Joint International Topical Meeting on Mathematics & Computation and Supercomputing in Nuclear Applications (M&C + SNA 2007)
Monterey, California, April 15-19, 2007, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2007)

KRYLOV ITERATIVE METHODS AND SYNTHETIC ACCELERATION FOR TRANSPORT IN BINARY STATISTICAL MEDIA

[†]Erin D. Fichtl, [*]James S. Warsa and [†]Anil K. Prinja

[†]Department of Chemical and Nuclear Engineering
The University of New Mexico
Albuquerque, NM 87131-1341
efichtl@unm.edu; prinja@unm.edu

[*]Transport Methods Group
Los Alamos National Laboratory
Los Alamos, NM 87545-0001
warsa@lanl.gov

ABSTRACT

Iterative solution of the Levermore-Pomraning equations for transport in binary statistical mixtures can be extremely slow in certain limits. We propose an iterative method that improves convergence by utilizing a combination of inner iterations, synthetic acceleration schemes and Krylov iterative methods. Spectral analysis and numerical results show that our new scheme outperforms simpler iterative methods for problems in either the diffusion or the atomic mix regimes, or both.

Key Words: Krylov iterative methods, synthetic acceleration, stochastic mixtures,

Levermore-Pomraning closure

1. INTRODUCTION

The development of transport methods for stochastic mixtures is necessary for the numerous physical constructs in which heterogeneities are not distributed in a predictable and ordered way —that is, when the material composition and distribution are known only statistically. We then describe the flux and other physical quantities of interest in terms of an ensemble average, or the average over all possible "realizations" of the material. A closed system of equations for the ensemble average of the flux was derived twenty years ago by Levermore, Pomraning and Vanderhaegen for binary, or two-state, mixtures using the Levermore-Pomraning (LP) closure [1, 2]. The model comprises two coupled equations for the conditional ensemble average of the flux in the two materials.

Unfortunately, iterative convergence of the discretized system can be very slow in diffusive and/or atomically mixed regimes. Recently, two distinct acceleration techniques have been devised to remedy different aspects of this problem. The first is a coupled diffusion synthetic acceleration (DSA) scheme, which was designed to accelerate iterative solution in optically thick, diffusive materials, although it was successful in a wider variety of regimes [3]. The second is an atomic mix synthetic acceleration (AMSA) scheme which exploits the asymptotic atomic mix limit to calculate a lower-order correction for the flux in the absence of scattering [4]. However, a scheme has not yet been devised to accelerate convergence when the material is both diffusive and atomically mixed. We propose a three-part acceleration scheme that

accomplishes this objective. Firstly, we divide the outer iteration into two inner iterations. One we call the "material iteration," which attenuates error modes in the atomic mix limit. The other we call the "species iteration," which attenuates error modes in the diffusion limit. Secondly, we apply atomic mix synthetic acceleration (AMSA) to the material iteration and S_2 synthetic acceleration (S_2SA) to the inner source iterations to offset the cost of doing three inner iterations per outer iteration. Finally, we wrap a Krylov iterative solver around both the inner and outer iterations to further accelerate convergence. We display the effectiveness and efficiency of the new iteration scheme by presenting a spectral analysis of the iteration operator and measuring the iteration count and computing cost compared against those for a basic one-step iteration (which can also be accelerated using a Krylov iterative method). Note that our methodology and some numerical results have been presented previously [5], but that this paper explores a more comprehensive set of test cases, complementary spectral analysis and an exploration of the behavior of the inner iterations.

2. ITERATIVE SOLUTION OF THE LEVERMORE-POMRANING EQUATIONS

2.1. One-Step Method

The transport of neutrons in one-dimensional, two-state statistical mixtures is described approximately by the following form of the S_N LP equations:

$$\hat{L}_{1,n}\psi_{1,n}^{(m+1)} = \frac{\Sigma_{s,1}}{4\pi}\phi_1^{(m)} + \frac{|\mu_n|}{\lambda_1}\psi_{2,n}^{(m)} + \frac{1}{4\pi}Q_{1,n} \tag{1a}$$

$$\hat{L}_{2,n}\psi_{2,n}^{(m+1)} = \frac{\Sigma_{s,2}}{4\pi}\phi_2^{(m)} + \frac{|\mu_n|}{\lambda_2}\psi_{1,n}^{(m+1)} + \frac{1}{4\pi}Q_{2,n}. \tag{1b}$$

The index m is used to denote the basic one-step iteration, the μ_n ($n = 1 \ldots N$) are Gauss quadrature angles and $\hat{L}_{\ell,n} = \left(\mu_n \frac{\partial}{\partial s} + \Sigma_\ell + \frac{|\mu_n|}{\lambda_\ell}\right)$. For material ℓ, Σ_ℓ and $\Sigma_{s,\ell}$ are total and scattering cross sections, λ_ℓ is the mean chord length, $\psi_{\ell,n}$ is the conditional ensemble average of the angular flux at angle μ_n, ϕ_ℓ is the conditional ensemble average of the scalar flux, and $Q_{\ell,n}$ is an isotropic volume source. The one-step iteration is robust but it converges slowly in either the atomic mix limit, $\lambda_\ell\Sigma_\ell \ll 1$, or in the diffusion limit, $c_\ell = \Sigma_{s,\ell}/\Sigma_\ell \approx 1$, or both.

