



PERGAMON

Journal of Quantitative Spectroscopy &
Radiative Transfer 69 (2001) 671–707

Journal of
Quantitative
Spectroscopy &
Radiative
Transfer

www.elsevier.com/locate/jqsrt

Fast and accurate discrete ordinates methods for multidimensional radiative transfer. Part I, basic methods

Dinshaw Balsara*

*NCSA, University of Illinois at Urbana-Champaign, 5600 Beckman Institute of Advanced Science and Technology,
Drawer 25, 405 North Mathews Avenue, Urbana, IL 61801, USA*

Received 17 April 2000; accepted 12 June 2000

Abstract

The ability to solve the radiative transfer equation in a fast and accurate fashion is central to several important applications in combustion physics, controlled thermonuclear fusion and astrophysics. Most practitioners see the value of using discrete ordinates methods for such applications. However, previous efforts at designing discrete ordinates methods that are both fast and accurate have met with limited success. This is especially so when parts of the application satisfy the free streaming limit in which case most solution strategies become unacceptably diffusive or when parts of the application have high absorption or scattering opacities in which case most solution strategies converge poorly. Designing a single solution strategy that retains second-order accuracy and converges with optimal efficiency in the free streaming limit as well as the optically thick limit is a challenge. Recent results also indicate that schemes that are less than second-order accurate will not retrieve the radiation diffusion limit.

In this paper we analyze several of the challenges involved in doing multidimensional numerical radiative transfer. It is realized that genuinely multidimensional discretizations of the radiative transfer equation that are second-order accurate exist. Because such discretizations are more faithful to the physics of the problem they help minimize the diffusion in the free streaming limit. Because they have a more compact stencil, they have superior convergence properties. The ability of the absorption and scattering terms to couple strongly to the advection terms is examined. Based on that we find that operator splitting of the scattering and advection terms damages the convergence in several situations. Newton–Krylov methods are shown to provide a natural way to incorporate the effects of nonlinearity as well as strong coupling in a way that avoids operator splitting. Used by themselves, Newton–Krylov methods converge slowly. However, when the Newton–Krylov methods are used as smoothers within a full approximation scheme multigrid method, the convergence is vastly improved. *The combination of a genuinely multidimensional, nonlinearly positive scheme that uses Full Approximation Scheme multigrid in conjunction with the Newton–Krylov method is shown to*

* Tel.: + 1-217-244-0072; fax: + 1-217-244-2909.

E-mail address: dbalsara@ncsa.uiuc.edu (D. Balsara).

result in a discrete ordinates method for radiative transfer that is highly accurate and converges very rapidly in all circumstances.

Several convergence studies are carried out which show that the resultant method has excellent convergence properties. Moreover, this excellent convergence is retained in the free streaming limit as well as in the limit of high optical depth. *The presence of strong scattering terms does not slow down the convergence rate for our method. In fact it is shown that without operator splitting, the presence of a strong scattering opacity enhances the convergence rate in quite the same way that the convergence is enhanced when a high absorption opacity is present!* We show that the use of differentiable limiters results in substantial improvement in the convergence rate of the method. By carrying out an accuracy analysis on meshes with increasing resolution it is further shown that the accuracy that one obtains seems rather close to the designed second-order accuracy and does not depend on the specific choice of limiter.

The methods for multidimensional radiative transfer that are presented here should improve the accuracy of several radiative transfer calculations while at the same time improving their convergence properties. Because the methods presented here are similar to those used for simulating neutron transport problems and problems involving rarefied gases, those fields should also see improvements in their numerical capabilities by assimilation of the methods presented here. © 2001 Elsevier Science Ltd. All rights reserved.

Keywords: Multidimensional radiative transfer; Discrete ordinates methods; Multidimensional upwinding; Newton–Krylov methods; Multigrid; Rapid convergence; Second-order accuracy

1. Introduction

Radiation plays a very important role in the transfer of energy in various scientific applications. This is especially true for applications involving high-temperature phenomena where the radiation can on occasion dominate the total energy in the system being studied. Such applications are typically drawn from combustion physics, controlled thermonuclear fusion and astrophysics. In such applications one typically wants to solve for the radiation field in multiple dimensions on a computational mesh. In principle, this should be an easy task since the photons that constitute the radiation field move along rays till they interact with matter. When they do interact with matter they either undergo absorption or scattering. The matter can also act as a source of photons and emit new photons. The rays along which photons move are always characteristic lines, which are always straight lines in all situations except when strong (and anisotropic) general relativistic effects become important.¹ We set aside the anisotropic general relativistic case for now because there is an abundance of applications that one can explore without the inclusion of general relativity. In practice, radiative transfer calculations have been rather difficult for several reasons. As a result, several approximate solution techniques have been tried. These include radiation hydrodynamics, P_n and simplified P_n methods and the discrete ordinates (i.e. S_n) method. We review each method briefly below.

