neutron transport, iteration methods, acceleration methods

1. INTRODUCTION

The CMFD method [1] is a very effective low-order acceleration scheme which reduces the number of iterations required to obtain a numerical solution to the higher-order neutron transport equation. Without any acceleration, source iteration methods are very slow to converge, and high fidelity two-and-three dimensional deterministic transport solutions would be computationally impractical. However, the fundamental CMFD scheme can sometimes be unstable [2], in particular if the coarse cell optical thickness becomes too large. To solve certain problems, the CMFD scheme must be modified to stabilize the iterative method.

In this paper we look at four distinct techniques, each of which is capable of resolving unstable CMFD cases:

1. Executing more than 1 inner transport sweep between each CMFD iteration
2. Under-relaxing the $\widehat{D}_{k+\frac{1}{2}}$ in the CMFD equations
3. Artificially increasing the diffusion coefficient $\widetilde{D}_{k+\frac{1}{2}}$
4. Refining the coarse mesh

The first two of these techniques are widely used in current Method of Characteristics (MOC) codes [3,4], while the authors are not aware of any current use of the third. The fourth

technique does not refer to any actual modification to the scheme (so there are no corresponding theoretical results), but rather is simply a minor change to the discretization of the standard CMFD method.

The principal motivation of this work is current interest in reducing the run time for MPACT, a 2D and 3D MOC neutron transport code being developed as part of the Consortium for Advanced Simulation of LWRs (CASL). [5] Because of its widespread success in previous applications, CMFD is used as the acceleration method in MPACT. However, both 2D and 3D benchmark problems tend to be unstable when a CMFD update is applied after each transport sweep. (Note: 3D problems in MPACT are typically solved with the 2D/1D method [6]). Therefore, multiple (usually two or three) transport sweeps must be performed before each CMFD update. Because transport sweeps are generally the most time-consuming portion of a solution, it is worthwhile to investigate alternate stabilization techniques to determine whether they are capable of achieving a faster rate of convergence.

In the next section an overview of the conventional CMFD method is first provided, along with the details of each of the stabilization techniques. Section 3 presents the Fourier analysis of each technique for a model problem and a detailed discussion of the results of the Fourier analysis. Finally in Section 4 the different stabilization techniques are tested for the C5G7 benchmark in 2D and 3D, and the Virtual Environment for Reactor Applications (VERA) benchmark problem 5 in 2D.

2. OVERVIEW OF CMFD AND THE STABILIZATION TECHNIQUES

2.1 Summary of CMFD Acceleration

The standard CMFD equations for an infinite-medium 1D planar S_N problem with isotropic scattering are presented in this section. First, we begin with the 1D planar S_N equations defining a transport sweep:

$$\frac{\mu_n}{h_j}\left(\psi_{n,j+1/2}-\psi_{n,j-1/2}\right)+\sigma_{t,j}\psi_{n,j}=\frac{\sigma_{s,j}}{2}\sum_{m=1}^{N}\psi_{m,j}w_m+\frac{q_j}{2} \qquad 1\le j\le p \qquad (1a)$$

$$\psi_{n,j}=\left(\frac{1+\alpha_{n,j}}{2}\right)\psi_{n,j+1/2}+\left(\frac{1-\alpha_{n,j}}{2}\right)\psi_{n,j-1/2} \qquad 1\le j\le p \qquad (1b)$$

$$\psi_{n,1/2}=\psi_n^b \quad , \quad \mu_n>0 \qquad (1c)$$

$$\psi_{n,p+1/2}=\psi_n^b \quad , \quad \mu_n<0 \qquad (1d)$$

Fine cell quantities are indexed by the subscript j, and include angular flux ψ, cross sections σ, thickness h, and source q. The number of fine cells per coarse cell is a fixed value p. The angular quadrature is defined by angles μ_n and weights w_n. The cell-average angular fluxes are built from the cell-edge values based on a weighting scheme defined by $\alpha_{n,j}$.

CASL-U-2015-0157-000

The CMFD equations for this problem are presented in Eq. (2). Subscripts denote spatial (or angular) indices and superscripts denote iterate indices. The letters *j* and *k* refer to fine and coarse mesh indices, respectively. Lower and upper case letters refer to fine and coarse cell quantities, respectively. The "half" iteration refers to the solution after a transport sweep but before a CMFD update. Coarse cell quantities include surface current J, scalar flux Φ, source Q, homogenized absorption cross section Σ_a, and thickness Δ.

