Journal of Computational Physics 228 (2009) 8413–8426

Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

Krylov iterative methods and synthetic acceleration for transport in binary statistical media

Erin D. Fichtl [b,*], James S. Warsa [b], Anil K. Prinja [a]

[a] Department of Chemical and Nuclear Engineering, The University of New Mexico, Albuquerque, NM 87131-1341, United States
[b] Computational Physics and Methods, Los Alamos National Laboratory, Los Alamos, NM 87545-0001, United States

ARTICLE INFO

Article history:
Received 30 September 2008
Received in revised form 15 August 2009
Accepted 18 August 2009
Available online 23 August 2009

Keywords:
Radiation transport
Synthetic acceleration
Krylov iterative methods
Levermore–Pomraning closure
Stochastic media

ABSTRACT

In particle transport applications there are numerous physical constructs in which heterogeneities are randomly distributed. The quantity of interest in these problems is the ensemble average of the flux, or the average of the flux over all possible material 'realizations.' The Levermore–Pomraning closure assumes Markovian mixing statistics and allows a closed, coupled system of equations to be written for the ensemble averages of the flux in each material. Generally, binary statistical mixtures are considered in which there are two (homogeneous) materials and corresponding coupled equations. The solution process is iterative, but convergence may be slow as either or both materials approach the diffusion and/or atomic mix limits. A three-part acceleration scheme is devised to expedite convergence, particularly in the atomic mix-diffusion limit where computation is extremely slow. The iteration is first divided into a series of 'inner' material and source iterations to attenuate the diffusion and atomic mix error modes separately. Secondly, atomic mix synthetic acceleration is applied to the inner material iteration and S_2 synthetic acceleration to the inner source iterations to offset the cost of doing several inner iterations per outer iteration. Finally, a Krylov iterative solver is wrapped around each iteration, inner and outer, to further expedite convergence. A spectral analysis is conducted and iteration counts and computing cost for the new two-step scheme are compared against those for a simple one-step iteration, to which a Krylov iterative method can also be applied.

1. Introduction

The development of particle transport methods for stochastic mixtures is necessary for the numerous physical constructs in which heterogeneities are not distributed in predictable and ordered ways, but are indeed random. These include the distribution of the pebbles and the internal structure of the pebbles themselves in pebble bed reactors, the composition of concrete used in reactor shielding, tissue variability in radiation therapy, and cloud composition for global climate modeling. It is useful to describe the composition of stochastic media in a statistical sense – i.e., the sizes of individual 'chunks' of material are described by an appropriate probability density function. Since the material composition is known only statistically, the flux and other physical quantities are described in terms of an ensemble average, or the average over all possible 'realizations' of the material.

The ensemble average can be found by generating large numbers of realizations using the statistical characterization, conducting transport calculations for each realization, and averaging the result over all realizations. While these 'numerical experiments' are straightforward, they are extremely time-consuming. It is desirable instead to ensemble average the trans-

* Corresponding author. Tel.: +1 505 665 1768.
E-mail addresses: efichtl@lanl.gov (E.D. Fichtl), warsa@lanl.gov (J.S. Warsa), prinja@unm.edu (A.K. Prinja).

0021-9991/$ - see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2009.08.013

port equation directly, thereby yielding an equation or system of equations for the ensemble average of the flux. One such description was derived twenty years ago by Levermore, Pomraning and Vanderhaegen for binary, or two-state, mixtures [1,2]. The model employs the Levermore–Pomraning (LP) closure and consists of two coupled equations for the conditional ensemble average of the flux in the two materials present. Their formalism is widely used, although it is not exact in time-dependent and scattering regimes since it assumes that the transport process is Markovian in nature, which is true only in time-independent, purely absorbing materials. It has been shown analytically that the model does not achieve the correct diffusion limit when the material is also atomically mixed and can produce extremely inaccurate results for a particular class of problems in which one of the materials is a void and the other highly diffusive [3]. The authors showed that even when the chunk sizes are large, the atomic mix approximation, in which volume-averaged cross sections are used in the deterministic transport equation, is much more accurate than the LP closure in the diffusion limit when a weak volume source is placed in the diffusive material. For many diffusive problems, however, the LP closure is not significantly more inaccurate than atomic mix. For instance, in [4], for the same solid–void mixtures with an incident source instead of a volume source, LP was shown to be competitive with atomic mix in diffusive materials and significantly more accurate in mildly diffusive materials. Despite its shortcomings, LP is the only robust model that attempts to account for the stochastic structure of media with alternating materials. It is widely applied because it is a useful tool with which to study the effects of stochasticity in transport applications.

While solution of the coupled equations is considerably more efficient than conducting transport calculations for numerous material realizations, convergence can still be slow in diffusive and/or atomically mixed regimes. Recently, three distinct acceleration techniques have been devised to remedy different aspects of this problem. The first is a 'two-grid' acceleration [5], analogous to a two-grid acceleration developed by Adams and Morel for multigroup transport iterations [6]. The scheme was shown to significantly reduce the spectral radius, indicating more rapid convergence, and virtually eliminate the dominant error mode when the cross sections of the two materials were highly contrasting. However, the equation for the correction is itself a transport equation which must also be solved iteratively and must additionally be accelerated using diffusion synthetic acceleration in diffusive regimes. The second method 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 [7]. We note that neither of these methods were developed for or tested in materials that were both highly diffusive and atomically mixed. The third method is an atomic mix synthetic acceleration (AMSA) scheme which exploits the asymptotic atomic mix limit to calculate a lower-order correction for the flux and was applied to electron energy-loss straggling computations for which there is no scattering, hence no diffusion limit [8]. When scattering is present, the correction equations are dependent on the scalar flux at each iteration, necessitating an intermediate source iteration to converge the scalar flux. Furthermore, there is no clear way in which to accelerate convergence in the diffusion limit. In conclusion, a scheme has not yet been devised to accelerate convergence when the material is both diffusive and atomically mixed.

To this end, a three-part acceleration scheme is proposed for the LP equations. As with the schemes previously described, the method is explored for mono-energetic particles in one-dimensional planar geometry. Firstly, the outer iteration is divided into a series of 'inner' material and species iterations to attenuate the diffusion and atomic mix error modes separately. Secondly, AMSA is applied to the inner material iteration and S_2 synthetic acceleration (S_2SA) to the inner species iteration to offset the cost of doing multiple inner iterations per outer iteration. And finally, a Krylov iterative solver is wrapped around each iteration, inner and outer, to further accelerate convergence. Krylov iterative methods have been shown to be effective acceleration tools for transport iterations in deterministic materials, particularly when preconditioned with DSA [9]. A Krylov iterative method can also be applied to a simple one-step iteration, against which the two-step iteration is compared.

In Section 2, the problem setup and pertinent equations are presented. Section 3 contains a discussion of the one- and two-step iterative methods, synthetic acceleration schemes and solution by Krylov iterative methods. A spectral analysis is presented in Section 4, detailing pertinent aspects of the eigenvalue spectra of the various transport operators and their implications for convergence. Section 5 contains numerical results, specifically iteration counts and computing costs for each method. And finally, conclusions are presented in Section 6.

2. Governing equations

The Levermore–Pomraning (LP) equations for binary statistical mixtures describe the transport of particles through a medium in which two immiscible materials are randomly distributed according to homogeneous Markov mixing statistics. Taking the conditional ensemble average of the transport equation and applying the LP closure yields a coupled system of equations for the conditional ensemble average of the flux in each material, $\psi_1(s,\mu)$ and $\psi_2(s,\mu)$:

$$\left(\mu\frac{\partial}{\partial s}+\Sigma_1\right)\psi_1(s,\mu)=\frac{\Sigma_{s,1}}{2}\phi_1(s)+\frac{|\mu|}{\lambda_1}(\psi_2(s,\mu)-\psi_1(s,\mu))+Q_1(s,\mu) \tag{1a}$$

$$\left(\mu\frac{\partial}{\partial s}+\Sigma_2\right)\psi_2(s,\mu)=\frac{\Sigma_{s,2}}{2}\phi_2(s)+\frac{|\mu|}{\lambda_2}(\psi_1(s,\mu)-\psi_2(s,\mu))+Q_2(s,\mu) \tag{1b}$$

where s is the spatial variable, μ is the cosine of the angular variable θ and, for material ℓ, $\phi_\ell(s)$ is the conditional ensemble average of the scalar flux, Σ_ℓ and $\Sigma_{s,\ell}$ are the total and scattering cross sections and $Q_\ell(s,\mu)$ is a volume source. The equations

are coupled through the $\frac{|\mu|}{\lambda_\ell}(\psi_k(s,\mu)-\psi_\ell(s,\mu))$ term, which is introduced by the Levermore–Pomraning closure and describes the diffusion of particles across material interfaces. In one-dimensional geometry, the material can be thought of as alternating slabs of the two materials where the individual slabs of material ℓ have mean length λ_ℓ. Since the mixing statistics are assumed to be Markovian, the slab widths, x, are distributed exponentially:

$$P_\ell(x)=\frac{1}{\lambda_\ell}\exp\left(-\frac{x}{\lambda_\ell}\right),\quad \ell=1,2. \tag{2}$$