¹ Note that this restriction is not so severe because we can still use the methods developed here to treat cosmological radiative transfer.

The fact that radiation can flow from one point to another suggests that a hydrodynamical model may be built for it. As a result, Castor [1] and Buchler [2,3] applied a closure approximation to the second moments of the transport equation to derive the equations of radiation hydrodynamics. As a result, the equations of radiation hydrodynamics describe the interaction of the radiation field, treated as a fluid, with matter. Balsara [4–6] showed that very powerful higher-order Godunov techniques can be applied to those equations. Furthermore, Balsara [7,8] showed that these advances extend to the equations of radiation magnetohydrodynamics. However, a fundamental shortcoming in the solution of the equations of radiation hydrodynamics stems from the need to provide accurate Eddington factors. In all problems other than optically thick ones the Eddington factors can become strongly anisotropic and can vary rapidly in space. This indicates that to a certain extent the use of a closure approximation at the level of second moments was an inadequate approximation.

The P_n approximation was first introduced by Krook [9] and Cheng [10]. It consists of taking moments of the radiative transport equation. This yields a system of equations with the angular dependence removed from them. Each P_n approximation yields a system of n^2 equations. In the asymptotic limit of large “ n ” the P_n approximation yields the solution of the actual transport equation. However, the positivity of the radiation field cannot be guaranteed by the P_n method. This led Gelbard [11], Liu and Gelbard [12], Morel [13] and Larsen et al. [14] to formulate the simplified P_n (i.e. SP_n) method. The positivity of the radiation field can be guaranteed by the SP_n methods. However, it cannot be guaranteed that in the limit of large “ n ” the SP_n method yields the solution of the actual transport equation. In practice only the lower moments are used and a wealth of practical experience only exists for the SP_1 and SP_3 methods. Thus the SP_1 approximation yields the familiar radiation diffusion approximation, which removes all angular dependence from the radiation field. The P_3 approximation retains some angular dependence, see Ratzell and Howell [15], as does the SP_3 approximation. However, the SP_3 approximation does not permit the same level of angular dependence to be retained in the radiation field as the higher order S_n approximations, see Larsen et al. [14]. Furthermore, each higher member in the family of P_n approximations needs to be individually formulated, making it difficult to obtain the higher-order members.

The discrete ordinates (S_n) method for radiative transfer was first formulated by Chandrasekhar [16] and later on developed by Lathrop and Carlson [17]. Each S_n approximation consists of discretizing the angular dependence in the transport equation along $n(n + 2)$ discrete ordinates. Numerous methods and formulations for solving the S_n transport equations have been devised. Thus Larsen and Morel [18] and Morel et al. [19] have devised a lumped linear-discontinuous (LLD) scheme. Their method is very much like the Galerkin formulations that are used in computational fluid dynamics (CFD) and even uses some concepts of upwinding to achieve second-order accuracy. Adams [20] and Adams and Novak [21] have designed and analyzed a simple corner balance (SCB) method that is very much like the lumped linear-discontinuous scheme in certain limits. Even parity, finite-element method (EP-FEM)-based solution techniques for the radiative transfer equation were explored in Fiveland and Jessee [22,23]. The EP methods have the advantage that they use a self-adjoint operator for the even parity transport equations and can, therefore, ensure positivity of the radiation field. However, for problems involving anisotropic scattering this advantage is lost. Also, in the limit of optically thick zones, the method loses

accuracy and converges poorly. First-order accurate control volume-based discrete ordinates methods were originally formulated by Fiveland [24] for unpolarized radiation and by Haferman et al. [25] for polarized radiation. Total variation diminishing (TVD) techniques drawn from the work of vanLeer [26] and Harten [27] were first applied to radiative transfer problems by Jessee and Fiveland [28]. The introduction of TVD schemes rectified several shortcomings in the earlier diamond discretization scheme of Fiveland [29]. Recently, Jessee et al. [30] and Howell et al. [31] have used first-order accurate formulations of these techniques in self-adaptive radiative transfer problems. Howell et al. [31] self-consistently calculated the radiation and chemistry of a combustion flow problem showing the great utility of these methods. For that reason we focus on control volume-based discrete ordinates methods in this work. Our more immediate goal is to arrive at a higher-order, accurate and rapidly convergent strategy for numerical radiative transfer that we wish to use in the author's RIEMANN framework for parallel adaptive mesh refinement (AMR), see Balsara and Norton [32].