$$J_{k+\frac{1}{2}}^{(l+1)} - J_{k-\frac{1}{2}}^{(l+1)} + \Sigma_{a,k}^{\left(l+\frac{1}{2}\right)} \Phi_k^{(l+1)} \Delta_k = Q_k \Delta_k \qquad 1 \le k \le K \qquad \text{(2a)}$$

$$J_{k+\frac{1}{2}}^{(l+1)} = -\widetilde{D}_{k+\frac{1}{2}}^{(l+1)} \left(\Phi_{k+1}^{(l+1)} - \Phi_k^{(l+1)}\right) + \widehat{D}_{k+\frac{1}{2}}^{(l+1)} \left(\Phi_{k+1}^{(l+1)} + \Phi_k^{(l+1)}\right) \qquad 1 \le k < K \qquad \text{(2b)}$$

$$2J_{\frac{1}{2}}^{+} = J_{\frac{1}{2}}^{\left(l+\frac{1}{2}\right)} + \left(\frac{\Sigma_{n=1}^{N} |\mu_n| \psi_{n,1}^{\left(l+\frac{1}{2}\right)} w_n}{\Sigma_{n=1}^{N} \Psi_{n,1}^{\left(l+\frac{1}{2}\right)} w_n} \right) \Phi_1^{\left(l+\frac{1}{2}\right)} \qquad k = 1 \qquad \text{(2d)}$$

$$2J_{K+\frac{1}{2}}^{+} = -J_{K+\frac{1}{2}}^{\left(l+\frac{1}{2}\right)} + \left(\frac{\Sigma_{n=1}^{N} |\mu_n| \psi_{n,J+\frac{1}{2}}^{\left(l+\frac{1}{2}\right)} w_n}{\Sigma_{n=1}^{N} \Psi_{n,K}^{\left(l+\frac{1}{2}\right)} w_n} \right) \Phi_K^{\left(l+\frac{1}{2}\right)} \qquad k = K \qquad \text{(2e)}$$

$$\phi_j^{(l+1)} = \frac{\Phi_k^{(l+1)}}{\Phi_k^{\left(l+\frac{1}{2}\right)}} \phi_j^{\left(l+\frac{1}{2}\right)} \qquad j \in k \qquad \text{(2f)}$$

The coarse mesh scalar fluxes $\Phi_k^{(l+1)}$ are obtained from solving Eqs. (1b) with $\widehat{D}_{k+\frac{1}{2}}^{\left(l+\frac{1}{2}\right)}$, which is obtained from Eq. (2b) at the $(l+\frac{1}{2})$ iterate. Since the $\widehat{D}_{k+\frac{1}{2}}^{(l+1)}$ terms are generally small, they are lagged from the half-iteration point in order to solve the equations. The ratio of the coarse cell flux is used to update the fine cell scalar fluxes. $\widetilde{D}_{k+\frac{1}{2}}^{(l+1)}$ is the standard P_1 diffusion coefficient at the surface between coarse cell *k* and *k*+1.

2.2 Techniques for Stabilizing CMFD

The first technique of performing multiple sweeps is straightforward, and this has been the traditional approach to stabilizing the CMFD acceleration in MPACT. [3] Techniques 2 and 3 are simple modifications to the standard CMFD equations. The last technique (4) is to simply change the definition of a coarse cell to reduce the coarse cell optical thickness. The feasibility of this last approach depends largely on the geometry and meshing infrastructure of the code.

2.2.1 $\widehat{D}$ Under-relaxation

Under-relaxation of the balance correction term $\widehat{D}$ is a technique that has been known for a few decades but has never appeared extensively in literature. It is investigated quite thoroughly for a few benchmark problems in another paper at this conference. [4] Our implementation was based upon correspondence with the authors of this paper. The only modification is that we redefine $\widehat{D}$:

$$\widehat{D}_{k+\frac{1}{2}}^{(l+1)} = \theta\, \widehat{D}_{k+\frac{1}{2}}^{(l+\frac{1}{2})} + (1-\theta)\, \widehat{D}_{k+\frac{1}{2}}^{(l)} \qquad 1 \le k < K \quad \text{(3a)}$$

where $0 < \theta \le 1$ is the relaxation coefficient. For a problem that is stabilized by this technique, the iteration is only stable for $0 < \theta \le \theta_{max} < 1$. For θ close to 0, the rate of convergence is very slow and approaches source iteration in effect.