2.2. Two-Step Method

Alternately, we could separate each outer iteration into a sequence of two inner iterations: A "material iteration" that attenuates the atomic mix error mode and a "species iteration" consisting of two source iterations, one for each material, that attenuates diffusive error modes. Since the prevalent error modes are being handled separately in this two-step outer iteration, it is expected that this scheme will converge in fewer iterations than the one-step outer iteration scheme (Eqs. 1). Consider the following two-step scheme, with outer iteration index m, comprising two inner iterations. The index i is assigned to the first inner iteration, or material iteration, which is given by:

$$\hat{L}_{1,n}\psi_{1,n}^{(i+\frac{1}{2})} = \frac{\Sigma_{s,1}}{4\pi}\phi_1^{(m)} + \frac{|\mu_n|}{\lambda_1}\psi_{2,n}^{(i)} + \frac{1}{4\pi}Q_{1,n} \tag{2a}$$

$$\hat{L}_{2,n}\psi_{2,n}^{(i+\frac{1}{2})} = \frac{\Sigma_{s,2}}{4\pi}\phi_2^{(m)} + \frac{|\mu_n|}{\lambda_2}\psi_{1,n}^{(i+\frac{1}{2})} + \frac{1}{4\pi}Q_{2,n}, \tag{2b}$$

where the values of $\phi_1^{(m)}$ and $\phi_2^{(m)}$ from the previous outer iteration are held constant.

The result of this iteration gives $\psi_{1,n}^{(m+\frac{1}{2})}$ and $\psi_{2,n}^{(m+\frac{1}{2})}$, which are then held constant in the second inner iteration, or species iteration. The species iteration comprises two independent iterations, each with index j:

$$\hat{L}_{\ell,n}\psi_{\ell,n}^{(j+\frac{1}{2})} = \frac{\Sigma_{s,\ell}}{4\pi}\phi_{\ell}^{(j)} + \frac{|\mu_n|}{\lambda_\ell}\psi_{\ell,n}^{(m+\frac{1}{2})} + \frac{1}{4\pi}Q_{\ell,n}, \tag{3}$$

for $l = 1, 2$. The result of this iteration gives $\phi_1^{(m+1)}$ and $\phi_2^{(m+1)}$ for the next outer iteration.

Since each two-step iteration requires three inner iterations while the one-step iteration requires only a single sweep, it is imperative that the inner iterations are accelerated. The structure of the inner iterations easily allows for the use of known synthetic acceleration schemes. Specifically, the material iteration is accelerated with atomic mix synthetic acceleration (AMSA) [4] and the species iteration is accelerated using S_2 synthetic acceleration (S_2SA) [6].

2.3. Operator Notation

The system is discretized using a linear discontinuous finite element method. The discretized form of the outer iteration schemes can then be written in general matrix notation as

$$\vec{\Psi}^{(m+1)} = \mathbf{T}\vec{\Psi}^{(m)} + \vec{\Psi}_0, \tag{4}$$

where $\vec{\Psi}$ is $\vec{\psi}$ for the one-step iteration, $\vec{\phi}$ for the two-step, $\vec{\psi}_2$ for the material iteration and $\vec{\phi}_\ell$ for the species iteration in material ℓ. Multiplication by $\mathbf{T}$ represents the inversion of a transport operator and $\vec{\Psi}_0$ represents an uncollided flux or source vector [7], which is computed before iteration begins.

Rewriting the accelerated forms of the inner iterations (Eqs. 2 and 3) in matrix notation yields:

$$\vec{\Psi}^{(i+\frac{1}{2})} = \mathbf{T}\vec{\Psi}^{(i)} + \vec{\Psi}_0 + \mathbf{E}\left(\vec{\Psi}^{(i+\frac{1}{2})} - \vec{\Psi}^{(i)}\right) = \mathbf{T}\vec{\Psi}^{(i)} + \vec{\Psi}_0 + \mathbf{E}(\mathbf{T}\vec{\Psi}^{(i)} + \vec{\Psi}_0 - \vec{\Psi}^{(i)}), \tag{5}$$

where $\mathbf{E}$ represents the inversion of the AMSA or S_2SA transport operator and is the lower-order atomic mix or S_2 approximation of $(\mathbf{I} - \mathbf{T})^{-1}\mathbf{T}$. It multiplies the residual, $\left(\vec{\Psi}^{(i+\frac{1}{2})} - \vec{\Psi}^{(i)}\right)$, to yield the AMSA or S_2SA correction of the angular flux.