Given that s is in material ℓ, $\frac{ds}{\lambda_\ell}$ is the probability of transition into material k in the distance ds [10]. The total ensemble average of the flux is then defined to be:

$$\langle\psi\rangle=p_1\psi_1+p_2\psi_2$$

where p_ℓ is the probability that a point in the domain is in material ℓ and is given by

$$p_1=\frac{\lambda_1}{\lambda_1+\lambda_2},\quad p_2=1-p_1.$$

For numerical solution, a discrete ordinates, or S_N, angular discretization is applied in which the angular derivative is approximated using Gauss–Legendre quadrature with weights w_n and abscissas μ_n [11] to yield the following form of the S_N LP equations:

$$\widehat{L}_{1,n}\psi_{1,n}(s)=\frac{\Sigma_{s,1}}{4\pi}\phi_1(s)+\frac{|\mu_n|}{\lambda_1}\psi_{2,n}(s)+\frac{1}{4\pi}Q_{1,n}(s) \tag{3a}$$

$$\widehat{L}_{2,n}\psi_{2,n}(s)=\frac{\Sigma_{s,2}}{4\pi}\phi_2(s)+\frac{|\mu_n|}{\lambda_2}\psi_{1,n}(s)+\frac{1}{4\pi}Q_{2,n}(s) \tag{3b}$$

where

$$\widehat{L}_{\ell,n}=\left(\mu_n\frac{\partial}{\partial s}+\Sigma_\ell+\frac{|\mu_n|}{\lambda_\ell}\right),\quad \ell=1,2$$

and the scalar flux is approximated by $\langle\phi\rangle=p_1\phi_1+p_2\phi_2$ where

$$\phi_\ell(s)=2\pi\sum_{n=1}^{N}w_n\psi_{\ell,n}(s).$$

3. Iterative solution method

In large, multi-dimensional systems, direct inversion of the transport operator is prohibitively expensive, therefore iterative methods are the standard solution method. Iterative solution of the LP system is complicated by the fact that the equation for $\psi_{1,n}$ relies not only on ϕ_1, but on $\psi_{2,n}$ as well, and vice versa. Thus any fully iterative method must have a source iteration component – i.e., ϕ_1 and ϕ_2 must be lagged – and a material iteration component – i.e., ψ_2 must be lagged. There are numerous combinations of these components, but we intuit that each will be influenced by the two separate asymptotic limits associated with source and material iteration – the diffusion and atomic mix limits, respectively.

3.1. One-step method

The most straightforward way to iterate between these two equations is to lag unknown quantities at each iteration so that the angular fluxes can be updated. This iteration can be written as

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

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

where m is the iteration index. Since there is only one complete transport calculation per iteration, this iteration is dubbed the 'one-step' iteration. It is robust but converges slowly as the material approaches the atomic mix limit, $\lambda_\ell\Sigma_\ell\ll 1$, the diffusion limit, $c_\ell=\Sigma_{s,\ell}/\Sigma_\ell\approx 1$, or both limits. Ideally, an acceleration method could be devised for this iteration that attenuates both the diffusive and atomic mix error modes. However, this would require an update to ψ_2,ϕ_1 and ϕ_2 and there is no clear way in which to do this.

3.2. Two-step method

Alternately, each iteration could be separated into 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 iteration scheme (Eq. (4b)). Consider the following two-step scheme, with outer iteration index m, comprising two inner iterations that are conducted in series:

1. *Material iteration*
   Throughout this iteration, the scalar flux is held constant. The index i is assigned to the material iteration, which for outer iteration $m+1$ is given by:

$$\widehat{L}_{1,n}\psi_{1,n}^{(i+\frac{1}{2})} = \frac{\Sigma_{s,1}}{4\pi}\phi_1^{(m)} + \frac{|\mu_n|}{\lambda_1}\psi_{2,n}^{(i)} + Q_{1,n} \tag{5a}$$

$$\widehat{L}_{2,n}\psi_{2,n}^{(i+\frac{1}{2})} = \frac{\Sigma_{s,2}}{4\pi}\phi_2^{(m)} + \frac{|\mu_n|}{\lambda_2}\psi_{1,n}^{(i+\frac{1}{2})} + Q_{2,n} \tag{5b}$$

   where the values of $\phi_1^{(m)}$ and $\phi_2^{(m)}$ from the previous outer iteration are held constant. The scattering terms, $\Sigma_{s,1}\phi_1$ and $\Sigma_{s,2}\phi_2$, are held constant throughout the material iteration and therefore resemble isotropic volume sources. When no synthetic acceleration scheme is applied, $\psi_{1,n}^{(i+1)} = \psi_{1,n}^{(i+\frac{1}{2})}$ and $\psi_{2,n}^{(i+1)} = \psi_{2,n}^{(i+\frac{1}{2})}$. The result of this iteration is $\psi_{1,n}^{(m+\frac{1}{2})}$ and $\psi_{2,n}^{(m+\frac{1}{2})}$.
2. *Species iteration*
   Throughout this iteration, the coupling terms (i.e., $\frac{|\mu_n|}{\lambda_1}\psi_{2,n}^{(m+\frac{1}{2})}$ for the source iteration for material 1 and $\frac{|\mu_n|}{\lambda_2}\psi_{1,n}^{(m+\frac{1}{2})}$ for the source iteration for material 2) are held constant. The species iteration comprises two independent source iterations, one for each material. For outer iteration $m+1$, the source iterations for material 1, denoted by iteration index j, and material 2, denoted by iteration index k, are given by:

$$\widehat{L}_{1,n}\psi_{1,n}^{(j+\frac{1}{2})} = \frac{\Sigma_{s,1}}{4\pi}\phi_1^{(j)} + \frac{|\mu_n|}{\lambda_1}\psi_{2,n}^{(m+\frac{1}{2})} + Q_{1,n} \tag{6}$$

$$\widehat{L}_{2,n}\psi_{2,n}^{(k+\frac{1}{2})} = \frac{\Sigma_{s,2}}{4\pi}\phi_2^{(k)} + \frac{|\mu_n|}{\lambda_2}\psi_{1,n}^{(m+\frac{1}{2})} + Q_{2,n}. \tag{7}$$

   The coupling terms are held constant throughout this iteration and therefore resemble angularly dependent volume sources. As in the material iteration, when no synthetic acceleration scheme is applied, $\psi_{1,n}^{(j+1)} = \psi_{1,n}^{(j+\frac{1}{2})}$ and $\psi_{2,n}^{(k+1)} = \psi_{2,n}^{(k+\frac{1}{2})}$. The results of this iteration are $\phi_1^{(m+1)}$ and $\phi_2^{(m+1)}$, which are then used to compute the 'sources' for the next material iteration.

Since each two-step iteration requires two inner iterations, with the species iteration comprising two 'sub-inner' iterations, while each one-step iteration requires only a single sweep through the coupled LP equations in which $\widehat{L}_{1,n}$ and $\widehat{L}_{2,n}$ are inverted, each two-step iteration will be more expensive than a one-step iteration. Fortunately, the structure of the inner iterations easily allows for the use of known synthetic acceleration schemes which should offset the cost of doing multiple inner iterations per outer iteration. Specifically, the material iteration is accelerated with atomic mix synthetic acceleration (AMSA) [8] and the species iteration is accelerated using S_2 synthetic acceleration (S_2SA) [12].

3.3. Synthetic acceleration

Synthetic acceleration schemes quickly generate corrections for the problem unknown at the end of each iteration by computing a low-order estimate of its error. While calculation of the correction increases the computational cost of each iteration, a suitable synthetic acceleration scheme will generally reduce the number of required iterations considerably, thereby reducing the required computational effort overall. For source iteration, S_2 synthetic acceleration provides a suitable update for the scalar flux by approximating its error using S_2 discrete ordinates. Computing the error estimate for ϕ at the end of each inner source iteration for the two-step method involves solving a system of coupled equations:

$$\widehat{L}_{\ell,1}f_{\ell,1}^{(j+\frac{1}{2})} = \frac{\Sigma_{s,1}}{4\pi}\left(F_\ell^{(j+\frac{1}{2})} + \phi_\ell^{(j+\frac{1}{2})} - \phi_\ell^{(j)}\right) \tag{8a}$$

$$\widehat{L}_{\ell,2}f_{\ell,2}^{(j+\frac{1}{2})} = \frac{\Sigma_{s,2}}{4\pi}\left(F_\ell^{(j+\frac{1}{2})} + \phi_\ell^{(j+\frac{1}{2})} - \phi_\ell^{(j)}\right) \tag{8b}$$

where $\widehat{L}_{\ell,1}$, $\widehat{L}_{\ell,2}$, $f_{\ell,1}$ and $f_{\ell,2}$ are given for the two S_2 quadrature angles. The S_2SA approximation of the error in $\phi_\ell^{(j+\frac{1}{2})}$ is given by

$$F_\ell^{(j+\frac{1}{2})} = f_{\ell,1}^{(j+\frac{1}{2})} + f_{\ell,2}^{(j+\frac{1}{2})} \approx \phi - \phi_\ell^{(j+\frac{1}{2})}$$

where ϕ is the exact solution. The scalar flux can then be updated before proceeding to the next iteration:

$$\phi_\ell^{(j+1)} = \phi_\ell^{(j+\frac{1}{2})} + F_\ell^{(j+\frac{1}{2})}. \tag{9}$$

The scheme is most efficient in diffusive regimes and is well documented [13,12].