2.2.2 Artificial Grid Diffusion (AGD)

Artificially increasing the diffusion coefficient, which may be less familiar to some, is derived from the linearization of the partial-current CMFD (pCMFD) method, which is unconditionally stable with a step-characteristics discretization. [7] By increasing the diffusion coefficient, we attempt to reconcile the difference between the regular CMFD method and the pCMFD method. The diffusion coefficient in the balance Eq. (2b) is redefined:

$$J_{k+\frac{1}{2}}^{(l+1)} = -\left(\widetilde{D}_{k+\frac{1}{2}}^{(l+1)} + \eta\Delta_k\right)\left(\Phi_{k+1}^{(l+1)} - \Phi_k^{(l+1)}\right) + \widehat{D}_{k+\frac{1}{2}}^{(l+1)}\left(\Phi_{k+1}^{(l+1)} + \Phi_k^{(l+1)}\right) \qquad 1 \le k < K \quad \text{(3b)}$$

The quantity η is a dimensionless constant. The theoretically suggested value is $\eta = \frac{1}{4}$, which comes from the linearization of pCMFD. However, it is not required to be exactly $\frac{1}{4}$. If η is increased, the spectral radius also increases and the CMFD solution becomes more flat and thus less useful as an acceleration of the transport solution. As η is decreased, the rate of convergence can increase. However, it reduces the region of stability for the iterative method and in the limit the iteration scheme will exhibit the same instabilities as traditional CMFD. Consequently, we may posit that there should exist some optimal η for a given problem and if a nearly optimal η applies to a large class of problems then this stabilization technique becomes more attractive.

2.2.3 Coarse mesh refinement

Refining the coarse mesh is an effective way to improve the stability of CMFD. Choosing a finer radial discretization allows the CMFD to be more effective in approximating/accelerating toward the fine mesh transport solution. When solving the neutron transport equation on the regular lattice of a typical LWR, the obvious choice is to line up the coarse grid with pin cell boundaries – which is what is done in MPACT. To refine the coarse mesh further, the most practical scheme is quarter-pin CMFD cells – this preserves the highly desirable rectilinear nature of the coarse mesh and does not require any changes to the fine mesh of the problem (the pin cells are typically at least divided into octants azimuthally). This will increase the memory and run time for the solution of the CMFD linear system, but this usually makes up only a small fraction of the overall run time which is dominated by the 2D MOC sweeps.

3. NUMERICAL FOURIER ANALYSIS

The theoretical analysis here very closely follows the procedure outlined by Larsen and Kelley [8]. Due to space constraints, most of the steps are omitted, but the interested reader should consult the literature to see the complete derivation of this system.

Since CMFD is a nonlinear method, it must first be linearized to facilitate spectral analysis. We expect the behavior of CMFD to be very similar to its linearized form, and this expectation is

verified experimentally by Larsen and Kelley [8]. To obtain a set of equations for the growth factors of the iterative methods, it is necessary to apply them to an idealized problem. We applied these methods to the monoenergetic infinite homogeneous medium fixed-source 1D planar S_N equations (Eqs. 1) with a step-characteristics discretization. The Fourier ansatz is:

$$q_j = 0 \tag{4a}$$

$$\psi_{n,j'-\frac{1}{2}}^{(l+\frac{1}{2})} = \omega^l d_{n,j} e^{i\lambda\Sigma_t x_k} \tag{4b}$$

$$\psi_{n,j'}^{(l+\frac{1}{2})} = \omega^l a_{n,j} e^{i\lambda\Sigma_t x_k} \tag{4c}$$

$$\phi_{j'}^{(l)} = \omega^l g_{n,j} e^{i\lambda\Sigma_t x_k} \tag{4d}$$

$$F_{0,k}^{(l+1)} = \omega^l F e^{i\lambda\Sigma_t x_k} \tag{4e}$$

The first subscript n denotes quadrature index and the second subscript j or k denotes fine or coarse mesh index. The superscript l denotes iteration index. ψ and ϕ are the fine cell angular and scalar fluxes, respectively. q is the source, which is assumed to be zero to simplify the analysis. $F_{0,k}$ is the additive correction to the scalar fluxes for fine cells within coarse cell k. ω is the growth mode, which depends on the error mode λ. After substituting this ansatz into the 1D planar S_N equations (Eq. 1) and a linearized form of Eq. 2, we have a system of equations for the coefficients $a_{n,j}$, $d_{n,j}$, g_j, and F. The coefficients of the matrix are dependent on both the frequency of the error mode λ and the coarse cell optical thickness $\Sigma_t ph$. By searching over the full frequency domain of λ while maintaining a fixed number of fine cells per coarse cell p and fine cell optical thickness $\Sigma_t h$, we can determine the theoretical spectral radius of the iterative method for that particular CMFD mesh. Because the problem is periodic, only a finite domain of $\lambda \in [0, \pi]$ must be searched. Due to the discrete nature of the solution, higher frequency error modes cannot exist on the grid and will be represented as a lower frequency which falls in this range. With some linear algebra the system can be reduced to an eigenvalue problem for growth factors ω_j. We compute these eigenvalues numerically as discrete functions of λ, and use a bisection method to calculate the spectral radius

$$\rho = \sup_{\lambda} \sup_{1 \le j \le p} |\omega_j(\lambda)| \tag{5}$$

The spectral radius is computed over a logarithmically-spaced set of coarse cell optical thicknesses to obtain a set of values that accurately describes the true spectral radius as a function of the mesh size.