2.4. Krylov Iterative Methods

While the two-step outer iteration scheme is expected to converge in fewer iterations than the one-step scheme, slow convergence is anticipated in the atomic mix-diffusion regime, where $c_\ell \approx 1$ and $\lambda_\ell\Sigma_\ell \ll 1$. Krylov iterative methods have proven to be efficient and effective in problems for which traditional source iteration converges slowly, particularly when preconditioned with DSA [8].

In order to solve our system using a Krylov iterative method, we rewrite Eq. 4 as

$$(\mathbf{I} - \mathbf{T})\vec{\Psi} = \vec{\Psi}_0. \tag{6}$$

Because a single transport sweep calculates the product $\mathbf{T}\vec{\Psi}$, the algorithm is very simply modified to compute $(\mathbf{I} - \mathbf{T})\vec{v}$, which represents the action of an operator corresponding to the one-step or two-step iteration, or to any of the inner iterations, on a vector $\vec{v}$ that is supplied by a Krylov method at every

iteration. For the inner iterations, it is possible to left precondition the system using AMSA and S_2SA in order to accelerate the convergence of the Krylov solvers. Rearranging Eq. 5 yields

$$(\mathbf{I}+\mathbf{E})(\mathbf{I}-\mathbf{T})\vec{\Psi} = (\mathbf{I}+\mathbf{E})\vec{\Psi}_0. \tag{7}$$

Because the synthetic acceleration algorithms compute the operation of $\mathbf{E}$ on the supplied residual, they are easily modified to return the action of $(\mathbf{I}+\mathbf{E})$ on a vector. We will use restarted Generalized Minimal RESidual (GMRES(m)) method as it was designed for use with nonsymmetric operators such as our transport operator.

3. SPECTRAL ANALYSIS

The spectral radius of the operator $\mathbf{T}$ determines the convergence rate of the simple iterative schemes. Thirty six test cases were chosen to explore a wide range of problem regimes and are delineated in Table I. As can be seen in Table II, in all cases, the spectral radii for the two-step operator were smaller than those for the one-step operator. For both the one-step and two-step schemes, the spectral radii were smallest when the material was not atomically mixed and/or diffusive, larger when only one material was atomically mixed and/or diffusive and largest when both materials were atomically mixed and/or diffusive.

Table I. Test Problem Parameters (width = 10.0 cm, 100 spatial cells and S-8 quadrature order)

Case #	λ_1	λ_2	Σ_1	Σ_2	$\Sigma_{s,1}$	$\Sigma_{s,2}$	Case #	λ_1	λ_2	Σ_1	Σ_2	$\Sigma_{s,1}$	$\Sigma_{s,2}$
1	1.0	5.0	5.0	1.0	5.0	1.0	19	0.01	0.005	5.0	1.0	4.5	0.9
2	0.1	0.5	5.0	1.0	5.0	1.0	20	0.001	5.0	5.0	1.0	4.5	0.9
3	0.01	0.05	5.0	1.0	5.0	1.0	21	0.001	0.5	5.0	1.0	4.5	0.9
4	0.001	0.005	5.0	1.0	5.0	1.0	22	0.001	0.05	5.0	1.0	4.5	0.9
5	1.0	5.0	5.0	1.0	4.95	0.99	23	0.01	0.05	5.0	1.0	4.95	1.0
6	0.1	0.5	5.0	1.0	4.95	0.99	24	0.01	0.05	5.0	1.0	4.5	1.0
7	0.01	0.05	5.0	1.0	4.95	0.99	25	0.01	0.05	5.0	1.0	2.5	1.0
8	0.001	0.005	5.0	1.0	4.95	0.99	26	0.01	0.05	5.0	1.0	5.0	0.99
9	1.0	5.0	5.0	1.0	4.5	0.9	27	0.01	0.05	5.0	1.0	5.0	0.9
10	0.1	0.5	5.0	1.0	4.5	0.9	28	0.01	0.05	5.0	1.0	5.0	0.5
11	0.01	0.05	5.0	1.0	4.5	0.9	29	1.0	5.0	5.0	0.1	5.0	0.1
12	0.001	0.005	5.0	1.0	4.5	0.9	30	0.1	0.5	5.0	0.1	5.0	0.1
13	1.0	5.0	5.0	1.0	2.5	0.5	31	0.01	0.05	5.0	0.1	5.0	0.1
14	0.1	0.5	5.0	1.0	2.5	0.5	32	0.001	0.005	5.0	0.1	5.0	0.1
15	0.01	0.05	5.0	1.0	2.5	0.5	33	1.0	5.0	5.0	0.01	5.0	0.01
16	0.001	0.005	5.0	1.0	2.5	0.5	34	0.1	0.5	5.0	0.01	5.0	0.01
17	1.0	0.005	5.0	1.0	4.5	0.9	35	0.01	0.05	5.0	0.01	5.0	0.01
18	0.1	0.005	5.0	1.0	4.5	0.9	36	0.001	0.005	5.0	0.01	5.0	0.01