In order to provide a suitable update for ψ_2 at the end of each inner material iteration, the asymptotic atomic mix limit is employed. In this limit, the mean chord lengths of the materials, λ_ℓ, are much smaller than the mean free paths of the

traversing particles, to which the material appears to be homogeneous – i.e., $\lambda_\ell \Sigma_\ell \ll 1$ where the mean free path is defined to be the inverse of Σ_ℓ. Equations for the error in the angular flux, ψ, for the material iteration (Eq. (5b)) are then given by:

$$\widehat{L}_{1,n} E_{1,n}^{(i+\frac{1}{2})} = \frac{|\mu_n|}{\lambda_1}\left(E_{2,n}^{(i+\frac{1}{2})} + \psi_{2,n}^{(i+\frac{1}{2})} - \psi_{2,n}^{(i)}\right) \tag{10a}$$

$$\widehat{L}_{2,n} E_{2,n}^{(i+\frac{1}{2})} = \frac{|\mu_n|}{\lambda_2} E_{1,n}^{(i+\frac{1}{2})} \tag{10b}$$

where, in material ℓ, $E_{\ell,n}^{(i+\frac{1}{2})} = \psi_{\ell,n} - \psi_{\ell,n}^{(i+\frac{1}{2})}$ and $\psi_{\ell,n}$ is the *exact* value of the angular flux for quadrature angle n.

Introducing a scaling parameter, $\epsilon \ll 1$, the mean chord lengths can be rewritten as $\lambda_\ell = \epsilon\tilde{\lambda}_\ell$ where $\tilde{\lambda}_\ell \sim O(1)$ and the errors, $E_{\ell,n}$, and fluxes, $\psi_{\ell,n}$, can be rewritten as series expansions in powers of ϵ:

$$E_{\ell,n} = \sum_{m=0}^{\infty} \epsilon^m E_{\ell,n}^{(m)} \tag{11a}$$

$$\psi_{\ell,n} = \sum_{m=0}^{\infty} \epsilon^m \psi_{\ell,n}^{(m)} \tag{11b}$$

where the superscript (m) in this case represents the mth expansion coefficient. Substituting into Eq. (10b) yields:

$$\left(\mu_n \frac{\partial}{\partial s} + \Sigma_1\right) \sum_{m=0}^{\infty} \epsilon^{m+1} E_{1,n}^{(i+\frac{1}{2})}, (m) = \frac{|\mu_n|}{\tilde{\lambda}_1} \sum_{m=0}^{\infty} \epsilon^m \left(E_{2,n}^{(i+\frac{1}{2}),(m)} - E_{1,n}^{(i+\frac{1}{2}),(m)} + \psi_{\ell,2}^{(i+\frac{1}{2}),(m)} - \psi_{\ell,2}^{(i),(m)}\right) \tag{12a}$$

$$\left(\mu_n \frac{\partial}{\partial s} + \Sigma_2\right) \sum_{m=0}^{\infty} \epsilon^{m+1} E_{2,n}^{(i+\frac{1}{2}),(m)} = \frac{|\mu_n|}{\tilde{\lambda}_2} \sum_{m=0}^{\infty} \epsilon^m \left(E_{1,n}^{(i+\frac{1}{2}),(m)} - E_{2,n}^{(i+\frac{1}{2}),(m)}\right). \tag{12b}$$

To lowest order, $\mathcal{O}(1)$, Eq. (12b) yields

$$E_{1,n}^{(i+\frac{1}{2}),(0)} = E_{2,n}^{(i+\frac{1}{2}),(0)}, \tag{13}$$

From Eq. (12a) we can then conclude that

$$\psi_{\ell,2}^{(i+\frac{1}{2}),(0)} - \psi_{\ell,2}^{(i),(0)} = 0. \tag{14}$$

To first order, $\mathcal{O}(\epsilon)$, the result is:

$$\left(\mu_n \frac{\partial}{\partial s} + \Sigma_1\right) E_{1,n}^{(i+\frac{1}{2}),(0)} = \frac{|\mu_n|}{\tilde{\lambda}_1}\left(E_{2,n}^{(i+\frac{1}{2}),(1)} - E_{1,n}^{(i+\frac{1}{2}),(1)} + \psi_{\ell,2}^{(i+\frac{1}{2}),(1)} - \psi_{\ell,2}^{(i),(1)}\right) \tag{15a}$$

$$\left(\mu_n \frac{\partial}{\partial s} + \Sigma_2\right) E_{2,n}^{(i+\frac{1}{2}),(0)} = \frac{|\mu_n|}{\tilde{\lambda}_2}\left(E_{1,n}^{(i+\frac{1}{2}),(1)} - E_{2,n}^{(i+\frac{1}{2}),(1)}\right). \tag{15b}$$

Multiplying Eq. (15a) by p_1 and Eq. (15b) by p_2, applying the relationship in Eq. (13) and summing yields:

$$\left(\mu_n \frac{\partial}{\partial s} + \langle\Sigma\rangle\right) E_{2,n}^{(i+\frac{1}{2}),(0)} = \frac{|\mu_n|}{\tilde{\lambda}_1 + \tilde{\lambda}_2}\left(\psi_{2,n}^{(i+\frac{1}{2}),(1)} - \psi_{2,n}^{(i),(1)}\right). \tag{16}$$

As can be seen, if both the numerator and denominator on the right-hand-side of this equation are multiplied by ϵ and the relationship in Eq. (14) applied, the final AMSA approximation for $\psi_{2,n}^{(i+\frac{1}{2})}$ is found to be:

$$\left(\mu_n \frac{\partial}{\partial s} + \langle\Sigma\rangle\right) E_{2,n}^{(i+\frac{1}{2})} = \frac{|\mu_n|}{\lambda_1 + \lambda_2}\left(\psi_{2,n}^{(i+\frac{1}{2})} - \psi_{2,n}^{(i)}\right). \tag{17}$$

As can be seen in Eq. (17), $E_{2,n}^{(i+\frac{1}{2})}$ can be obtained by conducting a single transport sweep to invert $(\mu_n \frac{\partial}{\partial s} + \langle\Sigma\rangle)$. This scheme was shown to be effective for electron energy-loss straggling calculations [8], for which there is no angular dependence, but this marks its first application in scattering regimes.

3.4. Operator notation

For numerical solution, a discrete ordinates representation has already been applied to the angular variable. A linear discontinuous finite element discretization is now applied to the spatial variable. All matrices used in this section are written out explicitly in Appendix. The discretized form of the one-step outer iteration can then be written in matrix notation as:

$$\mathbf{A}\vec{\psi}^{(m+1)} = \mathbf{B}\vec{\phi}^{(m)} + \mathbf{C}\vec{\psi}^{(m)} + \vec{Q} \tag{18}$$

where, for outer iteration $m+1$, Eq. (18) represents Eqs. (4a) and (4b) and can be rewritten as:

$$\vec{\psi}^{(m+1)} = \mathbf{T_1}\vec{\psi}^{(m)} + \vec{q}_0 \tag{19}$$

where $\mathbf{D}\vec{\psi} = \vec{\phi}$,

$$\mathbf{T_1} = \mathbf{A}^{-1}(\mathbf{BD} + \mathbf{C})$$

and

$$\vec{q}_0 = \mathbf{A}^{-1}\vec{Q}$$

which physically represents the uncollided angular flux and can be computed before iteration begins.

Likewise, the two-step outer iteration can be written in matrix notation as:

$$\mathbf{A_m}\vec{\psi}^{(m+\frac{1}{2})} = \mathbf{B_m}\vec{\phi}^{(m)} + \mathbf{C_m}\vec{\psi}^{(m+\frac{1}{2})} + \vec{Q}_m \tag{20a}$$

$$\mathbf{A_s}\vec{\psi}^{(m+1)} = \mathbf{B_s}\vec{\phi}^{(m+1)} + \mathbf{C_s}\vec{\psi}^{(m+\frac{1}{2})} + \vec{Q}_s \tag{20b}$$

where for outer iteration $m+1$, Eq. (20a) represents the unaccelerated material iteration (Eqs. (5a) and (5b)) and Eq. (20b) represents the unaccelerated species iteration (Eq. (7) for $\ell = 1, 2$). Collapsing Eq. (20b) into a single matrix equation for $\vec{\phi}^{(m)}$ and $\vec{\phi}^{(m+1)}$ yields:

$$\vec{\phi}^{(m+1)} = \mathbf{T_2}\vec{\phi}^{(m)} + \vec{q}_0 \tag{21}$$

where

$$\mathbf{T_2} = (\mathbf{I} - \mathbf{D}\mathbf{A_s^{-1}}\mathbf{B_s})^{-1}\mathbf{D}\mathbf{A_s^{-1}}\mathbf{C_s}\left(\mathbf{I} - \mathbf{A_m^{-1}}\mathbf{C_m}\right)^{-1}\mathbf{A_m^{-1}}\mathbf{B_m}$$

and

$$\vec{q}_0 = (\mathbf{I} - \mathbf{D}\mathbf{A_s^{-1}}\mathbf{B_s})^{-1}\mathbf{D}\mathbf{A_s^{-1}}(\vec{Q}_s + \mathbf{C_s}\left(\mathbf{I} - \mathbf{A_m^{-1}}\mathbf{C_m}\right)^{-1}\mathbf{A_m^{-1}}\vec{Q}_m),$$

which physically represents the uncollided scalar flux.

