M&C 2017 - International Conference on Mathematics & Computational Methods Applied to Nuclear Science & Engineering, Jeju, Korea, April 16-20, 2017, on USB (2017)

Multi-Level Acceleration of Scattering-Source Iterations with Application to Electron Transport

Clif Drumm and Wesley Fan

Radiation Effects Theory Dept.
Sandia National Laboratories[a], PO Box 5800, Albuquerque, NM 87185-1179
crdrumm@sandia.gov

Abstract – *Acceleration/preconditioning strategies available in the SCEPTRE radiation transport code are described. A flexible Transport Synthetic Acceleration (TSA) algorithm that uses a low-order discrete-ordinates (S_N) or spherical-harmonics (P_N) solve to accelerate convergence of a high-order S_N source-iteration (SI) solve is described. The convergence of the low-order solves can be further accelerated by applying off-the-shelf incomplete-factorization or algebraic-multigrid methods. Also available is an algorithm that uses a GMRES iterative method rather than SI for convergence, using a parallel sweep-based algorithm to build up a Krylov subspace. TSA has been applied as a preconditioner to accelerate the convergence of the GMRES iterations. The methods are applied to several problems involving electron transport and problems with artificial cross sections with large scattering ratio. The methods will be compared and evaluated by considering material discontinuities and scattering anisotropy. Observed accelerations obtained are highly problem dependent, but speedup factors around ten have been observed in typical applications.*

I. INTRODUCTION

Electron transport is characterized by a large number of collisions before removal from the scattering source, which makes the source iterations in a sweeps-based solution method slowly convergent. The SCEPTRE radiation transport code [1] has several alternative solvers that do converge rapidly for electron transport in many circumstances. In these alternative solvers, a discretized linear system is constructed involving both the spatial and angular variables, including the scattering source, which is solved in parallel by a Krylov iterative method. The convergence of the Krylov methods depends upon the condition number of the system matrix and the clustering of the eigenvalues, rather than upon the scattering ratio, so that rapid convergence may be obtained even when highly-scattering media is present in the problem. A main drawback of these methods is that they can be memory intensive.

In this work, a method is described that combines the two solution approaches, by using a coarse-level (low-order-in-angle) space/angle solve to accelerate a fine-level (high-S_N-order) source-iteration solve. For the coarse-level solves, it is possible to attain further acceleration by applying off-the-shelf Incomplete-Factorization (IF) [2] or Algebraic Multi-Grid (AMG) [3] algorithms. This work is basically generalized Transport Synthetic Acceleration (TSA) [4] with great flexibility in performing the coarse level solves.

Additionally, capability has recently been added to SCEPTRE to use a Generalized Minimum Residual (GMRES) algorithm for convergence rather than source iteration. Krylov iterative methods have been effectively applied to radiation transport problems for some time [5-7]. In this approach, sweep solves are used to build up a Krylov subspace that is used to minimize a residual and converge to a solution. This method is shown to be effective for electron-transport and other applications with large scattering ratio. TSA has been applied as a preconditioner to accelerate the convergence of the Krylov iterations. The Krylov algorithm implemented in SCEPTRE will be described below.

The next section will provide more details of the methods, and then results will be provided for modeling electron transport in a twisted-pair electrical cable, and a three-material-block cylinder with various combinations of uniform and non-uniform cross sections. Observed accelerations obtained are highly problem dependent, but speedup factors around ten have been observed in typical applications.

II. BACKGROUND

Acceleration methods such as Diffusion Synthetic Acceleration (DSA) and TSA have been very effective at reducing the number of iterations and/or solve time for a wide range of applications [8-11]. However, partially consistent DSA may perform poorly for certain applications, and fully-consistent DSA and TSA may effectively reduce the iteration count, while shifting much of the work to the coarse-level solves. The effectiveness of acceleration methods tends to be very problem dependent. Generally, for problems with highly-scattering regions, some type of acceleration is effective, but practical problems often have highly-scattering regions and streaming regions, so some flexibility is desired in applying a preconditioner.

The SCEPTRE radiation transport code has a unique ability of allowing for any number of different types of solvers to be defined within a single transport calculation [1]. The primary application of SCEPTRE is coupled photon/electron transport, so this flexibility was needed since the convergence properties for problems involving neutral particles is so different from that involving charged parti-

cles, and even among different energy groups, convergence behavior may be quite variable. As noted previously, SCEPTRE includes a sweeps-based source-iteration solver and also a number of solvers that use a Krylov iterative method to solve the space/angle dependence simultaneously.

A wide variety of space/angle solvers are available in SCEPTRE. Available space/angle solvers include those based on the first-order form of the transport equation using Discontinuous Finite Elements (DFE), for both the spherical-harmonics (P_N) and discrete ordinates (S_N) treatment of the angular variable. A DFE spatial discretization always results in a non-SPD linear system, so GMRES is used for these solvers. Also available are several solvers that result in a Symmetric Positive Definite (SPD) linear system, so that the highly-efficient Conjugate Gradients (CG) algorithm may be used. These methods include the Self-Adjoint Angular Flux (SAAF) method and Least-Squares (LS) method, using Continuous Finite Elements (CFE), again for both P_N and S_N [12-13]. The use of spatial DFE results in a non-symmetric linear system, regardless of the form of the transport equation used, so only CFE spatial approximation is available for use with the SAAF and LS solvers.

As noted earlier, due to huge memory requirements, the space/angle solvers are not generally practical except for fairly small problems. However, these methods are very effective as coarse-level accelerators for a source-iteration solver. Because the space/angle linear systems are built using Trilinos [14] data structures, the rich assortment of preconditioners included in the Trilinos project, including Algebraic Multi-Grid (AMG) and Incomplete Factorization (IF), are accessible.

Mapping from the fine-level problem to the coarse-level problem is basically Angular Multi-Level (AML) optionally combined with an algebraic preconditioner, either IF [2] or AMG [3]. For AML, the restriction and interpolation operators involve the moment-to-discrete and discrete-to-moment operators, as will be described in the next section. For AMG, the restriction and interpolation operators and number of levels are handled by the Trilinos/MueLu software, and for IF the preconditioning is handled by the Trilinos/Ifpack2 software. For the SAAF and LS coarse-level solves, which use CFE, the mapping back and forth between the DFE and CFE representations of the spatial dependence is handled by Trilinos/Tpetra tools [14].

1. **Description of the TSA Algorithm**

The mono-energetic transport equation is

$$\mathbf{\Omega}\cdot\nabla\psi+\sigma_t\psi(\mathbf{r},\mathbf{\Omega}) = \int\sigma_s(\mathbf{\Omega}'\rightarrow\mathbf{\Omega})\psi(\mathbf{r},\mathbf{\Omega}')d\mathbf{\Omega}'+Q(\mathbf{r},\mathbf{\Omega}) \quad (1a)$$

with an imposed surface-source boundary condition

$$\psi(\mathbf{r},\mathbf{\Omega})=\psi_b(\mathbf{r},\mathbf{\Omega}),\ \ \mathbf{r}\in\Gamma,\ \mathbf{\Omega}\cdot\mathbf{n}<0, \quad (1b)$$

where $\mathbf{r}$ is the spatial position, $\mathbf{\Omega}$ is the particle direction of motion, σ_t is the total cross section, σ_s is the scattering cross section, ψ is the angular flux, and Q is a known distributed fixed source. Γ is the external spatial boundary and $\mathbf{n}$ is the unit outward normal on Γ. ψ_b is the imposed surface boundary condition for incoming directions.