The spectral radii for the inner material iterations (see Table II) were large when the λs were small, and vice versa, but were unaffected by the magnitude of c. Not surprisingly, the AMSA method was

particularly effective at reducing the spectral radius when the material was atomically mixed. The spectral radii for the inner species iterations, on the other hand, were affected by both λ and c because Σ_ℓ is modified by the term $\frac{|\mu_n|}{\lambda_\ell}$ (see the definition of the operator $\hat{L}_{\ell,n}$ following Eq. 1). Therefore, the spectral radius is large whenever c and λ_ℓ are large, and small when *either* c or λ_ℓ are small. S_2SA always effectively reduced the spectral radius.

Table II. Spectral Radii

Case #	One-Step	Two-Step	Material Iteration		Source Iteration 1		Source Iteration 2	
			Unaccel.	Accel.	Unaccel.	Accel.	Unaccel.	Accel.
1	0.9966936	0.9622605	0.004604727	0.03615472	0.9102434	0.2091249	0.9056439	0.2106908
2	0.9980027	0.9935060	0.2192886	0.1164641	0.5386727	0.1482097	0.5380735	0.1488443
3	0.9995424	0.9968614	0.8135431	0.1216091	0.1230802	0.04713692	0.1230751	0.04713548
4	0.9999469	0.9972437	0.9792762	0.04595763	0.01449146	0.005956703	0.01449145	0.005956700
5	0.9872019	0.8670624	0.004604727	0.03615472	0.9011409	0.2058952	0.8965875	0.2075174
6	0.9913585	0.9722325	0.2192886	0.1164641	0.5332859	0.1460215	0.5326928	0.1466865
7	0.9978882	0.9855096	0.8135431	0.1216091	0.1218494	0.04662518	0.1218444	0.04662375
8	0.9997527	0.9871261	0.9792762	0.04595763	0.01434654	0.005896626	0.01434653	0.005896622
9	0.9016800	0.4358996	0.004604727	0.03615472	0.8192190	0.1789132	0.8150795	0.1809711
10	0.9314274	0.8007682	0.2192886	0.1164641	0.4848054	0.1274916	0.4842661	0.1283738
11	0.9832774	0.8847576	0.8135431	0.1216091	0.1107722	0.04205899	0.1107676	0.04205773
12	0.9980562	0.8962015	0.9792762	0.04595763	0.01304231	0.005356394	0.01304230	0.005356391
13	0.5186925	0.08109350	0.004604727	0.03615472	0.4551217	0.08562440	0.4528220	0.08787327
14	0.6634324	0.3137835	0.2192886	0.1164641	0.2693363	0.06153971	0.2690367	0.06266294
15	0.9274770	0.4657467	0.8135431	0.1216091	0.06154012	0.02259028	0.06153757	0.02258964
16	0.9919523	0.4949826	0.9792762	0.04595763	0.007245729	0.002965511	0.007245725	0.002965509
17	0.9089299	0.4501294	0.1239213	0.09597821	0.8192190	0.1789132	0.01304230	0.005356391
18	0.9502559	0.8072400	0.5797576	0.5309616	0.4848054	0.1274916	0.01304230	0.005356391
19	0.9903724	0.8880756	0.9215766	0.7975140	0.1107722	0.04205899	0.01304230	0.005356391
20	0.9032171	0.4299275	0.05371041	0.007998693	0.01304231	0.005356394	0.8150795	0.1809711
21	0.9467202	0.7959895	0.3822993	0.06235006	0.01304231	0.005356394	0.4842661	0.1283738
22	0.9897244	0.8824906	0.8693382	0.1923093	0.01304231	0.005356394	0.1107676	0.04205773
23	0.9987148	0.9911827	0.8135431	0.1216091	0.1218494	0.04662518	0.1230751	0.04713548
24	0.9913584	0.9406001	0.8135431	0.1216091	0.1107722	0.04205899	0.1230751	0.04713548
25	0.9609431	0.7254153	0.8135431	0.1216091	0.06154012	0.02259028	0.1230751	0.04713548
26	0.9987150	0.9911844	0.8135431	0.1216091	0.1230802	0.04713692	0.1218444	0.04662375
27	0.9913610	0.9406168	0.8135431	0.1216091	0.1230802	0.04713692	0.1107676	0.04205773
28	0.9609585	0.7254941	0.8135431	0.1216091	0.1230802	0.04713692	0.06153757	0.02258964
29	0.9834807	0.8522786	0.004380495	0.03138147	0.9102434	0.2091249	0.4999837	0.1439586
30	0.9941145	0.9777877	0.2396072	0.1111660	0.5386727	0.1482097	0.1225893	0.04699587
31	0.9991322	0.9902529	0.8530386	0.2175930	0.1230802	0.04713692	0.01449061	0.005956363
32	0.9999079	0.9915338	0.9837568	0.3495543	0.01449146	0.005956703	0.001476247	0.00061150055
33	0.9778270	0.7530704	0.004277812	0.03023650	0.9102434	0.2091249	0.1006161	0.03974493
34	0.9930783	0.9710386	0.24958820	0.1109693	0.5386727	0.1482097	0.01441389	0.005925595
35	0.9990532	0.9884460	0.85929287	0.2534003	0.1230802	0.04713692	0.001476158	0.0006114638
36	0.9999007	0.9900919	0.9841007	0.3922212	0.01449146	0.005956703	0.0001479024	0.00006131269