Rewriting the accelerated inner material iteration (Eqs. (5b) and (17)) in matrix notation using the matrices from Eq. (20a) yields:

$$\vec{\psi}^{(i+\frac{1}{2})} = \mathbf{T_m}\vec{\psi}^{(i)} + \vec{q}_m \tag{22a}$$

$$\vec{e}^{(i+\frac{1}{2})} = \mathbf{E_m}\left(\vec{\psi}^{(i+\frac{1}{2})} - \vec{\psi}^{(i)}\right) \tag{22b}$$

$$\vec{\psi}^{(i+1)} = \vec{\psi}^{(i+\frac{1}{2})} + \vec{E}^{(i+\frac{1}{2})} \tag{22c}$$

where

$$\mathbf{T_m} = \mathbf{A_m^{-1}}\mathbf{C_m},$$

which represents the inversion of the appropriate transport operator, and

$$\vec{q}_m = \mathbf{A_m^{-1}}\mathbf{B_m}\vec{\phi}^{(m)},$$

which does not physically represent the uncollided flux in this case, but is an invariant source term since $\vec{\phi}^{(m)}$ is held constant during the material iteration. Note that $\vec{Q}_m$ is *not* included in $\vec{q}_m$ since it is contained in $\vec{q}_0$ (Eq. (21)). The matrix $\mathbf{E_m}$ represents the inversion of the AMSA transport operator and is the lower-order atomic mix approximation of $(\mathbf{I} - \mathbf{T_m})^{-1}\mathbf{T_m}$.

Likewise, the inner species iteration (Eq. (7)) can be rewritten in matrix notation using the matrices from Eq. (20b) after first rewriting them in block form

$$\mathbf{M_s} = \begin{bmatrix} \mathbf{M_{s,1}} & 0 \\ 0 & \mathbf{M_{s,2}} \end{bmatrix}.$$

The species iteration is then

$$\vec{\phi}_1^{(j+\frac{1}{2})} = \mathbf{T_{s,1}}\vec{\phi}_1^{(j)} + \vec{q}_{s,1} \qquad \vec{\phi}_2^{(k+\frac{1}{2})} = \mathbf{T_{s,2}}\vec{\phi}_2^{(k)} + \vec{q}_{s,2} \tag{23a}$$

$$\vec{F}_1^{(j+\frac{1}{2})} = \mathbf{E_{s,1}}\left(\vec{\phi}_1^{(j+\frac{1}{2})} - \vec{\phi}_1^{(j)}\right) \qquad \vec{F}_2^{(k+\frac{1}{2})} = \mathbf{E_{s,2}}\left(\vec{\phi}_2^{(k+\frac{1}{2})} - \vec{\phi}_2^{(k)}\right) \tag{23b}$$

$$\vec{\phi}_1^{(j+1)} = \vec{\phi}_1^{(j+\frac{1}{2})} + \vec{F}_1^{(j+\frac{1}{2})} \qquad \vec{\phi}_2^{(k+1)} = \vec{\phi}_2^{(k+\frac{1}{2})} + \vec{F}_2^{(k+\frac{1}{2})} \tag{23c}$$

where

$$\mathbf{T_{s,1}} = \mathbf{D}\mathbf{A_{s,1}^{-1}}\mathbf{B_{s,1}} \quad \text{and} \quad \mathbf{T_{s,2}} = \mathbf{D}\mathbf{A_{s,2}^{-1}}\mathbf{B_{s,2}}$$

and

$$\vec{q}_{s,1} = \mathbf{D}\mathbf{A_{s,1}^{-1}}\mathbf{C_{s,1}}\vec{\psi}_1^{(m+\frac{1}{2})} \quad \text{and} \quad \vec{q}_{s,2} = \mathbf{D}\mathbf{A_{s,2}^{-1}}\mathbf{C_{s,2}}\vec{\psi}_2^{(m+\frac{1}{2})}.$$

Table 1
Spectral radii for a 10 cm slab with isotropic incidence on the left boundary and a reflecting boundary condition at the right ($I = 100$ spatial cells, $N = 8$ quadrature angles, $\Sigma_1 = 5.0$ cm^{-1}, $\Sigma_2 = 1.0$ cm^{-1}, $\tilde{\lambda}_1 = 1.0$ cm and $\tilde{\lambda}_2 = 5.0$ cm).

	ϵ_λ							
	1.0		0.1		0.01		0.001	
c	1S	2S	1S	2S	1S	2S	1S	2S
1.0	0.9967	0.9623	0.998	0.9935	0.9995	0.9969	0.9999	0.9972
0.99	0.9872	0.8671	0.9914	0.9722	0.9979	0.9855	0.9998	0.9871
0.9	0.9017	0.4359	0.9314	0.8008	0.9833	0.8848	0.9981	0.8962
0.5	0.5187	0.08109	0.6634	0.3138	0.9275	0.4657	0.992	0.495
	MI	AMI	MI	AMI	MI	AMI	MI	AMI
1.0	0.004605	0.03615	0.2193	0.1165	0.8135	0.1216	0.9793	0.04596
0.99	0.004605	0.03615	0.2193	0.1165	0.8135	0.1216	0.9793	0.04596
0.9	0.004605	0.03615	0.2193	0.1165	0.8135	0.1216	0.9793	0.04596
0.5	0.004605	0.03615	0.2193	0.1165	0.8135	0.1216	0.9793	0.04596
	SI1	ASI1	SI1	ASI1	SI1	ASI1	SI1	ASI1
1.0	0.9102	0.2091	0.5387	0.1482	0.1231	0.04714	0.01449	0.005957
0.99	0.9011	0.2059	0.5333	0.146	0.1218	0.04663	0.01435	0.005897
0.9	0.8192	0.1789	0.4848	0.1275	0.1108	0.04206	0.01304	0.005356
0.5	0.4551	0.08562	0.2693	0.06154	0.06154	0.02259	0.007246	0.002966
	SI2	ASI2	SI2	ASI2	SI2	ASI2	SI2	ASI2
1.0	0.9056	0.2107	0.5381	0.1488	0.1231	0.04714	0.01449	0.005957
0.99	0.8966	0.2075	0.5327	0.1467	0.1218	0.04662	0.01435	0.005897
0.9	0.8151	0.181	0.4843	0.1284	0.1108	0.04206	0.01304	0.005356
0.5	0.4528	0.08787	0.269	0.06266	0.06154	0.02259	0.007246	0.002966

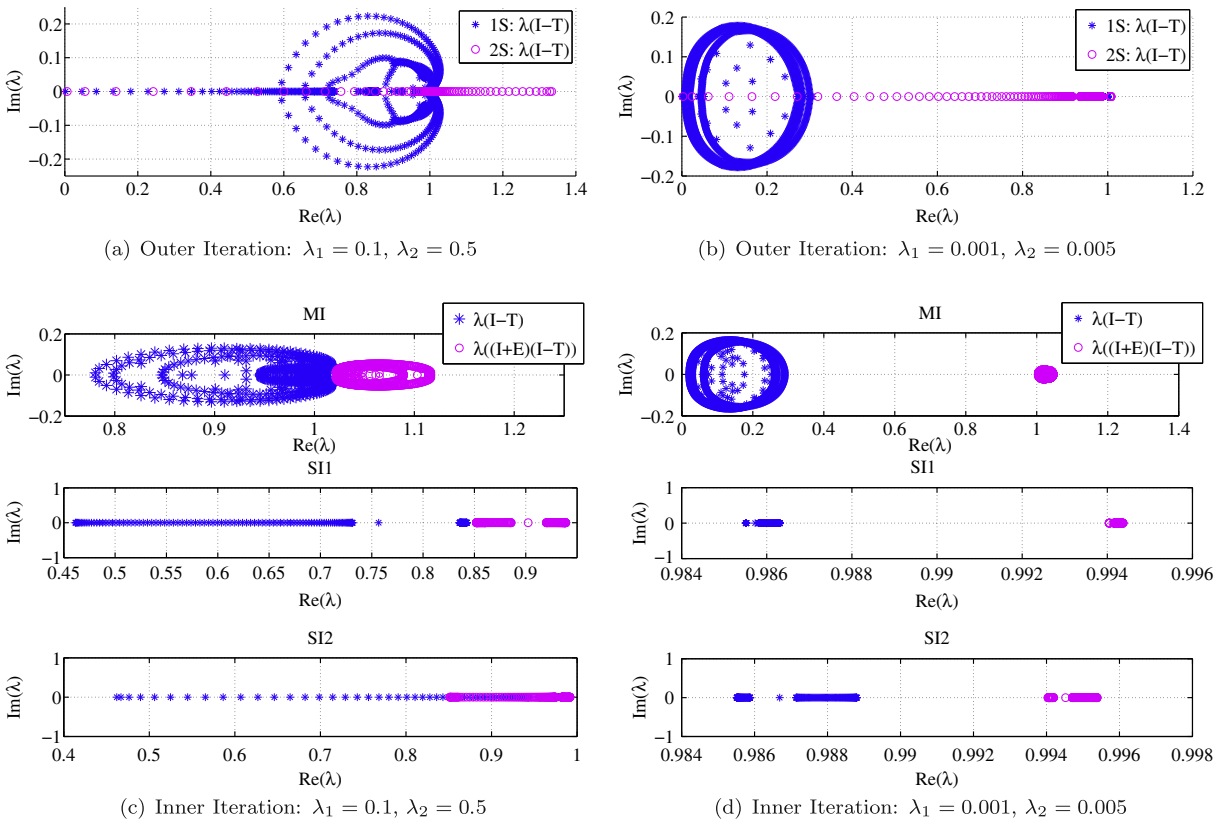

(a) Outer Iteration: $\lambda_1 = 0.1$, $\lambda_2 = 0.5$

(b) Outer Iteration: $\lambda_1 = 0.001$, $\lambda_2 = 0.005$

(c) Inner Iteration: $\lambda_1 = 0.1$, $\lambda_2 = 0.5$

(d) Inner Iteration: $\lambda_1 = 0.001$, $\lambda_2 = 0.005$

Fig. 1. Spectrum of (**I–T**) for a 10 cm slab with isotropic incidence on the left boundary and a reflecting boundary condition at the right ($I = 100, N = 8, \Sigma_1 = 5.0$ cm^{-1}, $\Sigma_2 = 1.0$ cm^{-1} and $c = 1.0$).