Defining the transport operator as

$$\mathcal{T}\circ=\mathbf{\Omega}\cdot\nabla\circ+\sigma_t\circ \quad (2a)$$

and the scattering operator as

$$\mathcal{S}\circ=\int d\mathbf{\Omega}'\sigma_s(\mathbf{\Omega}'\rightarrow\mathbf{\Omega})\circ, \quad (2b)$$

the transport equation can be written as

$$\mathcal{T}\psi(\mathbf{r},\mathbf{\Omega})=\mathcal{S}\psi(\mathbf{r},\mathbf{\Omega})+Q(\mathbf{r},\mathbf{\Omega}). \quad (2c)$$

The transport equation can be solved by source iteration,

$$\mathcal{T}\psi^{(k+1)}=\mathcal{S}\psi^{(k)}+Q, \quad (3)$$

where $\psi^{(k)}$ is the known k^{th} iterate of the angular flux, $\psi^{(k+1)}$ is the computed $(k+1)^{st}$ iterate, with the iterations continuing until convergence. With an S_N discretization in angle and a DFE discretization in space, the $(k+1)^{st}$ iterate of the angular flux is obtained by an efficient sweeping algorithm.

In the TSA algorithm, the $(k+1)^{st}$ iterate of the angular flux is obtained by means of a two-step process: 1) a sweep solve, and 2) computation and application of a correction term. Re-indexing the result of a sweep solve as the $(k+1/2)^{th}$ iteration,

$$\mathcal{T}\psi^{(k+1/2)}=\mathcal{S}\psi^{(k)}+Q, \quad (4)$$

the goal is to determine a correction to the $(k+1/2)^{th}$ iterate of the solution to accelerate convergence. Defining the error term, $\varepsilon^{(k+1/2)}$, as the difference between the actual solution, ψ, and the $(k+1/2)^{th}$ iterate,

$$\psi=\psi^{(k+1/2)}+\varepsilon^{(k+1/2)}, \quad (5)$$

the residual at the $(k+1/2)^{th}$ iteration is

$$r^{(k+1/2)}=Q-(\mathcal{T}-\mathcal{S})\psi^{(k+1/2)}, \quad (6a)$$

or

$$r^{(k+1/2)}=\mathcal{S}(\psi^{k+1/2}-\psi^{k}). \quad (6b)$$

Combining Eqs. (6a), (2c) and (5) results in an expression for computing the correction term from the residual

$$(\mathcal{T}-\mathcal{S})\varepsilon^{(k+1/2)}=r^{(k+1/2)}. \quad (7)$$

Rather than solving Eq. (7) directly, which may be as difficult as solving the original transport equation, the procedure is to map the fine-level linear system for computing the correction term, Eq. (7), to a coarse-level linear system, which is solved for a coarse-level correction term that is then mapped back to fine-level and applied to the $(k+1/2)^{th}$ iterate of the angular flux. Defining a coarse-level linear system as

$$(\mathcal{T}_c - \mathcal{S}_c)\varepsilon_c^{(k+1/2)} = r_c^{(k+1/2)}, \qquad (8a)$$

where a coarse-level residual is computed from

$$r_c^{(k+1/2)} = \mathcal{R}r^{(k+1/2)}, \qquad (8b)$$

where $\mathcal{R}$ is a restriction operator for mapping from the fine-level linear problem to a coarse-level linear problem. Eq. (8a) can be solved by any desired method for the coarse-level correction term. Through numerical experimentation, an optimal convergence criterion (for minimum solver time) may be determined for solving the coarse-level system, often much more relaxed than that for the fine-level solves. The coarse-level correction term is mapped back to the fine-level problem by applying a prolongation operator,

$$\varepsilon^{(k+1/2)} = \mathcal{P}\varepsilon_c^{(k+1/2)}, \qquad (8c)$$

where $\mathcal{P}$ is a prolongation operator, for mapping the coarse-level problem back to the fine-level problem. The $(k+1)^{st}$ iterate of the angular flux is obtained by applying the correction term,

$$\psi^{(k+1)} = \psi^{(k+1/2)} + \varepsilon^{(k+1/2)}. \qquad (9)$$

Specific restriction and prolongation operators for using either spherical harmonics or discrete ordinates as a coarse-level solver are described in the next sections.

A. Using P_N as a Coarse-Level Solver

For a coarse-level P_N problem, the coarse-level residual moments are computed from

$$r_c^{(k+1/2)} = \mathcal{R}r^{(k+1/2)} = \mathcal{D}_{fc}\, r^{(k+1/2)}, \qquad (10a)$$

where $\mathcal{R}$ is the restriction operator, which is $\mathcal{D}_{fc}$, the fine-to-coarse discrete-to-moment operator. The coarse-level correction term is mapped back to fine level by

$$\varepsilon^{(k+1/2)} = \mathcal{P}\varepsilon_c^{(k+1/2)} = \mathcal{M}_{cf}\, \varepsilon_c^{(k+1/2)}, \qquad (10b)$$

where $\mathcal{P}$ is the prolongation operator, which is $\mathcal{M}_{cf}$, the coarse-to-fine moment-to-discrete operator.

B. Using S_N as a Coarse-Level Solver

For a coarse-level S_N problem, the coarse-level residual is computed from

$$r_c^{(k+1/2)} = \mathcal{R}r^{(k+1/2)} = \mathcal{M}_{cc}\mathcal{D}_{fc}\, r^{(k+1/2)}, \qquad (11a)$$

where $\mathcal{R}$ is the restriction operator, which is the product of $\mathcal{M}_{cc}$, the coarse-to-coarse moment-to-discrete operator, and $\mathcal{D}_{fc}$, the fine-to-coarse discrete-to-moment operator. The coarse-level error term is then mapped back to the fine level by

$$\varepsilon^{(k+1/2)} = \mathcal{P}\varepsilon_c^{(k+1/2)} = \mathcal{M}_{cf}\mathcal{D}_{cc}\varepsilon_c^{(k+1/2)}, \qquad (11b)$$

where $\mathcal{P}$ is the prolongation operator, which is the product of $\mathcal{M}_{cf}$, the coarse-to-fine moment-to-discrete operator, and $\mathcal{D}_{cc}$, the coarse-to-coarse discrete-to-moment operator.