The discrete ordinate methods mentioned above have several attractive aspects. However, there are several situations where they do not converge rapidly, necessitating the use of convergence acceleration methods. Hence, Fiveland and Jessee [33] have proposed mesh rebalance and synthetic acceleration methods to accelerate the solution to convergence. Likewise, Ramone et al. [34] proposed a transport synthetic acceleration (TSA) method for the scattering terms which they show to be superior to the diffusion synthetic acceleration method of Larsen [35]. However, several problems persist. They are:

(1) Mesh rebalance fails to produce rapid convergence for problems with high optical thickness, see Fiveland and Jessee [33]. TSA has a similar shortcoming when the problem is optically thick and strongly scattering. Both mesh rebalance and TSA have difficulties in making the solution converge when the medium is strongly scattering.

(2) The use of higher-order, nonlinear operators to discretize the advection term in the transfer equation causes a degradation of the convergence properties of the scheme, see Jessee and Fiveland [28].

(3) Jessee and Fiveland [28] proposed the use of Gauss–Seidel relaxation which, because it uses directional sweeps, causes a partial serialization of the parallel implementation of these methods, see Burns and Christon [36]. This is so even when first-order methods are used. This partial serialization can prevent the radiative transfer problem from being highly parallelizable. It can also cause bottlenecks in the parallel processing of self-adaptive radiative transfer problems. For example, the method of Howell et al. [31] requires the grid patches in the AMR hierarchy to be processed in different sequences for different ordinates, making efficient load balancing very difficult if not impossible.

(4) It is also known that the convergence of Gauss–Seidel relaxation degrades when the mesh has zones with bad aspect ratios. The problem can be cured with semi-coarsening but only at the expense of introducing extra programming complexity in supporting the semi-coarsening techniques. It is also worthwhile to point out that in several radiative transfer applications, meshes that have zones with bad aspect ratio are the norm rather than the exception.

(5) Most methods for multidimensional radiative transfer used to date make use of a dimension-by-dimension discretization of the advection part of the transfer equation. This is a serious

shortcoming because the characteristics almost never point along the grid lines of the computational mesh.

(6) In anisotropic scattering problems or problems with moving media, the frequency groups can interact with each other. The above-mentioned acceleration methods do not provide a natural way to account for that mixing across frequency groups. As a result, when multiple frequency groups are being used, a linear multifrequency-gray acceleration method has to be invoked, see Morel et al. [37].

In this paper we examine two-dimensional radiative transfer. Insights gained in two dimensions will be applied to three dimensions in a subsequent paper. In this paper we put forward the viewpoint that superior, genuinely multidimensional, discretizations of the transport equation exist. Furthermore, we point out that the relaxation strategies that have been used in the past to iterate second-order accurate discretizations of the transfer equation to convergence are not well matched to the discretizations that have been used. This is because the relaxation strategies benefit maximally from diagonal dominance and compact structure of the discretized transfer equation even though the dimension-by-dimension upwind discretizations do not have these properties. The genuinely multidimensional, nonlinearly positive discretizations of Sidilkover and Roe [38] are shown to rectify this problem because they have a compact stencil. Previously used relaxation methods for iterating the transport equation to convergence have always used an operator splitting of the discretized transfer equation into an advection term and a source term which contains the scattering contribution. The acceleration schemes were then invoked within each iteration to correct for this operator splitting. However, principal component analysis shows that this operator splitting is not justified in several situations. For that reason, the Krylov-subspace technique of Saad and Schultz [39] is used to relax the discretized transfer equation without resort to operator splitting. This permits the scattering term to fully participate in the relaxation process. All TVD discretizations of the radiative transfer equation are nonlinear. For that reason, the Newton–Krylov technique of Brown and Saad [40] is used to factor in a Newton step into the Krylov-subspace iteration. When applied to the radiative transfer equation the Krylov-subspace techniques display very slow convergence for a lot of iterations. To accelerate convergence, the full approximation scheme (FAS) multigrid strategy of Brandt [41] is used with the Newton–Krylov method acting as a nonlinear multigrid smoother on each of the levels. *The combination of a genuinely multidimensional, nonlinearly positive scheme that uses FAS multigrid in conjunction with the Newton–Krylov method is shown to result in an S_n scheme for radiative transfer that is highly accurate and converges very rapidly in all circumstances.* It also permits the same solution strategy to be invoked on all levels in the multigrid hierarchy, further simplifying the solution strategy. Balsara and Norton [32] have shown that self-adaptive multigrid techniques can be parallelized extremely well. The Krylov-subspace techniques do not require ordered processing of the grid patches in an AMR hierarchy. As a result, the multidimensional radiative transfer strategy presented here is eminently well suited for use in parallel AMR applications.