There are several properties of an operator that influence the choice of a Krylov iterative method and yield some information about its rate of convergence, which is difficult to quantify precisely. Among these are the symmetry and positive-definiteness of the operator, and whether it is normal in the case that it is nonsymmetric. However, in general the eigenvalue spectrum of the matrix $(\mathbf{I}-\mathbf{T})$ yields qualitatively useful information about the rate of convergence. Loosely speaking, the convergence rate is determined by the distribution of eigenvalues. It has been shown that it is directly proportional to the radius of the circles bounding clusters of eigenvalues, relative to their centers, along with the relative distances between clusters [9]. Convergence is best for small clusters centered near unity and is worse for large clusters or clusters centered near zero. We examine spectra for some representative problems in Fig. 1.

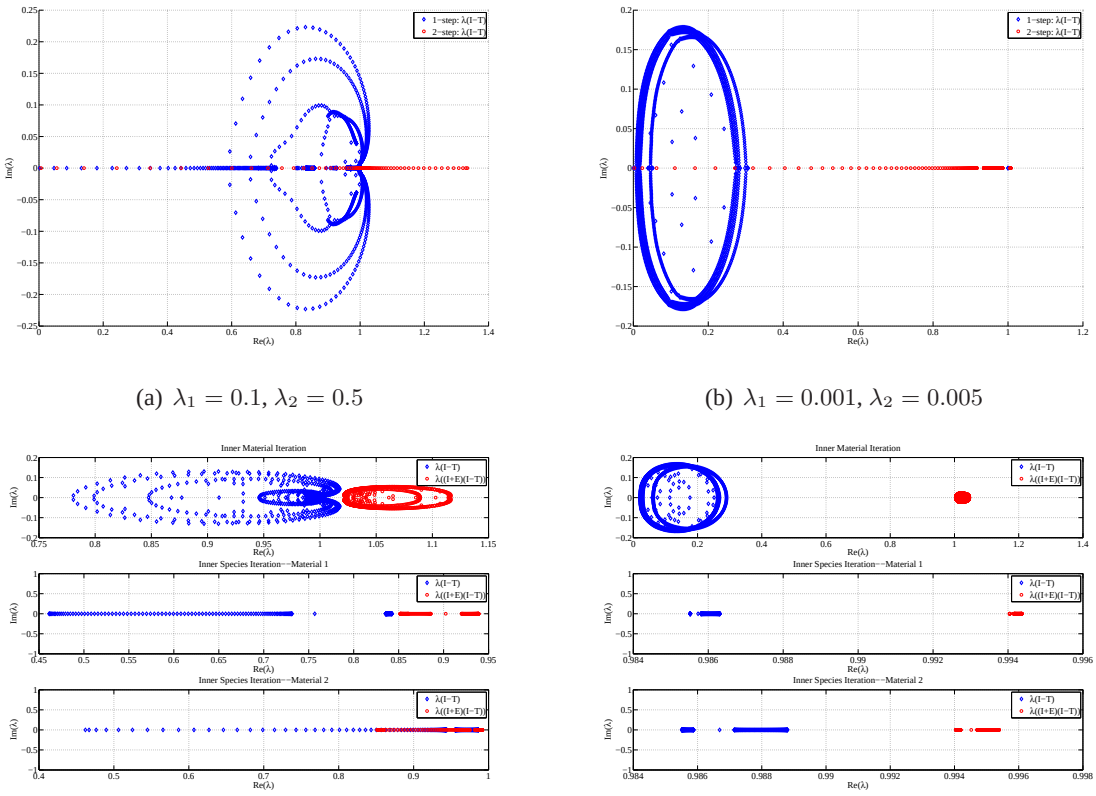

(a) $\lambda_1 = 0.1$, $\lambda_2 = 0.5$

(b) $\lambda_1 = 0.001$, $\lambda_2 = 0.005$

(c) $\lambda_1 = 0.1$, $\lambda_2 = 0.5$

(d) $\lambda_1 = 0.001$, $\lambda_2 = 0.005$

Figure 1. Spectrum of $(\mathbf{I}-\mathbf{T})$ for $\boldsymbol{\Sigma_1 = 5.0}$ cm^{-1}, $\boldsymbol{\Sigma_2 = 1.0}$ cm^{-1}, $\mathbf{c_1 = c_2 = 1.0}$ (Cases 2 and 4)

In each of the plots, $c_1 = c_2 = 1.0$, but the materials are more atomically mixed in 1(b) and 1(d). We see that the spectrum of the two-step method is tightly clustered at one with a few outliers, whereas the one-step spectrum is less tightly clustered and the clusters are farther from unity. In Fig. 1(b), the one-step spectrum has shifted closer to zero. Although the cluster is slightly smaller than in 1(a), its proximity to zero indicates that convergence could be slow. The spectrum for the two-step operator is also shifted away from one, although not as dramatically as for the one-step operator, and there are more outliers between zero and unity. The spectra indicate that both operators are well-suited for GMRES and convergence is slower in the atomic mix regime for both operators, but the two-step method should converge more rapidly in general. For smaller values of c, shapes of the spectra are similar but the outliers are closer to unity,

indicating more rapid convergence.

The preconditioned operators for all three of the inner iterations (1(c) and 1(d)) show better convergence properties than the unpreconditioned as their spectra are more tightly clustered and the centers of their clusters are closer to one. It can also be seen that the spectrum of the unpreconditioned material iteration operator in 1(c) shows much better convergence properties than that in 1(d), which is extremely close to zero since the material is atomically mixed. The preconditioned operators show the opposite behavior, however, demonstrating