2. **Description of the Preconditioned Krylov Algorithm**

Starting with the mono-energetic transport equation, Eq. (2c), and operating from the left with the inverse transport operator and rearranging results in [5-7]

$$(\mathcal{I} - \mathcal{T}^{-1}\mathcal{S})\psi(\mathbf{r},\mathbf{\Omega}) = \psi_{unc}(\mathbf{r},\mathbf{\Omega}), \qquad (12a)$$

where the uncollided flux is given by

$$\psi_{unc}(\mathbf{r},\mathbf{\Omega}) = \mathcal{T}^{-1}Q(\mathbf{r},\mathbf{\Omega}). \qquad (12b)$$

The linear system, Eq. (12a), is solved by a GMRES algorithm. We have implemented the GMRES(m) algorithm with restart, given by Saad [15] to solve a general linear system in SCEPTRE. The matrix-vector multiplication as required in constructing the Krylov subspace is replaced by the operation $(I\text{-}\mathcal{T}^{-1}S)\psi$. The computation of $\mathcal{T}^{-1}S$ is performed using a highly-efficient sweeping algorithm. The maximum dimension of the Krylov subspace is set to m so that the Krylov process is restarted after m iterations. This restart strategy can degrade the convergence of the GMRES iterations but is necessary to manage memory usage.

The GMRES iterations can be preconditioned by applying a left preconditioner to Eq. (12a),

$$\mathcal{M}^{-1}(\mathcal{I} - \mathcal{T}^{-1}\mathcal{S})\psi(\mathbf{r},\mathbf{\Omega}) = \mathcal{M}^{-1}\psi_{unc}(\mathbf{r},\mathbf{\Omega}), \quad (13)$$

where $\mathcal{M}^{-1}$ is an approximation to the system linear operator, I - $\mathcal{T}^{-1}S$.

A good approximation to the inverse of the linear-system operator may be obtained as follows. Using the uncollided flux as an approximate solution of the transport equation results in a residual of

$$r_{unc} = Q - (\mathcal{T} - \mathcal{S})\psi_{unc} = S\psi_{unc}. \qquad (14a)$$

The error term in using the uncollided flux as a approximate solution is

$$\varepsilon_{unc} = \psi - \psi_{unc} = (\mathcal{T} - \mathcal{S})^{-1}r_{unc}. \qquad (14b)$$

Combining Eqs. (14a) and (14b) results in an expression for the angular flux solution in terms of the uncollided flux,

$$\psi = \left[I + (\mathcal{T} - \mathcal{S})^{-1}S\right]\psi_{unc}. \qquad (15)$$

Comparing Eqs. (12a) and (15), the inverse linear operator can be written as

$$(\mathcal{I} - \mathcal{T}^{-1}\mathcal{S})^{-1} = \mathcal{I} + (\mathcal{T} - \mathcal{S})^{-1}\mathcal{S}. \qquad (16)$$

Mapping the transport solve in Eq. (16), $(\mathcal{T}-\mathcal{S})^{-1}$, to a coarse-level linear system results in an approximate inversion of the linear operator,

$$(\mathcal{I}-\mathcal{T}^{-1}\mathcal{S})^{-1} \simeq \mathcal{I}+\mathcal{P}(\mathcal{T}_c-\mathcal{S}_c)^{-1}\mathcal{R}\mathcal{S}, \quad (17a)$$

which indicates that an effective preconditioner is

$$\mathcal{M}^{-1} = \mathcal{I}+\mathcal{P}(\mathcal{T}_c-\mathcal{S}_c)^{-1}\mathcal{R}\mathcal{S}, \quad (17b)$$

In principle, any of the Krylov-method based TSA solvers, such as those mentioned in the previous section, can be used to perform the operation $(\mathcal{T}_c-\mathcal{S}_c)^{-1}$. So far, we have only implemented the GMRES(m) algorithm for this purpose.

III. RESULTS

In this section, the various preconditioners/accelerators will be applied to several test problems, using either typical cross sections for electron transport, or artificial cross sections with large scattering ratios. The effectiveness of the various methods will be evaluated by comparing the convergence rates of the error norm and also by comparing the solver times for the electron energy groups.

The TSA algorithms have many options available for controlling both the coarse-level solves and the coarse-level acceleration. One of the most important factors affecting overall runtime is specifying the convergence tolerance and/or maximum number of iterations for the coarse level solves. Often, running a small number of coarse-level iterations, resulting in a fairly crude approximation to the correction term, may result in a good speed up in the convergence of the fine-level source iterations. This is especially true for the DFE coarse-level solvers using GMRES, which is fairly expensive. For the SAAF and LS coarse-level solvers, which use CG, limiting the number of coarse-level iterations is not as crucial, since the CG algorithm is so fast.

For some of the results presented in this section, the coarse-level parameters are somewhat optimized for reducing overall runtime. For other of the results, the coarse-level parameters are set to provide a fairly well-converged coarse-level solve, which effectively speeds up convergence versus source iteration, while not necessarily providing optimal overall runtime.

1. Twisted-Pair Cable

The problem considered here is a segment of a braided-shield twisted-pair cable with copper conductors, shown in Fig. 1a. A planar source of 50-keV photons was incident approximately along the positive x-axis. The spatial mesh contained about a half-million linear hexahedral (hex8) elements, and the angular approximation used S_8 level-symmetric quadrature with P_3 scattering, including a δ-function down scatter correction. 10 linear photon and 10 log electron energy groups were used, and cross sections were obtained from the CEPXS code [16]. Fig. 1b shows the conductor pair with insulating and shielding materials removed.

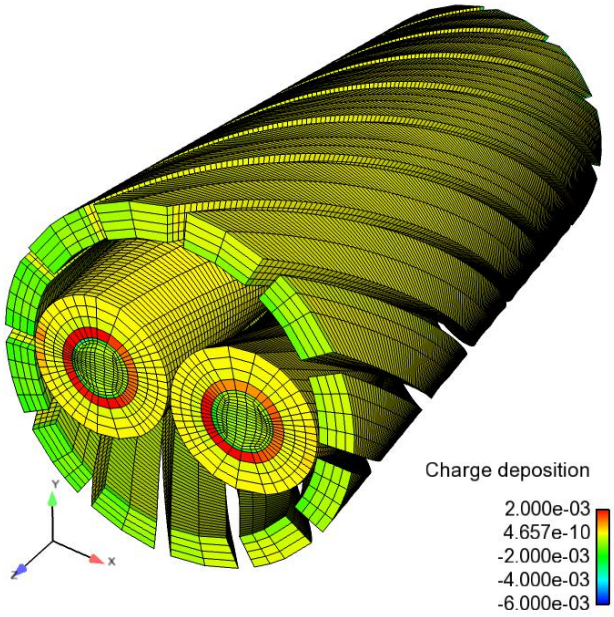

Fig. 1a. Charge-deposition distribution in braided-shield twisted-pair cable.