Once again, note that $\vec{q}_{s,1}$ and $\vec{q}_{s,2}$ are not uncollided fluxes, but invariant source terms containing $\vec{\psi}_1^{(m+\frac{1}{2})}$ and $\vec{\psi}_2^{(m+\frac{1}{2})}$, respectively, and $\vec{Q}_s$ is not included in $\vec{q}_s$ since it is already contained in $\vec{q}_0$. The matrix $\mathbf{E_s}$ represents the inversion of the S_2SA transport operator and is the lower-order S_2 approximation of $(\mathbf{I}-\mathbf{T_s})^{-1}\mathbf{T_s}$.

3.5. Krylov iterative methods

While the two-step outer iteration scheme is expected to converge in fewer iterations than the one-step scheme, the spectral radius is still very large in the atomic mix-diffusion regime, i.e., in the regime where $c_\ell \approx 1$ and $\lambda_\ell \Sigma_\ell \ll 1$ (see Section 4). In order to accelerate convergence, a Krylov iterative method, such as the restarted Generalized Minimal RESidual (GMRES(n)) method, which was designed for use with non-symmetric operators such as the transport operators, could be 'wrapped around' the existing algorithm to accelerate convergence. This approach has shown excellent acceleration when used in place of traditional source iteration, particularly when preconditioned with DSA [9].

Krylov iterative solvers solve the matrix equation $\mathbf{A}\vec{x} = \vec{b}$ or the left preconditioned matrix equation $\mathbf{M^{-1}A}\vec{x} = \mathbf{M^{-1}}\vec{b}$. In order to solve our system using a Krylov solver, Eqs. (19), (21), (22a) and (23a) are rewritten in the general form as

$$(\mathbf{I}-\mathbf{T})\vec{\Psi} = \vec{\Psi}_0. \tag{24}$$

Multiplication by $\mathbf{T}$ represents the inversion of the appropriate transport operator and its action is effected by conducting a transport sweep. $\mathbf{I}$ is the appropriately sized identity matrix and $\vec{\Psi}_0$ represents the appropriate uncollided flux or source vector [14]. Since a single transport sweep calculates the product $\mathbf{T}\vec{\Psi}$, the algorithm is very simply modified to output $(\mathbf{I}-\mathbf{T})\vec{\Psi}$, the result of which is returned to the Krylov solver.

For the inner iterations, it is possible to left precondition the system using AMSA and S_2SA. The synthetic acceleration update for inner iteration i can be written in general as:

$$\vec{\Psi}^{(i+1)} = \vec{\Psi}^{(i+\frac{1}{2})} + \mathbf{E}(\vec{\Psi}^{(i+\frac{1}{2})} - \vec{\Psi}^{(i)}) \tag{25}$$

where $\vec{\Psi}^{(i+\frac{1}{2})} = \mathbf{T}\vec{\Psi}^{(i)} + \vec{\Psi}_0$ and $\mathbf{E}$ is the appropriate synthetic acceleration approximation to $(\mathbf{I}-\mathbf{T})^{-1}\mathbf{T}$. Thus, in matrix notation, the accelerated systems for the inner iterations take the form:

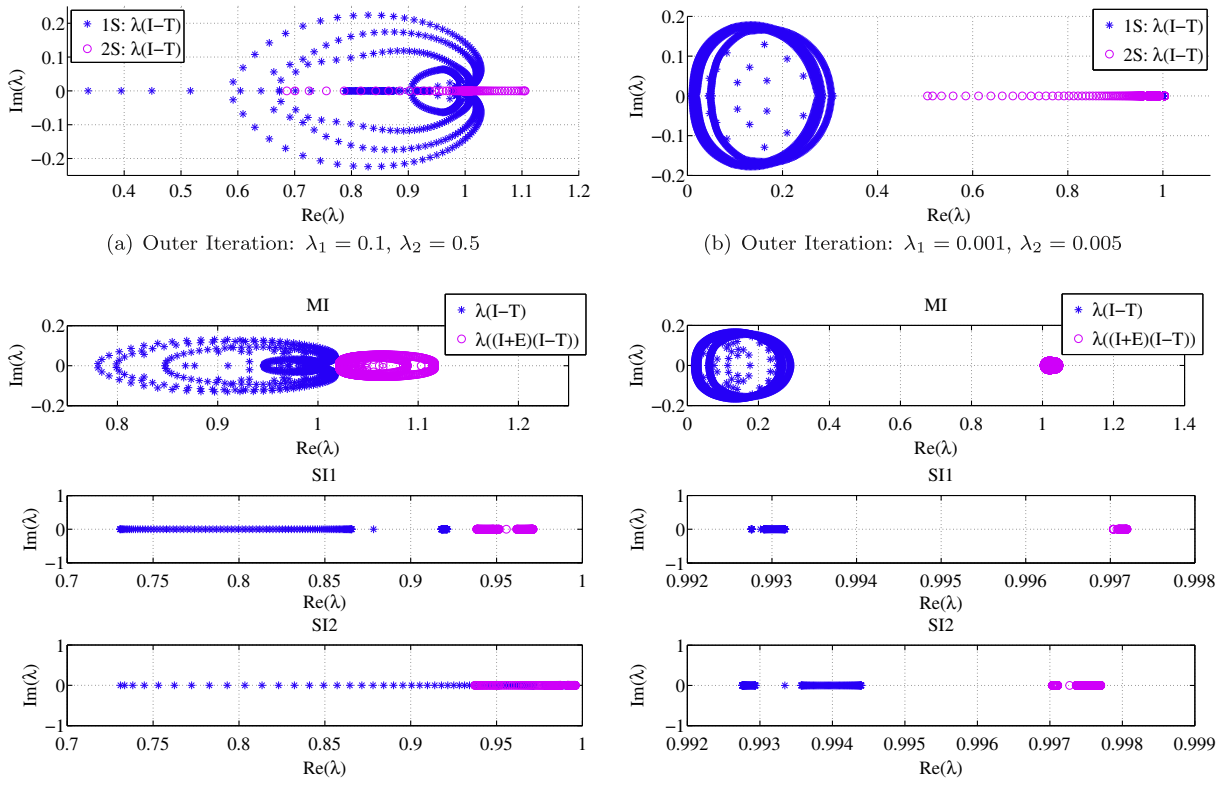

(a) Outer Iteration: $\lambda_1 = 0.1$, $\lambda_2 = 0.5$

(b) Outer Iteration: $\lambda_1 = 0.001$, $\lambda_2 = 0.005$

(c) Inner Iteration: $\lambda_1 = 0.1$, $\lambda_2 = 0.5$

(d) Inner Iteration: $\lambda_1 = 0.001$, $\lambda_2 = 0.005$

Fig. 2. Spectrum of (**I–T**) for a 10 cm slab with isotropic incidence on the left boundary and a reflecting boundary condition at the right ($I = 100, N = 8, \Sigma_1 = 5.0\ \text{cm}^{-1}, \Sigma_2 = 1.0\ \text{cm}^{-1}$ and $c = 0.5$).

$$\vec{\Psi}^{(i+1)} = \mathbf{T}\vec{\Psi}^{(i)} + \vec{\Psi}_0 + \mathbf{E}(\mathbf{T}\vec{\Psi}^{(i)} + \vec{\Psi}_0 - \vec{\Psi}^{(i)}). \tag{26}$$

Grouping terms yields:

$$\vec{\Psi}^{(i+1)} = (\mathbf{I} - (\mathbf{I}+\mathbf{E})(\mathbf{I}-\mathbf{T}))\vec{\Psi}^{(i)} + (\mathbf{I}+\mathbf{E})\vec{\Psi}_0. \tag{27}$$

Eq. (27) represents one possible iteration scheme, namely Richardson iteration, applied to the matrix equation:

$$(\mathbf{I}+\mathbf{E})(\mathbf{I}-\mathbf{T})\vec{\Psi} = (\mathbf{I}+\mathbf{E})\vec{\Psi}_0. \tag{28}$$

Clearly, if the equation is to be solved using a Krylov iterative method, (**I** + **E**) can be regarded as a preconditioner. Since the synthetic acceleration algorithms compute the operation of **E** on the supplied residual, they are easily modified to return the operation of (**I** + **E**) on an input vector. 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 $\lambda_\ell \Sigma_\ell \ll 1$.

4. Spectral analysis

4.1. Simple iterative schemes

The spectral radius, ρ, of the operator **T** is indicative of the convergence rate of the simple iterative schemes, with larger ρ indicating slower convergence. The spectral radii of each operator are shown in Table 1 for the one-step (1S) and two-step (2S) schemes, as well as the inner material iteration (MI) and source iterations (SI1 and SI2) and accelerated material (AMI) and source (ASI1 and ASI2) iterations. Note that the spectral radius given for the outer two-step iteration assumes fully converged inner material and species iterations. Accelerating these inner iterations does not alter their final result, and therefore does not affect the spectral radius of the outer two-step iteration. Results are shown for various scattering ratios, c, and values of ϵ_λ, where $\lambda = \epsilon_\lambda \tilde{\lambda}$ and $\tilde{\lambda} \sim \mathcal{O}(1)$.