3.1 Numerical Results

Fig. 1 shows the results of applying this analysis to the standard linearized CMFD equations. For lower scattering ratios c, the scheme is stable for all coarse cell sizes. For scattering ratios closer to unity, the scheme becomes unstable for cells of an intermediate optical thickness, beginning at a thickness of approximately 2 mean-free-paths.

Fig. 2 shows the effect of executing multiple sweeps between each CMFD update. Because the definition of an iteration includes many transport sweeps, the spectral radius decreases. The thin- and thick-cell asymptotic values decrease, as well as the wide peak in the region of intermediate cell sizes. The initial size of the peak is very large for c close to unity, and increases

with c. For any scattering ratio $c < 1$, there is some number of sweeps that will reduce the magnitude of the peak enough such that the iteration will be stable.

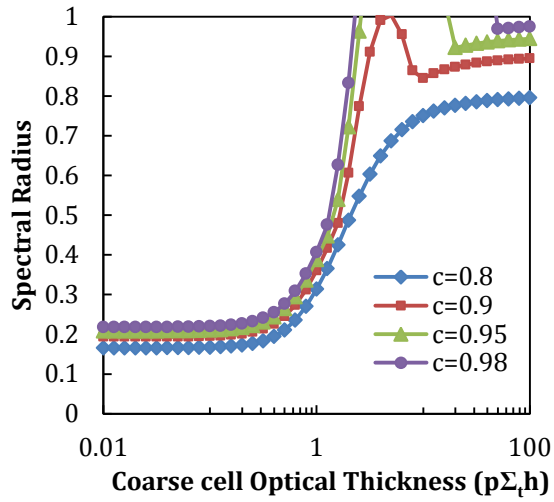

Figure 1: Spectral Radius for Standard CMFD, $p = 4$

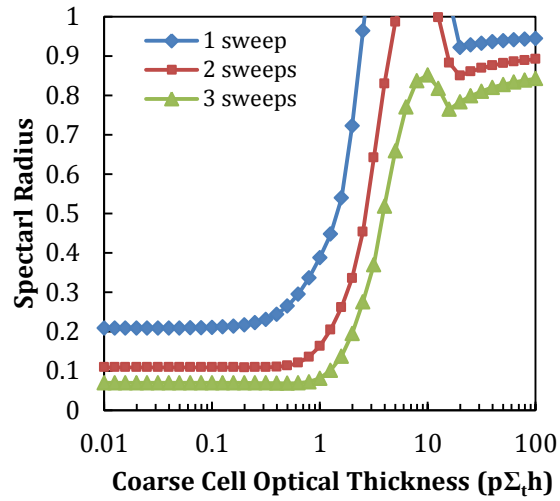

Figure 2: Spectral Radius for Technique 1, $p = 4$ and $c = 0.95$

Fig. 3 shows the smoothing effect of CMFD under-relaxation on the spectral radii. Under-relaxation reduces the magnitude of the spectral radius peak, but at the same time it also slows the convergence for optically thin cells. Much like multiple sweeps, there appears to be some level of under-relaxation $0 < \theta < 1$ that can reduce the peak enough to stabilize the iteration for any c. However, for c close to unity, the required θ may be so small that there will be little or no improvement compared to source iteration.