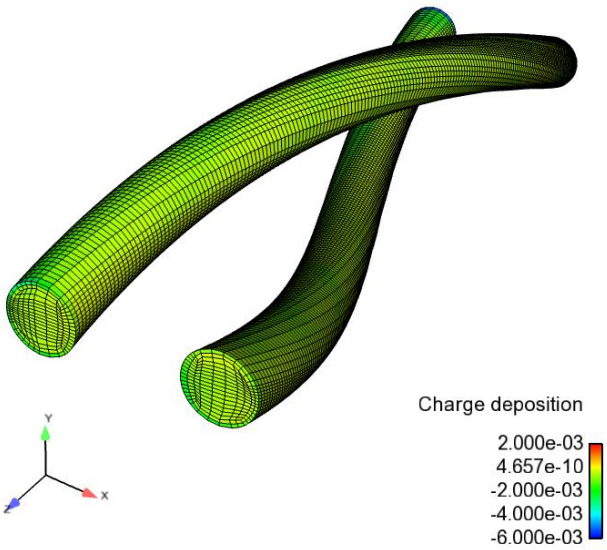

Fig. 1b. Charge-deposition distribution in conductor pairs.

Fig. 2 shows the scattering ratio for Cu as a function of electron energy, for either logarithmic or linear energy differencing and for P_3 or P_7 scattering order. The scattering ratio depends upon the scattering order due to the δ-function correction, which subtracts the $(L+1)^{st}$ order moment from the total and self-scatter cross sections. The scattering ratio also depends upon the energy grid, due to the manner in

which the electron multi-group cross sections are computed [16].

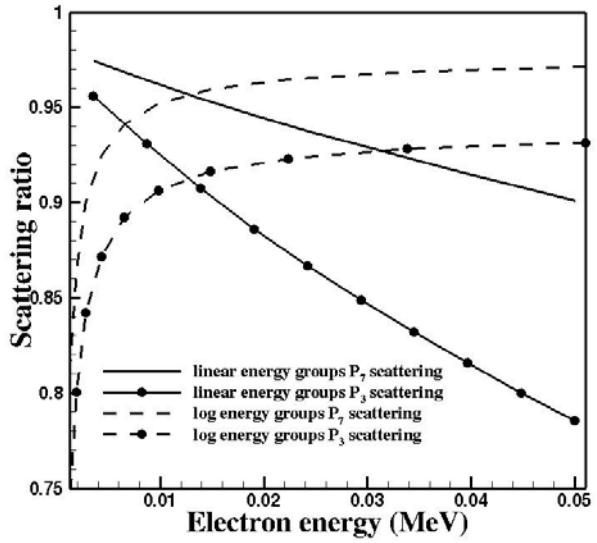

Fig. 2. Scattering ratio for electron transport in copper for different energy differencing and scattering order.

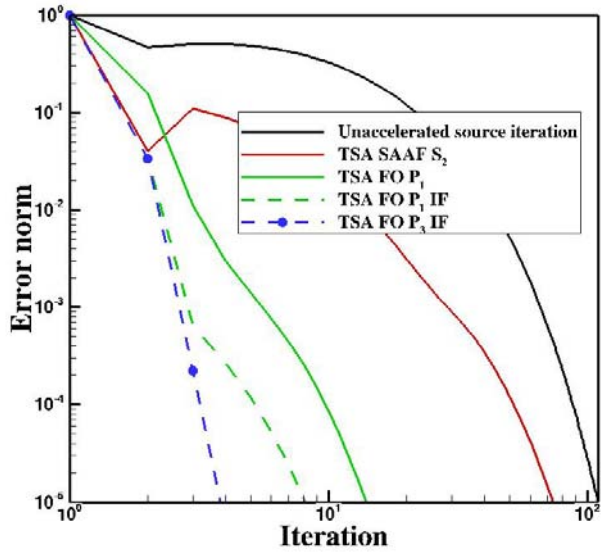

Fig. 3a. Convergence of the 5th electron group iterations for various TSA solvers.

Fig. 3a compares the convergence of the iterations for the 5th electron group for various TSA methods, including unaccelerated source iteration (SI), SAAF with CFE spatial differencing, and First-Order (FO) TSA with DFE spatial differencing, with and without Incomplete Factorization (IF) acceleration of the coarse-level solves. Solver time versus electron energy group is shown in Fig. 3b for the various acceleration methods. For this test problem, FO P_3 TSA with IF performs the best, in terms of both convergence rate and solver time. SAAF performed poorly for this test problem, being only slightly faster than unaccelerated SI. The timings were obtained by running on a TLCC2 machine, using 256 Intel Sandy Bridge 2.6 GHz cores. The convergence tolerance for the fine-level sweeps was an L_2 error norm of 10^{-5}.

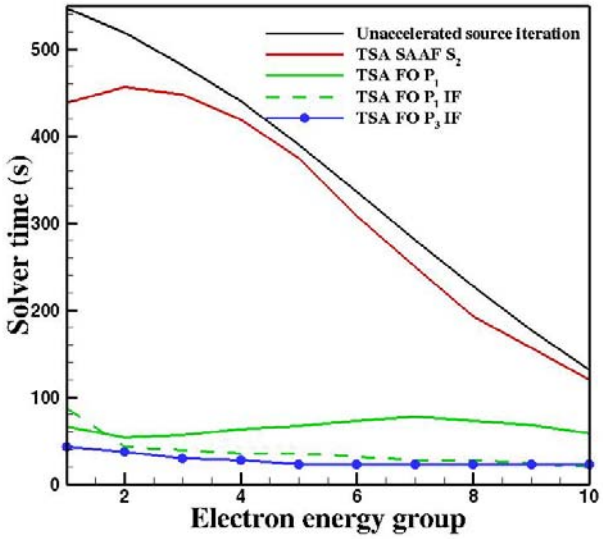

Fig. 3b. Solve times for the electron groups for various TSA solvers.

Table I lists the convergence metrics for the TSA solves and the combined solve time for the electron energy groups. Using P_3 FO TSA with IF resulted in a factor of 13 speedup compared with unaccelerated SI. Convergence metrics for the coarse-level solves were somewhat optimized experimentally for minimum solve time. For the FO DFE coarse solves, a fairly loose convergence tolerance of 0.01 was used, since the GMRES iterations are fairly expensive, providing minimal additional benefit for a tighter convergence. The use of IF to accelerate the coarse-level solves results in faster convergence, so a lower maximum iteration count of 10 was used for these runs. A tighter convergence tolerance of 10^{-4} was used for the SAAF TSA run, since the CG iterations are relatively inexpensive.

Table I. TSA convergence metrics and total solve time for electron groups for various acceleration methods.

TSA method	Coarse-level acceleration	Coarse-level convergence tolerance	Coarse-level max iterations	Total solve time of electron groups (s)
none	-	-	-	3530
FO P_1	none	0.01	100	653
FO P_1	IF	0.01	10	365
FO P_3	IF	0.01	10	272
SAAF S_2	none	10^{-4}	1000	3160

2. Three-Region Cylinder

The test problems in this section model transport in a three-region cylinder. For the first two tests, a plane-parallel source of 1-MeV electrons is incident on the side of the cylinder. The first test, described in Sec. 2.A, has lead in all three material regions, to test the performance of the accelerators in uniform material. In the test described in Sec. 2.B, the middle material region is replaced by void, to test the performance of the accelerators for problems with non-uniform material regions. The scalar flux in the inner- and outer-lead regions for the lead/void/lead test problem is shown in Fig. 4.