Table 2
Simple iteration: outer iteration counts for the one-step iteration and two-step iteration with accelerated inner iterations (10 cm slab, isotropic incidence on the left boundary and reflecting boundary condition on the right, $I = 100, N = 8, \tilde{\Sigma}_1 = 5.0$ cm^{-1}, $\tilde{\Sigma}_2 = 1.0$ cm^{-1}, $\tilde{\lambda}_1 = 1.0$ cm and $\tilde{\lambda}_2 = 5.0$ cm). Unless otherwise specified, $c = 0.99$ and $\epsilon_\lambda = 0.01$.

	ϵ_λ							
	1.0		0.1		0.01		0.001	
c	1S	2S	1S	2S	1S	2S	1S	2S
1.0	3637	347	5991	1953	24770	3977	199360	4523
0.99	936	94	1381	452	5364	857	42792	964
0.9	117	16	170	58	674	103	5442	114
0.5	19	6	31	11	154	17	1342	18
	ϵ_{λ_1}							
	1.0		0.1		0.01		0.001	
ϵ_{λ_2}	1S	2S	1S	2S	1S	2S	1S	2S
1.0	936	94	837	110	816	85	813	79
0.1	1144	156	1381	452	1456	403	1453	362
0.01	1189	130	2021	579	5364	857	7178	699
0.001	1200	122	2129	554	10077	1006	42792	964
	c_1							
	1.0		0.99		0.9		0.5	
c_2	1S	2S	1S	2S	1S	2S	1S	2S
1.0	24770	3977	8816	1412	1308	205	287	40
0.99	8817	1412	5364	857	1194	186	281	39
0.9	1308	205	1194	186	674	103	241	32
0.5	287	40	281	39	240	32	154	17
	ϵ_{Σ_1}							
	10.0		1.0		0.1		0.01	
ϵ_{Σ_2}	1S	2S	1S	2S	1S	2S	1S	2S
10.0	1590	526	2063	584	2153	550	2163	545
1.0	2029	581	5364	857	6307	621	6320	573
0.1	2111	546	6188	610	3481	95	2639	49
0.01	2120	541	6179	562	2635	49	1531	13

As can be seen, in all cases, the spectral radii for the 2S operator are smaller than those for the 1S operator. For both the 1S and 2S schemes, the spectral radii are smallest when the materials are *not* atomically mixed or diffusive, larger when they are either atomically mixed *or* diffusive and largest when they are atomically mixed *and* diffusive. The spectral radii for the MI are large when ϵ_λ is small, but are unaffected by the magnitude of c. As demonstrated previously [8], the addition of AMSA is particularly effective at reducing the spectral radius when the material is atomically mixed while increasing it slightly when the material is far from atomic mix. The spectral radii for SI1 and SI2, on the other hand, are affected by both ϵ_λ and c. This can be seen by examining the operator

$$\widehat{L}_{\ell,n} = \mu_n \frac{\partial}{\partial s} + \Sigma_\ell + \frac{|\mu_n|}{\lambda_\ell}$$

in Eq. (7). In effect, the material always looks less diffusive than it actually is because the *effective* total cross section is $\Sigma_\ell + \frac{|\mu_n|}{\lambda_\ell}$. Thus, the smaller λ_ℓ, the larger the effective total cross section and the less diffusive the material appears. 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 (ASI1 and ASI2) always effectively reduces the spectral radius.

4.2. GMRES

There are several properties of an operator that influence the choice of 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 non-symmetric. However, in general the eigenvalue spectrum of the matrix (**I**–**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 [15]. Convergence is fastest for small clusters centered near unity and is slowest for large clusters or clusters centered near zero. Spectra for some representative problems are examined in Figs. 1 and 2.

Table 3
GMRES(10): outer iteration counts for the one-step iteration and two-step iteration with accelerated inner iterations (10 cm slab, isotropic incidence on the left boundary and reflecting boundary condition on the right, $I = 100, N = 8, \widetilde{\Sigma}_1 = 5.0\ \text{cm}^{-1}, \widetilde{\Sigma}_2 = 1.0\ \text{cm}^{-1}, \tilde{\lambda}_1 = 1.0$ cm and $\tilde{\lambda}_2 = 5.0$ cm). Unless otherwise specified, $c = 0.99$ and $\epsilon_\lambda = 0.01$.

	ϵ_λ							
	1.0		0.1		0.01		0.001	
c	1S	2S	1S	2S	1S	2S	1S	2S
1.0	150	17	378	156	6559	203	40473	196
0.99	85	15	127	60	1551	83	8905	86
0.9	25	9	38	18	252	22	1231	24
0.5	9	5	16	7	91	8	345	8
	ϵ_{λ_1}							
	1.0		0.1		0.01		0.001	
ϵ_{λ_2}	1S	2S	1S	2S	1S	2S	1S	2S
0.01	85	15	66	14	68	9	66	8
0.1	93	33	127	60	116	27	113	21
0.01	96	27	249	65	1551	83	2395	59
0.001	97	26	238	61	2633	88	8905	86
	c_1							
	1.0		0.99		0.9		0.5	
c_2	1S	2S	1S	2S	1S	2S	1S	2S
1.0	6559	203	2532	115	432	33	134	14
0.99	2539	116	1551	83	398	31	132	14
0.9	431	33	396	31	252	22	119	13
0.5	132	14	131	14	118	13	91	8
	ϵ_{Σ_1}							
	10.0		1.0		0.1		0.01	
ϵ_{Σ_2}	1S	2S	1S	2S	1S	2S	1S	2S
10.0	136	66	225	65	246	62	247	61
1.0	217	64	1551	83	2395	42	2477	35
0.1	235	60	2349	36	1569	8	1261	6
0.01	235	60	2339	33	1267	6	855	4

In each of the plots in Fig. 1, $c = 1.0$, but the materials are more atomically mixed in Fig. 1(b) and (d). The spectrum of the 2S operator in Fig. 1(a) is confined to the real line and is tightly clustered at unity with a few outliers. The 1S spectrum is less tightly clustered with some imaginary eigenvalues, and the clusters are farther from unity. In Fig. 1(b), the 1S spectrum has shifted closer to zero and, although the cluster is slightly smaller than in Fig. 1(a), its proximity to zero indicates that convergence will be slow. Compared with Fig. 1(a), the spectrum for the 2S operator has also shifted away from unity, although not as dramatically as the 1S operator, and there are more outliers between zero and unity. The spectra indicate that both operators are well-suited for GMRES and that convergence will be slower in the atomic mix regime for both operators, but the 2S method should converge more rapidly in general.

The preconditioned operators for the inner iterations (Fig. 1(c) and (d)) show better convergence properties than the unpreconditioned as their spectra are more tightly clustered and the centers of their clusters are closer to unity. The spectrum of the unpreconditioned MI operator in Fig. 1(c) also shows much better convergence properties than that in Fig. 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. As discussed previously and as can be seen in comparing Fig. 1(c) and (d), varying λ affects the convergence rates of the SI. While c is the same in both cases, the spectra in Fig. 1(d) are tightly clustered close to unity while the spectra in Fig. 1(c) are far from unity and are much more spread out, indicating that the SI should converge more quickly in the atomic mix case.

Fig. 2 shows spectra for the same parameters as Fig. 1 with the exception of the scattering ratio, which is now $c = 0.5$. Comparing Figs. 1(a) and 2(a) reveals that while the shapes of the 1S and 2S spectra are similar for different c, the outliers on the real line have shifted towards unity in the $c = 0.5$ case. Since the material is less diffusive, it is expected that GMRES will converge more rapidly. In Figs. 1(b) and 2(b), the same effect is observed, although it is not as drastic for the 1S spectrum. Since only c has changed, the spectra for the material iteration are unchanged and since c is now smaller, the spectra for the species iteration are smaller and have shifted towards unity.

It is also known that restarted GMRES is guaranteed to converge if the matrix is positive definite (PD) [16,17]. The operator for the 1S scheme is generally not PD, while the 2S operator usually is. The operator for the MI is not PD in those cases where $\epsilon_\lambda \sim \mathcal{O}(0.001)$, but the *preconditioned* operator is PD in all cases. The SI operators are PD in all cases, both with and without preconditioning. Although restarted GMRES is guaranteed to converge whenever the operator is PD, it is also not

Table 4
Simple iteration: MFLOP counts for the one-step iteration and two-step iteration with accelerated inner iterations (10 cm slab, isotropic incidence on the left boundary and reflecting boundary condition on the right, $I = 100, N = 8, \tilde{\Sigma}_1 = 5.0\ \text{cm}^{-1}, \tilde{\Sigma}_2 = 1.0\ \text{cm}^{-1}, \tilde{\lambda}_1 = 1.0$ cm and $\tilde{\lambda}_2 = 5.0$ cm). Unless otherwise specified, $c = 0.99$ and $\epsilon_\lambda = 0.01$.