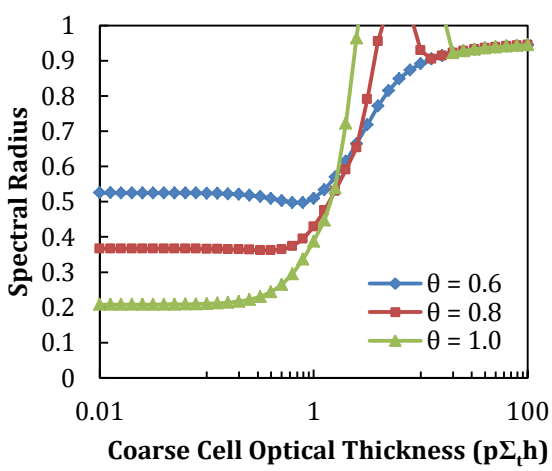

Figure 3: Spectral Radius for Technique 2, $p = 4$ and $c = 0.95$

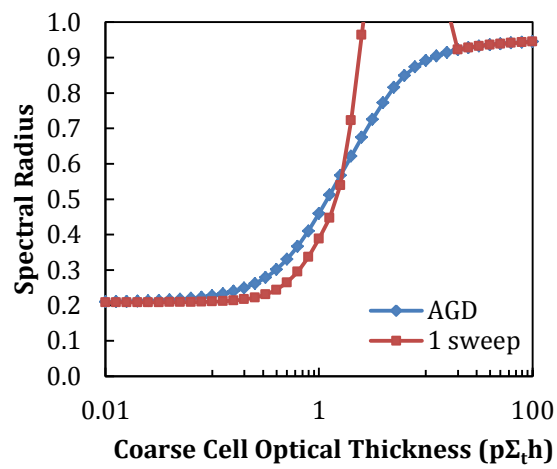

Figure 4: Spectral Radius for Technique 3, $p = 4$ and $c = 0.95$

Fig. 4 shows the stabilizing effect of the AGD modification (η = 0.25). The unstable peak in spectral radius is almost completely eliminated. The convergence rate is slightly slower for optically thin cells. Fig. 5 shows how the AGD technique mitigates the negative growth modes, which are responsible for the instabilities, while slightly increasing the positive growth factors.

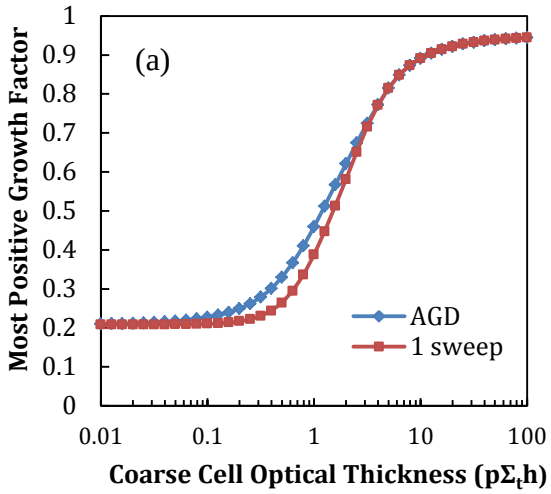

Figure 5a: Largest Magnitude Positive Growth Factors for p = 4 and c = 0.95

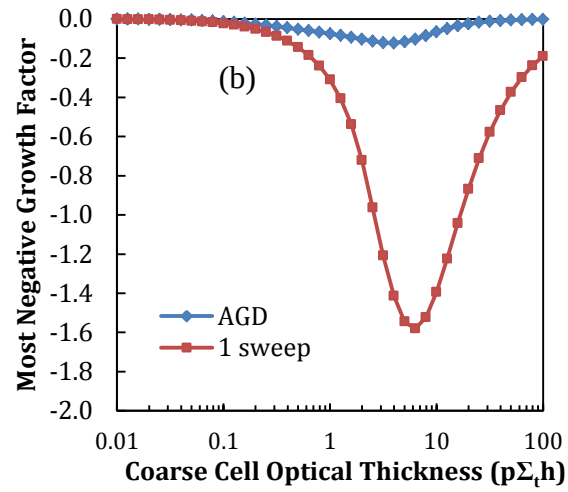

Figure 5b: Largest Magnitude Negative Growth Factors for p = 4 and c = 0.95

3.2 Discussion

3.2.1 Multiple Transport Sweeps

Executing multiple sweeps inside each outer iteration improves the effectiveness (i.e. reduces the spectral radius) of each outer iteration. For optically thin cells, the improvement does not make up for increased cost, so it is more effective to perform the CMFD update after each transport sweep. For cells greater than approximately one mean-free-path thick, it is more effective to perform multiple transport sweeps between each CMFD update. This technique reduces the magnitude of the large peak for intermediate-to-large optical thickness CMFD cells, and increases the range of stability for the iterative method. Note, however, that the method does not become unconditionally stable for scattering ratios arbitrarily close to unity.

3.2.2 $\widehat{D}$ Under-relaxation

Under-relaxing the $\widehat{D}_{k+\frac{1}{2}}^{(l+1)}$ terms increases the positive growth factors while decreasing the magnitude of the negative growth factors. This reduces the magnitude of the peak in the unstable region, but at the cost of pushing the positive growth factors higher and thus slowing down convergence in those regions.

3.2.3 Artificial Grid Diffusion (AGD)

Artificially increasing the diffusion coefficient virtually eliminates the negative growth factors, making the iteration unconditionally stable. However, it also makes the CMFD solution

more diffusive (i.e. a poorer approximation to the transport solution), which slightly reduces its effectiveness. This decrease in effectiveness shows up as a shift towards smaller cells in the transition region from rapid convergence to slow convergence. For optically thin cells, CMFD is highly effective and the positive growth factors asymptotically approach a minimum. For optically thick cells, CMFD is ineffective and the positive growth factors approach that of a regular source iteration (the scattering ratio c). The transition from high to low effectiveness occurs for cells that are between 1 and 10 mean-free-paths thick, approximately. Increasing the diffusion coefficient shifts this transition toward thinner cells, so that for any cells between about 0.1 and 5 mean-free-paths thick (most practical cases), the largest positive growth factor is slightly larger.