the effectiveness of AMSA as a preconditioner in the atomic mix limit. Interestingly, as discussed previously, varying λ does effect the convergence rates of the species iterations. As can be seen in comparing 1(c) and 1(d), while c is the same in both cases, the spectra in 1(d) are tightly clustered close to one while the spectra in 1(c) are far from one and are much more spread out, indicating that the species iterations should converge more quickly in the atomic mix case. While decreasing c has no effect on the inner material iterations, it does cause the clusters in the spectrum of the species iterations to become smaller and shift toward unity.

Restarted GMRES is guaranteed to converge if the matrix is positive definite (PD) [10, 11]. The operator $(\mathbf{I}-\mathbf{T})$ and, for the inner iterations, the corresponding preconditioned operators $(\mathbf{I}+\mathbf{E})(\mathbf{I}-\mathbf{T})$ were tested for positive definiteness. We found that the operator for the one-step scheme was PD in only five of the thirty-six cases examined and the two-step operator was PD in all but eight of the cases examined. The inner material iteration is not PD in those cases where both λ_1 and λ_2 are $\mathcal{O}(10^{-3})$, but the *preconditioned* operator is PD in all cases. The inner species iteration operators are PD in all cases, both with and without preconditioning. Note that although restarted GMRES is guaranteed to converge whenever the operator is PD, it is also not guaranteed to fail when the operator is not PD. Indeed, for all cases examined, GMRES(10) always converges (see Section 4). However, the fact that the outer iteration operator is not positive definite in certain regimes emphasizes the need to find a preconditioner for the outer iterations.

4. NUMERICAL RESULTS

Numerical results were obtained iteratively and in all cases $\vec{\Psi}_0$ is used as the initial guess for $\vec{\Psi}$. Iteration proceeds until

$$\frac{\|\vec{\Psi}^{(m+1)}-\vec{\Psi}^{(m)}\|_2}{\|\vec{\Psi}_0\|_2} \leq \varepsilon \tag{8}$$

where ε is a small, user-defined convergence criterion. As can be seen in Table III, for both the simple scheme and GMRES(10), the two-step method always converged in fewer iterations than the one-step method. In most cases the iteration counts differed by one to two orders of magnitude. Similarly, in comparing the simple and GMRES(10) iterative schemes for one- and two-step, GMRES(10) always converged in fewer iterations than the simple scheme, and the counts differed by at least one order of magnitude in most cases. For both outer iteration schemes, the greatest differences were seen in the atomic mix-diffusion limit. For the inner iterations, GMRES(10) always converged in fewer iterations than traditional material or source iteration. Acceleration (preconditioning) decreased the iteration count in most cases.