	ϵ_λ							
	1.0		0.1		0.01		0.001	
c	1S	2S	1S	2S	1S	2S	1S	2S
1.0	360.92	597.96	594.45	3976.82	2457.48	6045.07	19778.2	4990.6
0.99	92.96	169.9	137.1	877.98	532.25	1311.39	4245.41	1130.7
0.9	11.7	31.91	16.96	115.67	66.96	167.42	539.99	220.29
0.5	1.98	10.26	3.17	21.78	15.37	35.03	133.23	104.93
	ϵ_{λ_1}							
	1.0		0.1		0.01		0.001	
ϵ_{λ_2}	1S	2S	1S	2S	1S	2S	1S	2S
1.0	92.96	169.9	83.13	187.23	81.05	123.03	80.75	97.57
0.1	113.59	262.49	137.1	877.98	144.54	629.38	144.25	445.49
0.01	118.06	184.67	200.6	970.22	532.25	1311.39	712.21	846.47
0.001	119.15	147.17	211.31	683.08	999.82	1329.24	4245.41	1130.7
	c_1							
	1.0		0.99		0.9		0.5	
c_2	1S	2S	1S	2S	1S	2S	1S	2S
1.0	2457.48	6045.07	874.71	2153.44	129.86	322.18	28.57	69.73
0.99	874.81	2153.44	532.25	1311.39	118.55	293.35	27.97	68.26
0.9	129.86	322.18	118.55	293.35	66.96	167.42	24.01	57.99
0.5	28.57	69.73	27.97	68.26	23.91	57.99	15.37	35.03
	ϵ_{Σ_1}							
	10.0		1.0		0.1		0.01	
ϵ_{Σ_2}	1S	2S	1S	2S	1S	2S	1S	2S
10.0	157.84	1021.0	204.76	910.22	213.69	802.33	214.68	767.69
1.0	201.39	1041.38	532.25	1311.39	625.8	826.31	627.09	736.58
0.1	209.52	926.75	614.0	883.92	345.44	150.65	261.91	93.57
0.01	210.42	891.61	613.1	789.67	261.51	93.48	151.98	51.83

guaranteed to fail when the operator is not PD. Indeed, for all cases examined, GMRES(10) always converges (see Section 5). However, positive-definiteness is a desirable quality that occurs more frequently in the 2S operator than the 1S and the preconditioned MI than the unpreconditioned.

5. Numerical results

Numerical results were obtained for the 1S and 2S schemes using both the simple iterative method and GMRES(10). For the 2S scheme, atomic mix and S_2 synthetic accelerations are always applied to the inner material and species iterations, respectively. Note, once again, that accelerating the inner iterations does not affect the convergence of the outer 2S iteration since the final results of each inner iterations should be roughly the same regardless of whether or not acceleration is used. In all cases $\bar{\Psi}_0$ is used as the initial guess for $\bar{\Psi}$ and the solution is converged when the relative residual is less than or equal to 10^{-6}. Four different studies are conducted in Tables 2–5. In the first, the scattering ratio, c, and order of magnitude of $\lambda, \epsilon_\lambda$, in the two materials are varied simultaneously. Thus the atomic mix parameter, $\Sigma\lambda$, is the same in both materials as well. In the second study, ϵ_{λ_1} and ϵ_{λ_2} vary separately for $c = 0.99$. In the third, c_1 and c_2 vary independently for $\epsilon_\lambda = 0.01$. And finally, in the fourth study ϵ_{Σ_1} and ϵ_{Σ_2} are allowed to vary for $\epsilon_\lambda = 0.01$ and $c = 0.99$.

Tables 2 and 3 show outer iteration counts for the simple iterations and GMRES(10), respectively. A restart of 10 was selected for GMRES because it achieved rapid convergence for the cases examined while requiring relatively little computer storage. The 2S method always converges in fewer iterations than the 1S method for both the simple iteration and GMRES(10) and, in most cases, the iteration counts differ by one to two orders of magnitude. Similarly, in comparing the simple and GMRES(10) iteration counts for the 1S and 2S iterations, GMRES(10) always converges in fewer iterations than the simple scheme, and the counts differ by at least one order of magnitude in most cases. Although the data for the inner iterations are not shown here, GMRES(10) always converges in fewer iterations than traditional MI or SI. Also, accelerating/preconditioning decreases the iteration count in most cases, therefore all data for the 2S scheme are shown for accelerated/preconditioned inner iterations.

Each GMRES(10) iteration requires more work than a single simple iteration, and a single 2S iteration requires more work than a 1S iteration, particularly in diffusive and/or atomically mixed regimes in which the species and/or material iterations

Table 5
GMRES(10): MFLOP counts for the one-step iteration and two-step iteration with accelerated inner iterations (10 cm slab, isotropic incidence on the left boundary and reflecting boundary condition on the right, $I = 100, N = 8, \widetilde{\Sigma}_1 = 5.0\ \text{cm}^{-1}, \widetilde{\Sigma}_2 = 1.0\ \text{cm}^{-1}, \tilde{\lambda}_1 = 1.0$ cm and $\tilde{\lambda}_2 = 5.0$ cm). Unless otherwise specified, $c = 0.99$ and $\epsilon_\lambda = 0.01$.

	ϵ_λ							
	1.0		0.1		0.01		0.001	
c	1S	2S	1S	2S	1S	2S	1S	2S
1.0	26.68	35.91	67.35	218.69	1169.99	293.16	7266.77	304.54
0.99	15.11	31.89	22.56	86.52	276.78	126.84	1598.6	180.45
0.9	4.4	18.7	6.68	28.14	44.99	41.93	219.63	110.51
0.5	1.59	8.55	2.76	11.99	16.18	21.98	61.46	92.31
	ϵ_{λ_1}							
	1.0		0.1		0.01		0.001	
ϵ_{λ_2}	1S	2S	1S	2S	1S	2S	1S	2S
1.0	15.11	31.89	11.88	25.4	12.11	16.44	11.88	14.7
0.1	16.5	50.0	22.56	86.52	20.6	36.06	20.1	27.52
0.01	17.07	39.59	44.34	100.35	276.78	126.84	427.19	78.89
0.001	17.29	32.3	42.38	73.83	469.61	136.37	1598.6	180.45
	c_1							
	1.0		0.99		0.9		0.5	
c_2	1S	2S	1S	2S	1S	2S	1S	2S
1.0	1169.99	293.16	451.66	171.32	77.01	57.24	23.88	30.68
0.99	452.91	172.72	276.78	126.84	70.98	54.32	23.65	30.68
0.9	76.87	56.97	70.74	54.16	44.99	41.93	21.21	29.29
0.5	23.65	30.46	23.41	30.46	21.0	29.07	16.18	21.98
	ϵ_{Σ_1}							
	10.0		1.0		0.1		0.01	
ϵ_{Σ_2}	1S	2S	1S	2S	1S	2S	1S	2S
10.0	24.18	90.96	40.04	88.1	43.8	90.39	43.97	88.08
1.0	38.66	108.5	276.78	126.84	427.19	74.46	441.85	66.47
0.1	41.84	100.08	418.94	72.45	279.81	46.84	224.87	44.97
0.01	41.84	96.21	417.17	69.26	226.0	44.82	152.47	42.66

are extremely time-consuming. Therefore, while iteration count indicates the effectiveness of a method, it says little about its computational efficiency. The computational work and relative efficiency are therefore quantified using floating point operation (FLOP) counts given in mega-FLOPS (MFLOPS). The results are displayed in Tables 4 and 5 for the simple iteration and GMRES(10), respectively.

GMRES(10) always requires less computational effort than the simple iteration despite the larger cost per iteration and, although data are only shown for accelerated/preconditioned inner iterations, in all cases examined it is always more efficient to accelerate or precondition the inner iterations. For the simple iteration scheme, the 1S is actually more efficient than the 2S in most cases, the exceptions being those materials in which $\epsilon_\lambda \sim \mathcal{O}(0.001)$. However, for GMRES(10), the 2S scheme is often more efficient than the 1S. This can be attributed to the fact that the 2S operator is much better suited to GMRES than the 1S operator, as discussed in Section 4. Furthermore, the 2S scheme with GMRES(10) is much more efficient in the atomic mix-diffusion regime, which is the regime in which the 1S scheme is particularly inefficient. Conversely, it is only slightly less efficient when the material is not atomically mixed or diffusive where the 1S scheme is quite efficient anyway.

6. Conclusion

A new two-step iterative scheme for solving the Levermore–Pomraning transport equations for binary statistical materials has been devised and demonstrated. In this two-step scheme, each iteration is divided into a series of inner material and source iterations that attenuate the diffusion and atomic mix error modes separately. Atomic mix and S_2 synthetic accelerations are applied to the inner material and source iterations, respectively, to offset the cost of doing several inner iterations per outer iteration. A Krylov iterative solver – specifically, restarted GMRES– is then wrapped around each iteration, inner and outer. The method is demonstrated for a wide variety of combinations of physical parameters.

As predicted by spectral analysis, the new two-step scheme always converges in fewer iterations than the one-step scheme for both simple iterations and GMRES(10), but it is not always more computationally efficient since each two-step iteration is more expensive than a one-step iteration. However, the two-step is far more efficient in the atomic mix-diffusion regime in which the one-step iteration is extremely time-consuming, while it is only slightly less efficient in regimes where the one-step scheme converges rapidly anyway. Accelerating/preconditioning the inner iterations with S_2SA and AMSA improves their convergence rate and results in better overall computational efficiency in most cases. Thus, the new restarted GMRES two-step iteration scheme with preconditioned inner iterations has been shown to be an effective and efficient solution technique in materials that are atomically mixed and diffusive.