3.2.4 Coarse Mesh Refinement

In all cases examined here, the spectral radius is strictly increasing from the thinnest coarse cells up to the peak from negative growth factors. Since this peak occurs for very optically thick cells, in most applications an unstable iterative method is on the increasing (with thickness) side of that peak. Decreasing the coarse cell size effectively moves the iterative method to the left on these curves, potentially stabilizing unstable growth modes, while also increasing the rate of convergence of the modes which were already stable.

4. EXPERIMENTAL RESULTS

The stabilization techniques described in this paper were evaluated using the C5G7 Benchmark problem which is a well-established transport benchmark problem consisting of 17x17 UO_2 and MOX fuel lattices [9]. A 2D version of the benchmark was analyzed on a single processor and the 3D C5G7 problems (unrodded and rodded configurations) were decomposed spatially into 18 axial planes and run on 18 processors.

Since the CMFD solution is relatively inexpensive, the total number of transport sweeps required to reach the converged solution is an appropriate and simple metric for the numerical effectiveness of a stabilization technique. It should also be noted that the respective computational kernels for 2D MOC and for the solution of the CMFD linear system are not fully optimized in MPACT. Therefore, the benefit of trading off MOC sweeps for CMFD solves may not necessarily be reflective of other codes.

4.1 2D C5G7 Results

The 2D C5G7 results are given in Table I, and as indicated the use of multiple sweeps was faster than the under-relaxation and AGD for 2D problems. Under-relaxation and AGD both required more overall transport sweeps than the first technique. This makes sense since both of these techniques achieve stabilization at the cost of convergence rate. While multiple sweeps is less effective for optically thin coarse cells, CMFD is most effective for thin coarse cells, so they are unlikely to dictate the convergence of the problem. The MOC time improvement from reducing the number of inner sweeps from 2 to 1 was not very significant. Refining the coarse grid reduced the number of transport sweeps from 22 to 10, but the MOC time was only reduced by approximately 25%. As a result, the increase in CMFD solution time and initialization time

mostly counteracted the MOC savings. Fig. 6 shows the convergence of the norm of the fission source error for the runs summarized in Table I. The refined coarse mesh significantly improves the convergence rate. Under-relaxation and AGD decrease the asymptotic convergence rate compared to multiple sweeps, but their initial convergence is more rapid because the CMFD updates are being applied more frequently. Techniques 1 and 2 had essentially identical run times, while 3 and 4 were slightly slower and faster, respectively.

Table I. Run Time Comparisons for C5G7 2D Benchmark Problem

Technique	Θ	η	Inners	Outers	Run Time (s)	Speedup	MOC (s)	CMFD (s)	k-eff
1	1.00	0.00	2	11	588	1.00	543	42	1.186112
2	0.72	0.00	1	18	589	1.00	515	72	1.186113
3	1.00	0.25	1	19	614	0.96	542	69	1.186138
4	1.00	0.00	1	10	562	1.05	413	138	1.186301

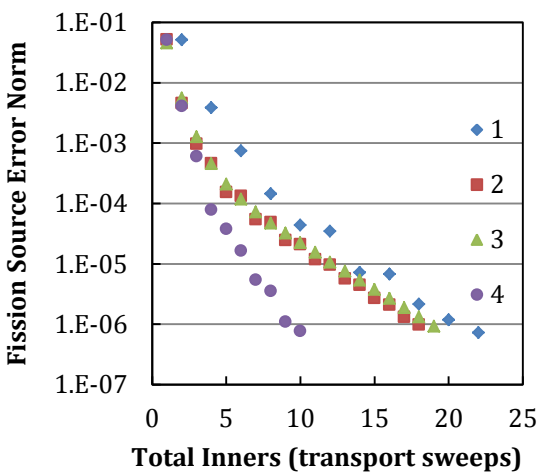

Figure 6: Fission Source Convergence C5G7 2D

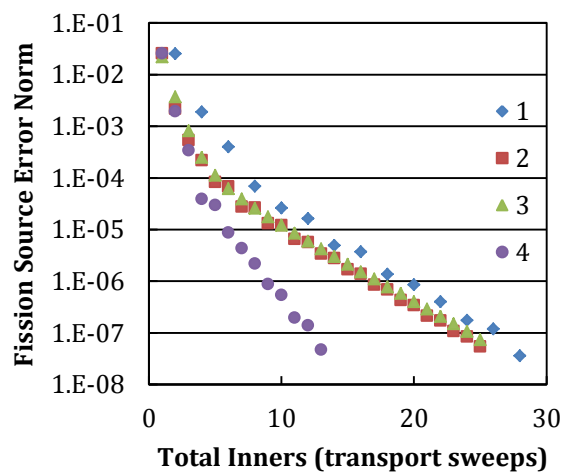

Figure 7: Fission Source Convergence C5G7 3D (unrodded)