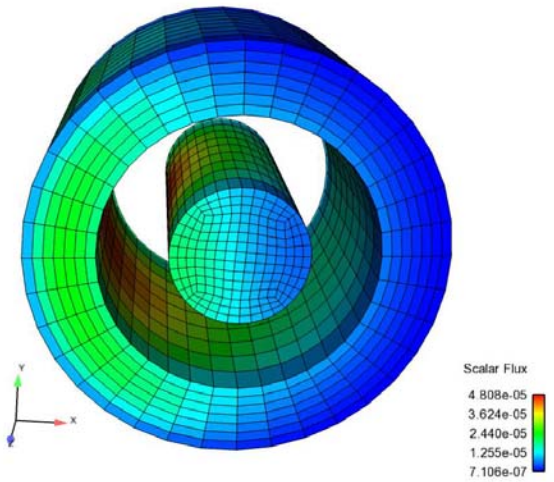

Fig. 4. Scalar flux in lead regions for lead/void/lead test problem.

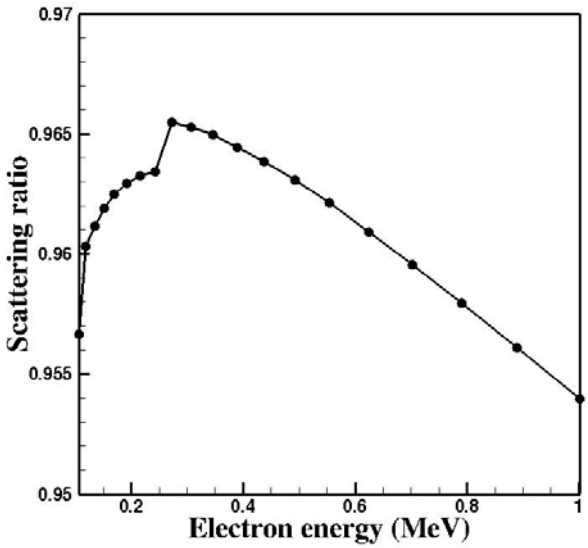

Fig. 5. Scattering ratio for electron transport in lead for 20 logarithmically-distributed energy groups and P_5 scattering order.

Fig. 5 shows the scattering ratio for lead for 20 logarithmically-distributed electron groups, with a maximum energy of 1 MeV and P_5 scattering. The third test problem in Sec. 3.C uses artificial cross sections with scattering ratio of unity, and compares the performance of the accelerators for P_3 and P_7 scattering, to investigate the performance of the accelerators with scattering anisotropy.

A. Performance of TSA in uniform material

In this section, a plane-parallel source of 1-MeV electrons is incident on a uniform lead cylinder. S_{10} level-symmetric quadrature is used with P_5 scattering, with 20 logarithmically-distributed energy groups. The figures compare the convergence of the error norm vs. iteration and solver time vs. energy group. In each of the plots, the performance of various TSA solvers is compared with unaccelerated source-iteration (SI) results and Krylov/GMRES solver results.

Fig. 6a compares the error-norm convergence of four different TSA solvers using CFE spatial differencing, using either a SAAF or LS transport solve and using a CG iterative solve. The TSA solves use either an S_2 discrete-ordinates solve or a P_1 spherical harmonics solve. The results shown in Fig. 6a are for the first energy group (1 MeV). The SPD accelerators for this test problem are very effective at reducing the error for the first few iterations, then tend to stall and then appear to become more effective as the convergence tolerance of 10^{-4} is approached. The iteration count is comparable to the Krylov/GMRES results for convergence tolerance of 10^{-4}.

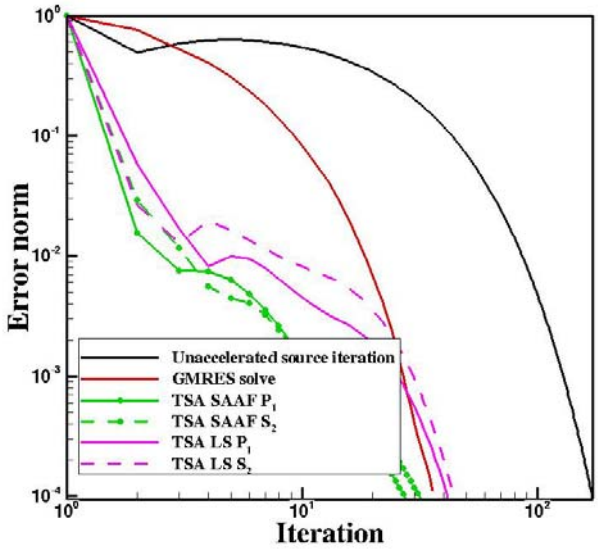

Fig. 6a. Error convergence of various CFE SPD TSA solvers compared with SI and Krylov/GMRES solvers for uniform-material test problem.

Fig. 6b compares solver times of various SPD CFE TSA solvers with SI and GMRES solvers. The solver time includes the time for the sweep solves and the TSA solves. The Krylov/GMRES solver generally performs better than the SPD TSA solves, especially for the lower-energy electron groups.

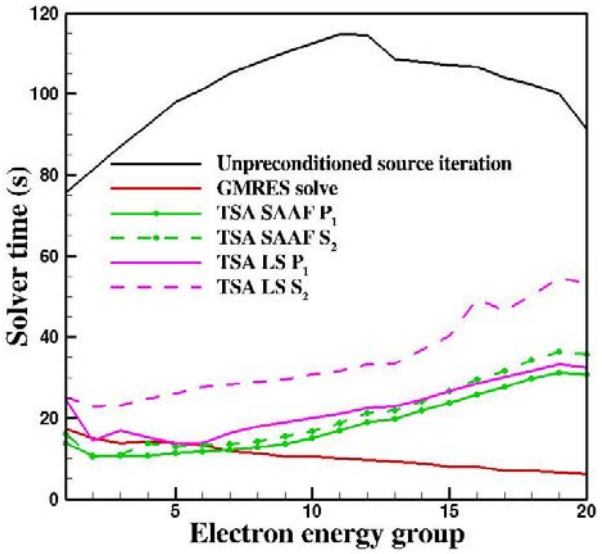

Fig. 6b. Solver time of various CFE SPD TSA solvers compared with SI and Krylov/GMRES solvers for uniform-material test problem.

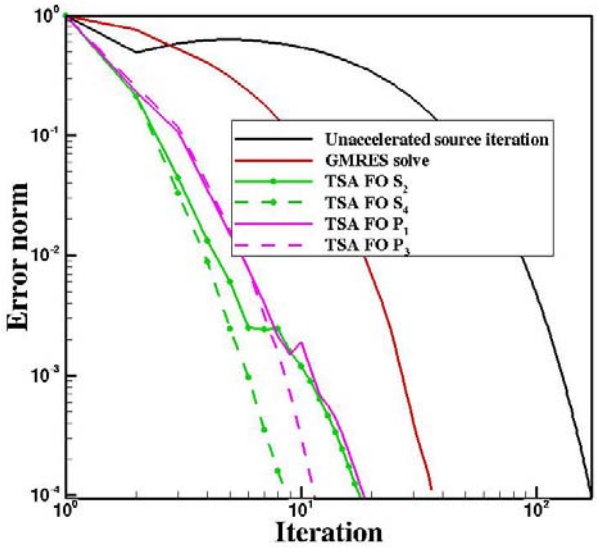

Fig. 7a. Error convergence of various DFE FO TSA solvers compared with SI and Krylov/GMRES solvers for uniform-material test problem.