Appendix. The matrices given in Section 3.4 are written explicitly in this Appendix. In all cases, angularly dependent quantities are written in the following vector form:

$$\vec{\psi} = \begin{bmatrix} \vec{\psi}_1 \\ \vec{\psi}_2 \end{bmatrix} \quad \vec{\psi}_\ell = \begin{bmatrix} \vec{\psi}_{\ell,1} \\ \vec{\psi}_{\ell,2} \\ \vdots \\ \vec{\psi}_{\ell,N} \end{bmatrix} \quad \vec{\psi}_{\ell,n} = \begin{bmatrix} \vec{\psi}_{\ell,n,1} \\ \vec{\psi}_{\ell,n,2} \\ \vdots \\ \vec{\psi}_{\ell,n,I} \end{bmatrix} \quad \vec{\psi}_{\ell,n,i} = \begin{bmatrix} \psi^+_{\ell,n,i} \\ \psi^-_{\ell,n,i} \end{bmatrix}$$

where ℓ is the material identity, n is index of the angular abscissa, i is the spatial cell (where 1 is the index of the left-most cell and I is the index of the right-most cell) and $\psi^+_{\ell,n,i}$ and $\psi^-_{\ell,n,i}$ are the right and left cell edge quantities, respectively, in cell i at angle n in material ℓ. The scalar flux is given by $\vec{\phi} = \mathbf{D}\vec{\psi}$ where

$$\mathbf{D} = [w_1\, w_2, \ldots, w_N] \otimes \mathbf{I}_{2I},$$

w_n is the nth quadrature weight, $\otimes$ denotes the Kronecker product and $\mathbf{I}_{2I}$ denotes the $2I \times 2I$ identity matrix.

In order to define **A**, **B** and **C**, we first define the matrices **L**, which contains the streaming $(\mu \frac{\partial \psi}{\partial x})$ and removal $(\Sigma\psi)$ terms, **P**, which contains the coupling terms, and **S**, which contains the scattering terms. These matrices are written in terms of 2 sub-matrices corresponding to the two materials and are composed of 2×2 sub-matrices arranged in a larger square matrix:

$$\mathbf{L} = \begin{bmatrix} \mathbf{L_1} & 0 \\ 0 & \mathbf{L_2} \end{bmatrix}, \quad \mathbf{S} = \begin{bmatrix} \mathbf{S_1} & 0 \\ 0 & \mathbf{S_2} \end{bmatrix} \quad \text{and} \quad \mathbf{P} = \begin{bmatrix} 0 & \mathbf{P_2} \\ \mathbf{P_1} & 0 \end{bmatrix}$$

$\mathbf{L_1}$ and $\mathbf{L_2}$ are composed of IN blocks in each row and column which contain 2×2 sub-matrices. The non-zero 2×2 sub-matrices are given by:

$$\mu_n > 0 : \mathbf{L}_{\ell,n,i} = \begin{bmatrix} \frac{4\mu_n}{\Delta x} + 2\left(\Sigma_\ell + \frac{|\mu_n|}{\lambda_\ell}\right) & 2\left(\Sigma_\ell + \frac{|\mu_n|}{\lambda_\ell}\right) \\ -2\left(\Sigma_\ell + \frac{|\mu_n|}{\lambda_\ell}\right) + & \frac{12\mu_n}{\Delta x} + 2\left(\Sigma_\ell + \frac{|\mu_n|}{\lambda_\ell}\right) \end{bmatrix} \quad \text{and} \quad \mathbf{L}_{\ell,n,i-1} = \begin{bmatrix} -\frac{4\mu_n}{\Delta x} & 0 \\ -\frac{12\mu_n}{\Delta x} & 0 \end{bmatrix}$$

$$\mu_n < 0 : \mathbf{L}_{\ell,n,i} = \begin{bmatrix} 2\left(\Sigma_\ell + \frac{|\mu_n|}{\lambda_\ell}\right) & -\frac{4\mu_n}{\Delta x} + 2\left(\Sigma_\ell + \frac{|\mu_n|}{\lambda_\ell}\right) \\ \frac{12\mu_n}{\Delta x} - 2\left(\Sigma_\ell + \frac{|\mu_n|}{\lambda_\ell}\right) & 2\left(\Sigma_\ell + \frac{|\mu_n|}{\lambda_\ell}\right) \end{bmatrix} \quad \text{and} \quad \mathbf{L}_{\ell,n,i+1} = \begin{bmatrix} 0 & \frac{4\mu_n}{\Delta x} \\ 0 & -\frac{12\mu_n}{\Delta x} \end{bmatrix}$$

where Δx is the cell width. $\mathbf{B_1}$ and $\mathbf{B_2}$ are composed of 2×2 sub-matrices arranged in a larger $I \times I$ matrix:

$$\mathbf{B} = \begin{bmatrix} \mathbf{B_1} & 0 \\ 0 & \mathbf{B_2} \end{bmatrix} \quad \text{where} \quad \mathbf{B}_\ell = \mathbf{I}_I \otimes \begin{bmatrix} \Sigma_{s,\ell} & \Sigma_{s,\ell} \\ -\Sigma_{s,\ell} & \Sigma_{s,\ell} \end{bmatrix}.$$

$\mathbf{P_1}$ and $\mathbf{P_2}$ have 2×2 sub-matrices on their diagonals and take the form:

$$\mathbf{P}_\ell = \begin{bmatrix} \mathbf{P}_{\ell,1} & 0 & \dots & 0 \\ 0 & \mathbf{P}_{\ell,2} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \mathbf{P}_{\ell,N} \end{bmatrix} \quad \text{and} \quad \mathbf{P}_{\ell,n} = \mathbf{I}_I \otimes \begin{bmatrix} 2\frac{|\mu_n|}{\lambda_\ell} & 2\frac{|\mu_n|}{\lambda_\ell} \\ -2\frac{|\mu_n|}{\lambda_\ell} & 2\frac{|\mu_n|}{\lambda_\ell} \end{bmatrix}.$$

It is also helpful to define the upper and lower triangular parts of $\mathbf{P}$:

$$\mathbf{P_U} = \begin{bmatrix} 0 & \mathbf{P_2} \\ 0 & 0 \end{bmatrix} \quad \text{and} \quad \mathbf{P_L} = \begin{bmatrix} 0 & 0 \\ \mathbf{P_1} & 0 \end{bmatrix}.$$

We can now write $\mathbf{A}$, $\mathbf{B}$ and $\mathbf{C}$ as:

$$\mathbf{A} = \mathbf{A_m} = \mathbf{I}_{4IN} - \mathbf{L}^{-1}\mathbf{P_L}$$
$$\mathbf{B} = \mathbf{B_m} = \mathbf{L}^{-1}\mathbf{S}$$
$$\mathbf{C} = \mathbf{C_m} = \mathbf{L}^{-1}\mathbf{P_U}$$
$$\mathbf{A_s} = \mathbf{L}$$
$$\mathbf{B_s} = \mathbf{S}$$
$$\mathbf{C_s} = \mathbf{P}.$$

Note that $\mathbf{L}^{-1}$ represents the inversion of the operator left-hand-side of the transport equation and is often referred to as a 'transport sweep.'

References

[1] C.D. Levermore, G.C. Pomraning, D.L. Sanzo, J. Wong, Linear transport theory in a random medium, Journal of Mathematical Physics 27 (1986) 2526–2536.
[2] D. Vanderhaegen, Radiative transfer in statistically heterogeneous mixtures, Journal of Quantitative Spectroscopy and Radiative Transfer 36 (1986) 557–561.
[3] E.W. Larsen, R. Vasques, M.T. Vilhena, Particle transport in the 1-d diffusive atomic mix limit, in: International Conference on Mathematics, Computational Methods and Reactor Physics, Avignon, France, 2005.
[4] I.M. Davis, T.S. Palmer, E.W. Larsen, A comparison of binary stochastic media transport models in "solid–void mixtures, in: PHYSOR 2004-The Physics of Fuel Cycles and Advanced Nuclear Systems: Global Developments, Chicago, IL, 2004.
[5] T.S. Palmer, B.S. Ching, A "two-grid" acceleration of binary stochastic mixture deterministic transport iterations in slab geometry, Annals of Nuclear Energy 35 (2008) 68–74.
[6] M.L. Adams, J.E. Morel, A "two-grid" acceleration scheme for the multigroup s_n equations with neutron upscattering, Nuclear Science and Engineering 115 (1993) 253–264.
[7] T.S. Palmer, A coupled diffusion synthetic acceleration for binary stochastic mixture transport iterations in slab geometry, Nuclear Science and Engineering 158 (2008) 40–48.
[8] A.K. Prinja, E.D. Fichtl, Atomic mix synthetic acceleration of dose computations in binary statistical media, Nuclear Science and Engineering 155 (2007) 441–448.
[9] J.S. Warsa, T.A. Wareing, 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 (2004) 218–248.
[10] G.C. Pomraning, Linear Kinetic Theory and Particle Transport in Stochastic Mixtures, World Scientific, Singapore, 1991.
[11] E.E. Lewis, W.F. Miller, Computational Methods of Neutron Transport, American Nuclear Society, La Grange Park, Illinois, 1993.
[12] L.J. Lorence, J.E. Morel, 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 (1989) 341–351.
[13] M.L. Adams, E.W. Larsen, Fast iterative methods for discrete-ordinates particle transport calculations, Progress in Nuclear Energy 40 (2002) 3–159.
[14] B. Guthrie, J.P. Holloway, B.W. Patton, GMRES as a multi-step transport sweep accelerator, Transport Theory and Statistical Physics 28 (1999) 83–102.
[15] S.L. Campbell, I.C.F. Ipsen, C.T. Kelley, C.D. Meyer, GMRES and the minimal polynomial, BIT Numerical Mathematics 36 (1996) 664–675.
[16] Y. Saad, Iterative Methods for Sparse Linear Systems, PWS Publishing Company, Boston, 1996.
[17] A. Greenbaum, Iterative Methods for Solving Linear Systems, SIAM, Philadelphia, 1997.