4.2 3D C5G7 Results

For 3D problems, the 2D MOC solution is coupled to a 1D axial solution, which is only updated once every outer iteration. This increases the penalty of executing multiple sweeps per outer iteration because the axial solution needs a certain number of iterations to converge. This implies that if the code performs multiple 2D MOC sweeps per axial solution update, it may end up performing many more 2D MOC sweeps than are really necessary to sufficiently converge to a solution. If the only one inner 2D MOC sweep is performed per outer iteration, the code will only have to perform the minimum number of 2D sweeps necessary to converge the 2D/1D solution. Thus, there is a benefit to performing only one inner sweep per outer iteration. This results in an advantage to techniques 2, 3, and 4 compared to technique 1 for 2D/1D problems.

Note that the quarter-pin cell solution provides a slightly different k-eff because the axial part of the 2D/1D solution is also refined, which provides more spatial shape to the axial transverse leakage and therefore an improved overall 3D solution.

Table II. Run Time Comparisons for C5G7 3D Benchmark Problems

Case	Technique	Inners	Outers	Run Time (min)	Speedup	MOC (min)	CMFD (min)	keff
unrodded	1	2	14	87	1.00	82	1	1.142310
unrodded	2	1	25	81	1.07	74	3	1.142311
unrodded	3	1	25	79	1.09	73	3	1.142311
unrodded	4	1	13	57	1.53	40	5	1.142142
rodded A	1	2	14	77	1.00	71	1	1.127395
rodded A	2	1	25	79	0.97	73	2	1.127395
rodded A	3	1	25	81	0.95	73	4	1.127395
rodded A	4	1	13	59	1.32	40	12	1.127230
rodded B	1	2	14	78	1.00	72	1	1.076568
rodded B	2	1	26	88	0.89	79	4	1.076568
rodded B	3	1	26	83	0.94	76	3	1.076568
rodded B	4	1	13	60	1.30	41	12	1.076421

Fig. 7 shows the convergence of the norm of the fission source error for the 3D C5G7 unrodded case. The 2 inners case converges more slowly, presumably because of the lower rate of axial solution updates. As a result, the performance of technique 1 is poorer than the other three techniques. Under-relaxation and AGD converge very similarly in all three cases. The quarter-pin coarse mesh case converges at a higher asymptotic rate. Although techniques 2 and 3 slightly reduced the total number of transport sweeps, the total MOC time was not reduced in all cases. This is likely due to the fact that the last inner of any outer iteration is slightly more expensive because of the need to store certain quantities, such as the surface currents for CMFD. Technique 4 reduced the total number of inners by about 50%, which produced a significant speedup, even with the increased CMFD time.

4.3 VERA Problem 5 2D results

The stabilization techniques were also tested on VERA benchmark problem 5 which, unlike C5G7, is an accurately and explicitly modeled PWR geometry. [10] The problem is a 2D slice of the Watts Bar Nuclear 1 startup core, with Westinghouse 17x17-type fuel assemblies at beginning-of-life and hot zero power isothermal conditions. Technique 4 was not tested on this problem because the mesh is generated automatically by MPACT, so refining the coarse mesh would be a significant undertaking, and would most likely be fruitless considering the C5G7 2D results showed only a small improvement by using a finer coarse mesh. In fact, it is apparent that all of the single-sweep stabilization techniques are more effective for 2D/1D problems than they are for simpler 2D problems. The timing results are given in Table III; while techniques 2 (unrodded) and 3 (rodded) reduced the overall number of transport sweeps, the MOC time was still lower for technique 1, and as a result techniques 2 and 3 produced a significant slowdown.

Note that this benchmark problem is stable without any of the techniques being applied (just 1 inner transport sweep) and converges in 14 outer iterations. However, the techniques were still applied to study the improvement (or detriment) of the convergence rate. For technique 2, the under-relaxation factor was 0.9.

Table III. Run Time Comparisons for VERA Benchmark Problem 5-2D

case	Technique	Inners	Outers	Run Time (hr)	Speedup	MOC (hr)	CMFD (hr)	keff
unrodded	1	2	8	12.0	1.00	6.6	4.8	1.003059
unrodded	2	1	16	16.4	0.73	7.3	8.4	1.003059
unrodded	3	1	15	17.6	0.68	6.8	10.0	1.003059
rodded	1	2	8	12.7	1.00	6.4	5.6	0.990783
rodded	2	1	14	15.3	0.83	6.9	7.8	0.990783
rodded	3	1	19	20.2	0.63	8.0	11.5	0.990783

5. CONCLUSIONS AND FUTURE WORK

This paper evaluated several different possible techniques for stabilizing the CMFD method to accelerate MOC solutions of the neutron transport equation. Fourier analysis was performed on a linearized version of each scheme and applied to an idealized (1D infinite homogeneous medium planar S_N) problem to characterize the stability and rate of convergence. The effectiveness of the techniques was then evaluated numerically by applying each to a 2D benchmark problem and a 2D/1D solution of a standard 3D benchmark problem using the MPACT code.