Figs. 7a and 7b include performance information for the first-order (FO) DFE non-SPD TSA solvers, showing generally better performance than the SPD TSA solvers, both in convergence rate and solver times. The S_4 FO TSA solver shows the best convergence behavior, while the P_3 TSA solver using Incomplete Factorization has the best solver-time performance.

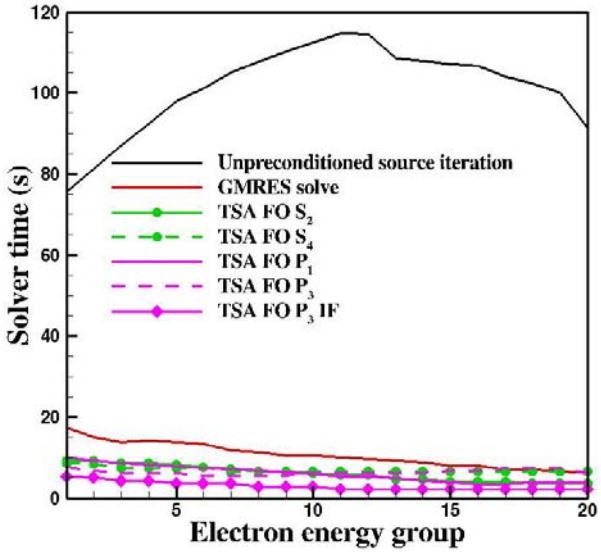

Fig. 7b. Solver time of various DFE FO TSA solvers compared with SI and Krylov/GMRES solvers for uniform-material test problem.

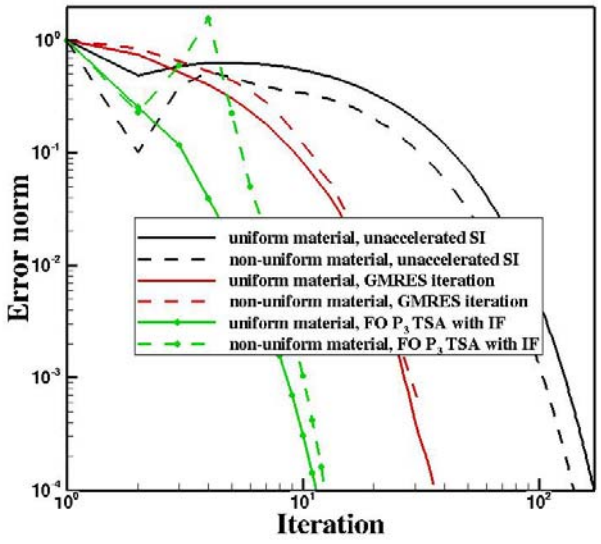

Fig. 8a. Effect of material discontinuity on convergence of various solvers.

B. Effect of Material Discontinuity

In this section, the effect of material discontinuity on the effectiveness of the methods is investigated, by replacing the middle material region by void, as shown in Fig. 4. Fig. 8a. compares convergence behavior for several solvers both for uniform lead and for the lead/void/lead configura-

tion. The Krylov/GMRES results are only slightly affected by the void region. The P_3 TSA method is negatively affected by the introduction of the void region. The method performs poorly for the first few iterations and then recovers with later iterations. Solver times are compared in Fig. 8b, showing an increase in global solver time of from 60 s for the uniform-material test problem to 143 s for the test problem including the void region, an increase in solver time by a factor of 2.4.

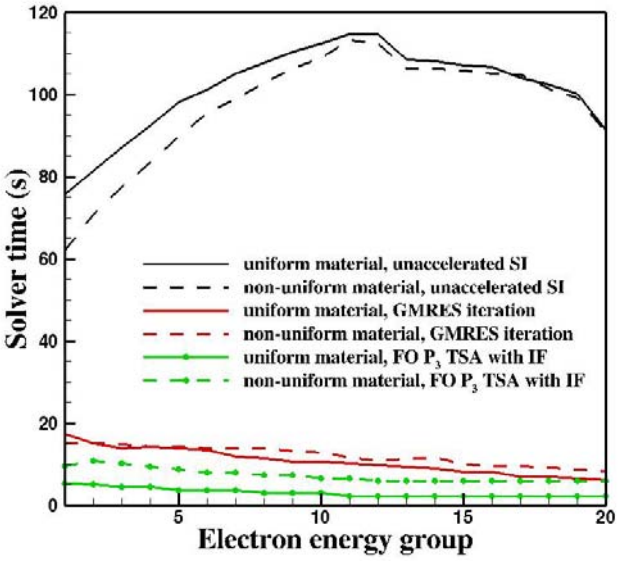

Fig. 8b. Effect of material discontinuity on solver time for various solvers/preconditioners.

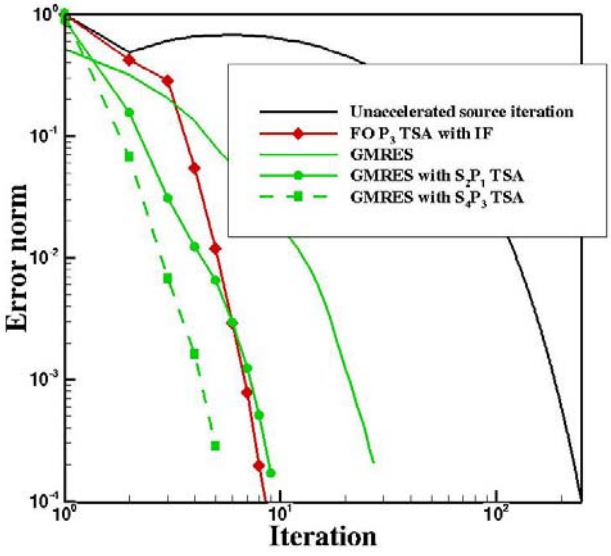

Fig. 9a. Effectiveness of TSA preconditioning on the iterative convergence of Krylov/GMRES algorithm for the lead/void/lead test problem.