Because GMRES(m) requires more work than simple iterative methods, and because a single two-step iteration requires more work than a one-step iteration, iteration count can be used to judge the effectiveness but not the efficiency of the various solution methods. We measured the computational work and relative efficiency by collecting floating point operation (FLOP) counts for the thirty-six test cases. The results are

displayed in Table IV. As can be seen, GMRES(10) always requires less computational effort than the simple iteration and it is always more efficient to accelerate (precondition) the inner iterations. The two-step scheme is not always more efficient than the one-step, but it is much more efficient in the atomic-mix regime. We found that for the cases in which the two-step method with GMRES(10) is slower, it is 2.5 times slower than the one-step method with GMRES(10) on average and it is 5.4 times slower at worst (when the material is *neither* atomically mixed nor diffusive). On the other hand, in cases where the one-step method with GMRES(10) is slower, it is 6.8 times slower than the two-step method GMRES(10) on average and nearly 24 times slower at worst (when the material is *both* atomically mixed and diffusive). We also note that case 4 corresponds to a highly diffusive and atomically mixed material while case 13 corresponds to a material that is neither diffusive nor atomically mixed. Thus, compared to the one-step method, the two-step method becomes increasingly effective as the material approaches the atomic mix-diffusion limit, and is extremely effective once in this limit, but it is less effective in materials that are neither diffusive nor atomically mixed.

Table III. Outer Iteration Counts ($\varepsilon = 10^{-6}$)

Case #	Simple		GMRES(10)		Case #	Simple		GMRES(10)	
	One-Step	Two-Step	One-Step	Two-Step		One-Step	Two-Step	One-Step	Two-Step
1	3637	347	150	17	19	1086	103	330	22
2	5991	1953	378	156	20	126	17	27	7
3	24770	3977	6559	203	21	229	58	46	11
4	199360	4523	40473	196	22	1117	103	427	21
5	936	94	85	15	23	8816	1412	2532	115
6	1381	452	127	60	24	1308	205	432	33
7	5364	857	1551	83	25	287	40	134	14
8	42792	964	8905	86	26	8817	1412	2539	116
9	117	16	25	9	27	1308	205	431	33
10	170	58	38	18	28	287	40	132	14
11	674	103	252	22	29	720	85	55	8
12	5442	114	1231	24	30	2075	583	114	19
13	19	6	9	5	31	13258	1315	4854	55
14	31	11	16	7	32	116346	1511	25329	76
15	154	17	91	8	33	521	49	46	6
16	1342	18	345	8	34	1768	450	113	11
17	120	17	26	9	35	12182	1114	4066	36
18	218	58	48	18	36	108093	1295	22786	55

Table IV. MFLOPS Counts ($\varepsilon = 10^{-6}$)

Case #	One-Step		Two-Step			
	Simple	GMRES(10)	Simple		GMRES(10)	
			Basic	Accelerated	Basic	Precond.
1	360.916290	26.6785793	4155.41113	597.989380	65.1575012	35.9485092
2	594.451904	67.3484192	6426.21191	3976.85400	334.790314	218.728699
3	2457.47900	1169.98901	40448.0781	6045.10303	1763.67297	293.194214
4	19778.1992	7266.76904	410404.094	4990.63379	6117.55420	304.579315
5	92.9554977	15.1098499	1032.17896	169.938004	46.3967896	31.9226799
6	137.103104	22.5619698	1453.44202	878.010803	133.021103	86.5515289
7	532.248474	276.782990	8715.21289	1311.42700	734.696106	126.872002
8	4245.40576	1598.60303	87148.5625	1130.72900	2745.72803	180.483704
9	11.7041397	4.40220404	101.378700	31.9402103	23.6467609	18.7313995
10	16.9621696	6.68456411	173.397705	115.708000	41.8453407	28.1793499
11	66.9629974	44.9920616	1041.27600	167.455002	204.527695	41.9601212
12	539.986694	219.633896	10269.3301	220.321503	829.377991	110.544998
13	1.98176003	1.59166396	13.7103004	10.2938499	8.30333042	8.58576584
14	3.17225599	2.76002789	27.5219498	21.8148403	16.4994297	12.0264101
15	15.3748398	16.1758099	175.516205	35.0615196	80.9392776	22.0135498
16	133.233902	61.4616203	1670.88696	104.960297	331.555206	92.3481216
17	12.0017700	4.55584478	65.3858032	22.7831497	18.5888805	12.9921799
18	21.7241497	8.48038101	200.730392	81.4227066	52.2051010	24.7933407
19	107.836700	58.9712486	1801.93298	164.310303	262.170990	48.1348801
20	12.5970201	4.72229624	64.8807526	22.8353500	14.7867804	11.2373199
21	22.8154392	8.18590927	200.782593	74.7462234	33.2801285	16.2323093
22	110.912102	76.1578522	1787.41199	140.317398	300.819489	39.8402214
23	874.714478	451.658508	14361.3496	2153.47510	1010.40198	171.354797
24	129.860901	77.0061188	2075.85498	322.209808	300.649109	57.2717094
25	28.5695000	23.8840504	404.423096	69.7607269	133.804794	30.7156391
26	874.813721	452.913300	14361.4297	2153.47510	1019.09100	172.751694
27	129.860901	76.8652878	2076.48291	322.209808	300.869293	56.9997902
28	28.5695000	23.6504192	404.815186	69.7607269	133.784698	30.4981098
29	71.5265732	9.91454792	630.284424	161.007507	23.0833797	18.4276409
30	205.953400	20.2539902	2007.50696	1073.84497	56.7671890	33.1353493
31	1315.39600	865.837891	22416.4492	1882.75806	834.929504	101.493797
32	11542.5498	4518.31006	240265.000	1707.54797	4236.63184	232.572296
33	51.7841797	8.18590927	340.734711	85.8660583	18.1486301	15.4801302
34	175.496597	20.1003494	1567.65100	785.656677	33.8015594	21.8378906
35	1208.64905	725.245972	20360.9492	1544.03699	592.642395	73.6422424
36	10723.7900	4064.65503	222279.797	1495.76599	3336.94702	222.839905