All four CMFD stabilization techniques demonstrated the ability to numerically stabilize CMFD consistent with the Fourier analysis. For the 2D C5G7 problem, the multiple sweeps technique was effective because the 2D/1D penalty of multiple sweeps only applies to 3D problems. However, for the 3D C5G7 benchmark problems, the penalty was significant and technique 4 performed significantly better than technique 1. The results were mixed for techniques 2 and 3.

While under-relaxation and AGD performed better than the traditional technique for all of the 3D C5G7 problems, it is not clear from the fission source convergence plots that they would generally converge more quickly. The asymptotic convergence for these problems is distinctly less steep than it is for the traditional technique. The principle reason these techniques perform better is that the traditional technique struggles early in the iteration because the coupling between the 2D and 1D solutions is not yet well converged at this point.

Measuring performance by the total number of transport sweeps, techniques 2 and 3 were comparable to technique 1 for VERA Problem 5-2D. However, they were about 20-35% slower in overall run time because of the increased CMFD time. Even when the overall number of transport sweeps was reduced, the time required for MOC was not.

For the C5G7 problems, the quarter-pin cell coarse mesh is the most effective technique for stabilizing and improving convergence of a 2D/1D solution in MPACT. However, the inputs to test this technique were specifically designed for this work without changing the code in any way. To modify the code to automatically generate the quarter-pin cell CMFD mesh would require significant reworking of the input processing and auto-meshing routines. Further numerical testing of all of the techniques in MPACT is necessary to fully characterize the advantages and disadvantages of each and eventually come to a defendable conclusion about which technique is most efficient. This will likely involve developing a figure-of-merit which considers the run time, memory requirements, relative stability, and other performance metrics of the techniques with an appropriate importance assigned to each metric. Such a thorough analysis should allow for an empirical determination of which technique requires the shortest run time across all relevant problems, or at least a subset of the most relevant problems.

6. 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 U.S. Department of Energy Contract No. DE-AC05-00OR22725. This material is also based upon work supported under a Department of Energy Nuclear Energy University Programs Graduate Fellowship.

7. REFERENCES

1. K.S. Smith, “Nodal Method Storage Reduction by Nonlinear Iteration,” Trans. Am. Nucl. Soc., 44, 265 (1983).
2. .Z. Cho, C.J. Park, “A Comparison of Coarse Mesh Rebalance and Coarse Mesh Finite Difference Accelerations for the Neutron Transport Calculations,” Int. Conf. Nuclear Mathematical and Computational Sciences (M&C 2003), Gatlinburg, USA, April 6-11, 2003, CD-ROM, American Nuclear Society (2003).
3. J.Y. Cho, H.G. Joo, K.S. Kim, and S.Q. Zee, "Cell Based CMFD Formulation for Acceleration of Whole-core Method of Characteristics Calculations", J. of Kor. Nucl. Soc. 34, 3, pp. 250-258, (2002).
4. L. Li, K. Smith, and B. Forget, “Techniques for Stabilizing CMFD in MOC”, *Proc. M&C 2015*,Nashville, TN, USA. April 19-23 (2015).
5. B. Kochunas, B.S. Collins, D. Jabaay, T.J. Downar, W.R. Martin, “Overview of Development and Design of MPACT: Michigan Parallel Characteristics Transport Code,” *Proc. M&C 2013*, Sun Valley, ID, USA. May 5-9 (2013).
6. B.S. Collins, S.G. Stimpson, A.T. Godfrey, B. Kochunas, T.J. Downar, “Assessment of the 2D/1D Implementation in MPACT,” *Proc. PHYSOR 2014*, Kyoto, Japan. Sep 28 – Oct 3 (2014). (Paper 1126966).
7. N.Z. Cho, G.S. Lee, C.J. Park, “Partial Current-Based CMFD Acceleration of the 2D/1D Fusion method for 3D Whole-Core Transport Calculations,” Trans. Am. Nucl. Soc., 88, 594, (2003).
8. E.W. Larsen and B.W. Kelley, “The Relationship Between the CMFD and the Coarse-Mesh DSA Methods,” *Nuclear Science & Engineering*, **Volume 178**, pp.1-15 (September 2014).
9. https://www.oecd-nea.org/science/docs/2003/nsc-doc2003-16.pdf
10. A. Godfrey, “VERA Core Physics Benchmark Progression Problem Specifications,” CASL Technical Report: CASL-U-2012-0131-002. Revision 2, pp.78-86. March 29, 2013.