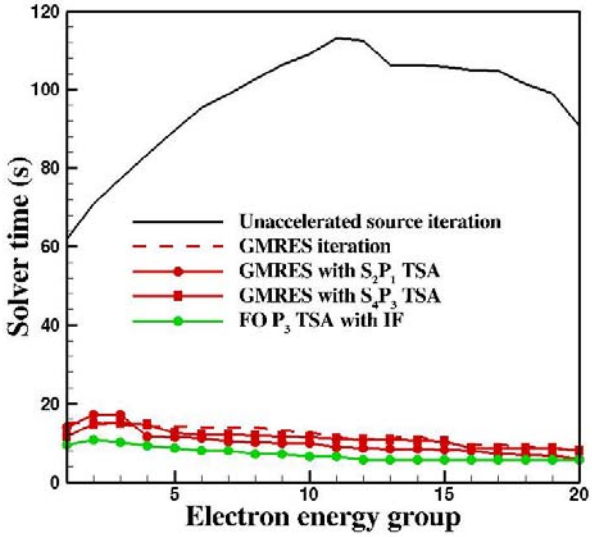

Fig. 9b. Comparison of solver times of various Krylov and TSA algorithms.

Figs. 9a and 9b show the effectiveness of TSA preconditioning applied to the Krylov/GMRES solver for the non-uniform material test problem. Preconditioning the GMRES iterations with S_4P_3 TSA effectively reduces the iterative convergence, resulting in a moderate decrease in solver time. The TSA solves for the GMRES preconditioning are currently performed by application of the GMRES(*m*) algorithm applied to the coarse-mesh linear problem. Other methods for performing this step in the computation are under investigation and may increase efficiency of the method.

C. Effect of Scattering Order

In this section, the effect of the scattering anisotropy on the effectiveness of the methods is investigated. Total and scattering cross sections for this test problem are shown in Table II. The scattering cross section is large, forward peaked, and highly scattering, analogous to electron transport.

Table II. Cross Sections for Test Problem

Cross Section	Inner	Middle	Outer
σ_t	$1.0\text{x}10^4$	0	$1.0\text{x}10^4$
σ_0	$1.0\text{x}10^4$	0	$1.0\text{x}10^4$
σ_1	$0.9\text{x}10^4$	0	$0.9\text{x}10^4$
σ_2	$0.8\text{x}10^4$	0	$0.8\text{x}10^4$
σ_3	$0.7\text{x}10^4$	0	$0.7\text{x}10^4$
σ_4	$0.6\text{x}10^4$	0	$0.6\text{x}10^4$
σ_5	$0.5\text{x}10^4$	0	$0.5\text{x}10^4$
σ_6	$0.4\text{x}10^4$	0	$0.4\text{x}10^4$
σ_7	$0.3\text{x}10^4$	0	$0.3\text{x}10^4$

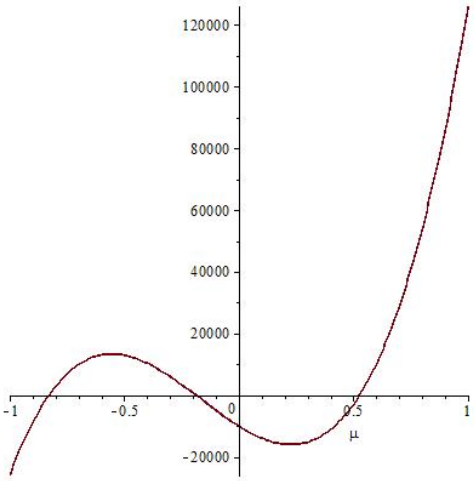

Fig. 10a. P_3 scattering cross section for test problem.

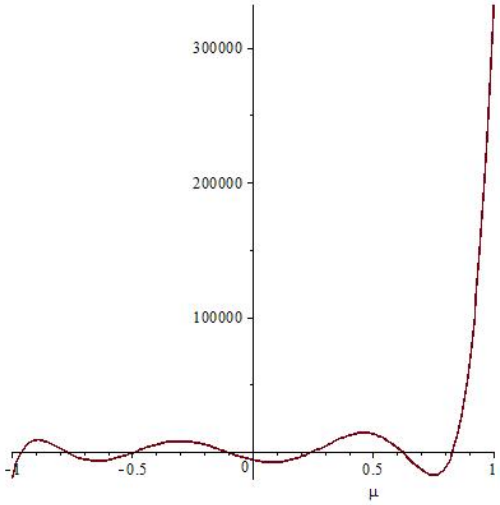

Fig. 10b. P_7 scattering cross section for test problem.

Results are compared for S_{16} level-symmetric quadrature and either P_3 or P_7 scattering. Figs. 10a and 10b show the angular dependence of the P_3 and P_7 scattering cross sections, respectively. Fig. 11 compares the convergence for various solvers for P_3 and P_7 scattering, indicating that the scattering order has little effect on any of the solvers. Unaccelerated SI diverges for this test problem. The SAAF TSA solver performs well for the first few iterations, then stalls for an error norm around 0.05, and again converges rapidly for later iterations as the convergence tolerance of 10^{-4} is approached. The Krylov/GMRES results are without TSA preconditioning and would be improved with the application of TSA. These results were obtained before TSA was implemented into the Krylov solver.

Fig. 12 shows the effectiveness of increasing the order of the TSA solver for increasing convergence rate, showing moderate improvement in going from P_1 to P_3 or S_2 to S_4 TSA. Table III compares global solver times for the two test problems, using P_3 and P_7 scattering, respectively. The solvers are essentially unaffected by the order of the scattering, with the P_1 TSA solver showing the best performance. For the SAAF TSA solver, the cross section was modified slightly to a scattering ratio of 0.99, since the SAAF algorithm is singular for pure scattering.

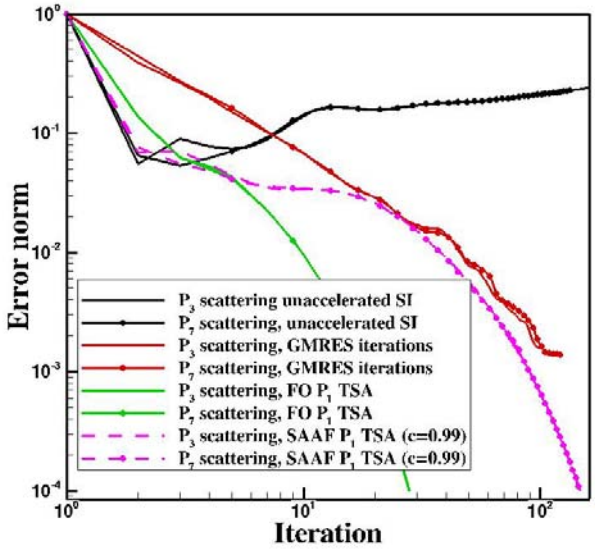

Fig. 11. Comparison of the various solvers/preconditioners for P_3 and P_7 scattering.