5. CONCLUSIONS

As expected, the new two-step scheme that we have proposed always converges in fewer iterations than the one-step scheme with for both simple iterations and GMRES(10). Acceleration (preconditioning) with S_2SA and AMSA improves the convergence rate of the inner iterations and results in better overall computational efficiency in most cases. The two-step method combined with a Krylov iterative method appears to be effective and efficient for problems that are atomically-mixed and diffusive, for which the simple one-step scheme is impractical. Development of a preconditioner for the outer iteration is essential in order to guarantee the convergence of restarted GMRES and remains an open issue.

REFERENCES

[1] C. D. Levermore, G. C. Pomraning, D. L. Sanzo and J. Wong. "Linear Transport Theory in a Random Medium." *Journal of Mathematical Physics*.**27**, 2526-2536 (1986).

[2] D. Vanderhaegen. "Radiative Transfer in Statistically Heterogeneous Mixtures." *Journal of Quantitative Spectroscopy and Radiative Transfer*. **36**, 557-561 (1986).

[3] T. S. Palmer. "A Coupled Diffusion Synthetic Acceleration for Binary Stochastic Mixture Transport Iterations in Slab Geometry." *ANS M&C Topical Meeting.* Avignon, France, September 12-15, 2005.

[4] A. K. Prinja and E. D. Fichtl. "Atomic Mix Synthetic Acceleration of Dose Computations in Binary Statistical Media." *ANS M&C Topical Meeting.* Avignon, France, September 12-15 (2005). To appear in *Nuclear Science and Engineering.*

[5] E. D. Fichtl, J. S. Warsa and A. K. Prinja. "Acceleration for Transport Iterations in Binary Statistical Media." *Proceedings of the American Nuclear Society.* Albuquerque, NM, November 12-16 (2006).

[6] L. J. Lorence, J. E. Morel and E. W. Larsen. "An S_2 Synthetic Acceleration Scheme for the One-Dimensional S_N Equations with Linear Discontinuous Spatial Differencing." *Nuclear Science and Engineering*. **101**, 341-351 (1989).

[7] B. Guthrie, J. P. Holloway and B. W. Patton. "GMRES as a Multi-Step Transport Sweep Accelerator." *Transport Theory and Statistical Physics*. **28**, 83-102 (1999).

[8] J. S. Warsa, T. A. Wareing and J. E. Morel. "Krylov Iterative Methods and the Degraded Effectiveness of Diffusion Synthetic Acceleration for Multidimensional S_N Calculations in Problems with Material Discontinuities." *Nuclear Science and Engineering*. **147**, 218-248 (2004).

[9] S. L. Campbell, I. C. F. Ipsen, C. T. Kelley and C. D. Meyer. "GMRES and the Minimal Polynomial." *BIT*. **36**, 664-675 (1996).

[10] Y. Saad. *Iterative Methods for Sparse Linear Systems.* PWS Publishing Company: Boston (1996).

[11] A. Greenbaum. *Iterative Methods for Solving Linear Systems.* SIAM: Philadelphia (1997).