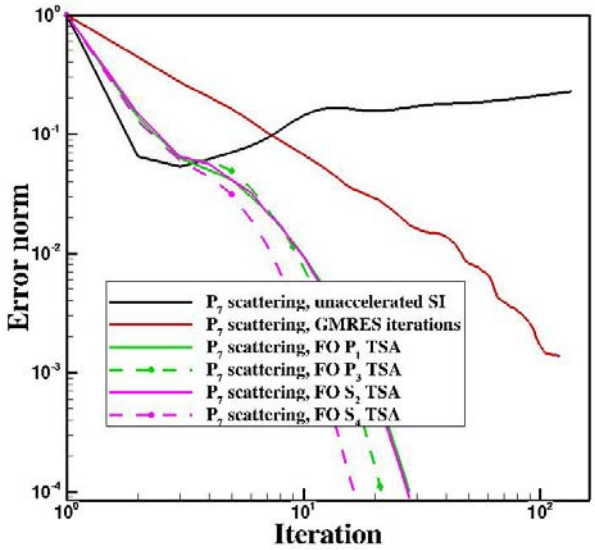

Fig. 12. Comparison of the effectiveness of TSA solvers of various orders.

Table III. Comparison of solver times for various solvers/preconditioners.

Acceleration method	Solver time (s)	
	P_3 scattering	P_7 scattering
none	-	-
GMRES	110	140
DFE P_1 TSA	36	38
DFE P_3 TSA	52	58
DFE S_2 TSA	58	48
DFE S_4 TSA	130	110
SAAF P_1 TSA (c=0.99)	180	190

IV. CONCLUSIONS

Accelerating/preconditioning of transport applications is a very rich and fascinating topic of study, due to the extreme variability in transport applications, often containing highly streaming regions adjacent to highly scattering regions, complex geometries, and a wide range of physics for different particle types and energies. In this work is presented a flexible ML approach for accelerating the convergence of source-iterations for problems involving highly-scattering media. The method attempts to efficiently compute a correction term for the fine-level source-iteration solve by solving a coarse-level level transport equation and mapping the coarse-level correction term to the fine-level space.

The method attempts to further accelerate the coarse-level solves by making use of existing off-the-shelf AMG and IF packages. In the application presented here, IF effectively reduced both iteration count and solve time when applied to the coarse-level linear system. AMG effectively reduced the iteration count, but due to the overhead cost was less effective at reducing solve time. Multi-grid methods tend to be most effective for elliptic problems, and the efficient use of multi-grid methods for other problem types needs further investigation.

The coarse-level solves can be either based on a 1st-order DFE discretization using GMRES or an SPD discretization using CG. For the application presented here, the DFE discretization with a small number of GMRES iterations (5-10) was the most effective at reducing solve time. The SPD CG solves were very fast but were generally less effective at reducing fine-level source-iteration count. This warrants further investigation; e.g. CG-based TSA methods may work well for problems with spatially smooth solutions, where a CFE differencing provides an accurate simulation.

The flexibility of the method could also be a liability, since there are many options available for the coarse level solves and many additional knobs to turn for the IF and AMG preconditioning, and it is not always clear for a particular application how to set the many options.

A GMRES iterative method has also been implemented that uses a parallel, sweep-based algorithm to build up a Krylov subspace, with TSA applied as a preconditioner to accelerate convergence of the GMRES iterations. Currently, the TSA solves use the GMRES(*m*) algorithm for the coarse-level solves. Other methods may provide better performance, but this warrants further investigation.

Finally, since the coarse-level linear system is built using Trilinos tools, Kokkos [17] capabilities should be effective for porting to next-generation processors, but that is for future investigations.

ACKNOWLEDGMENT

[a]Sandia National Laboratories is a multi-mission laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

REFERENCES

1. S. PAUTZ, C. DRUMM, W. BOHNHOFF AND W. FAN, "Software Engineering in the SCEPTRE Code," *Proc. M&C 2009,* Saratoga Springs, New York, May 3-7, 2009, American Nuclear Society (2009).
2. A. PROKOPENKO, C. SIEFERT, J. HU, M. HOEMMEN, AND A. KLINVEX, "Ifpack2 User's Guide 1.0 (Trilinos version 12.6), SAND2016-5338, Sandia National Laboratories (2016).
3. A. PROKOPENKO, J. HU, T. WIESNER, C. SIEFERT, AND R. TUMINARO, "MueLu User's Guide 1.0," SAND2014-18874, Sandia National Laboratories (2014).
4. G. L. RAMONÉ, M. ADAMS AND P. NOWAK, "A Transport Synthetic Acceleration Method for Transport Iterations," *Nucl. Sci. Eng.*, **125**, 257 (1997).
5. B. PATTON and J. HOLLOWAY, "Some Remarks on GMRES for Transport Theory," *Proc. M&C 2003*, Gatlinburg, Tennessee, April 6-11, 2003, American Nuclear Society (2003).
6. J. WARSA, T. WAREING, and J. MOREL, "Krylov Iterative Methods and the Degraded Effectiveness of Diffusion Synthetic Acceleration for Multidimensional S_N Calculations in Problems with Material Discontinuities," *Nucl. Sci. Eng.*, **147**, 218 (2004).
7. J. MOREL, "Basic Krylov Methods with Application to Transport," *Proc. M&C 2005*, Avignon, France, September 12-15, 2005, American Nuclear Society (2005).
8. J. WARSA, T. WAREING AND J. MOREL, "Fully Consistent Diffusion Synthetic Acceleration of Linear Discontinuous S_N Transport Discretizations on Three-

Dimensional Unstructured Tetrahedral Meshes," *Nucl. Sci. Eng.*, **141,** 236-251 (2002).
9. M. ADAMS AND E. LARSEN, "Fast Iterative Methods for Discrete-Ordinates Particle Transport Calculations," *Prog. Nucl. Ener*., **40**, 3 (2002).
10. B. TURCKSIN AND J. RAGUSA, "Discontinuous Diffusion Synthetic Acceleration for Sn Transport on 2D Arbitrary Polygonal Meshes," *J. Comp. Phys.* **274**, 356-369 (2014).
11. B. TURCKSIN, J. RAGUSA, AND J. MOREL, "Angular Multigrid Preconditioner for Krylov-Based Solution Techniques Applied to the Sn Equations with Highly Forward-Peaked Scattering," *Trans. Theor. And Stat. Phys.*, **41**, 1-22 (2012).
12. C. DRUMM, "Spherical Harmonics (P_N) Methods in the SCEPTRE Radiation Transport Code," *Proc. M&C 2015*, Nashville, Tennessee, April 19-23, 2015, American Nuclear Society (2015) (CD-ROM).
13. C. DRUMM, W. FAN AND S. PAUTZ, "Phase-Space Finite Elements in a Least-Squares Solution of the Transport Equation," *Proc. M&C 2013*, Sun Valley, Idaho, May 5-9, 2013, American Nuclear Society (2013).
14. M. HEROUX, et. al, "An Overview of Trilinos", SAND2003-2927, Sandia National Laboratories (2003).
15. Y. SAAD and M. Schultz, "GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems," *SIAM Sci. Stat. Comput.*, **7**, 856 (1986).
16. L. LORENCE, J. MOREL AND G. VALDEZ, "Physics Guide to CEPXS: A Multigroup Coupled Electron-Photon Cross-Section Generating Code Version 1.0," SAND89-1685, Sandia National Laboratories (1989).
17. J. AMELANG, C. TROTT AND H. EDWARDS, "Kokkos Tutorial," SAND2015-9620C, Sandia National Laboratories (2015).