In Section 2 we catalogue and discuss various discretizations of the radiative transfer equation. In Section 3 we briefly review Newton–Krylov techniques and FAS multigrid. In Section 4 we present a sequence of model problems which we use to demonstrate why the combination of techniques invoked here yield extremely powerful strategies for multidimensional radiative transfer. In Section 5 we present several realistic test problems. In Section 6 we present some conclusions.

2. The radiative transfer equation and its discretization

The radiative transfer equation (RTE) for a moving medium was written out in its most general form by Mihalas and Klein [42]. For the purposes of this paper it is adequate to focus on the following form of the RTE:

$$(\hat{\Omega} \cdot \nabla)I(\mathbf{r}, \hat{\Omega}) = -(\kappa + \sigma)I(\mathbf{r}, \hat{\Omega}) + \frac{\sigma}{4\pi} \int_{4\pi} I(\mathbf{r}, \hat{\Omega}') d\hat{\Omega}' + \kappa I_b(\mathbf{r}), \quad (1)$$

where $\hat{\Omega} = \Omega_x \hat{x} + \Omega_y \hat{y} + \Omega_z \hat{z}$ is the ordinate direction; $I(\mathbf{r}, \hat{\Omega})$ is the radiation intensity at the location \mathbf{r} in the ordinate direction $\hat{\Omega}$; $I_b(\mathbf{r})$ is the intensity of blackbody radiation at the temperature of the medium and κ and σ are the absorption and scattering opacities of the medium. For computational domains that are bounded by gray surfaces that reflect diffusively, the above equation has to be solved with the boundary condition:

$$I(\mathbf{r}, \hat{\Omega}) = \varepsilon I_b(\mathbf{r}) + \frac{\rho}{\pi} \int_{\hat{n} \cdot \hat{\Omega}' < 0} |\hat{n} \cdot \hat{\Omega}'| I(\mathbf{r}, \hat{\Omega}') d\hat{\Omega}', \quad (2)$$

where \hat{n} is the inwardly pointing normal at the surface; ε is the surface emissivity; $\rho = 1 - \varepsilon$ is the surface reflectivity and $I_b(\mathbf{r})$ is the intensity being emitted by the boundary. Eq. (2) for the boundaries is general enough to accommodate the free space boundary conditions that are common in astrophysics because one would then use $\varepsilon = 1$ and set the intensity at the boundary to be equal to the incident intensity that is shining into the computational domain from the outside.

Eq. (1) can be discretized for a given ordinate set S_n . For more on the choice of ordinate sets see the appendix in this paper and Fiveland [43] and Lathrop and Carlson [17]. Thus the RTE is replaced by a discrete set of equations for a finite number of ordinates $\hat{\Omega}_m$. The integrals in Eqs. (1) and (2) are replaced by a quadrature so that each ordinate has a weight w_m . There are $M = n(n + 2)$ such ordinates in three dimensions and half as many for two dimensions. Thus the discrete ordinates approximation of the RTE in Eq. (1) becomes.

$$(\hat{\Omega}_m \cdot \nabla)I(\mathbf{r}, \hat{\Omega}_m) = -(\kappa + \sigma)I(\mathbf{r}, \hat{\Omega}_m) + \frac{\sigma}{4\pi} \sum_{k=1}^M w_k I(\mathbf{r}, \hat{\Omega}_k) + \kappa I_b(\mathbf{r}). \quad (3)$$

A general control volume approximation can be made for the above equation. For the purpose of illustration, we use a Cartesian mesh with zones of size h_x and h_y along the x - and y -directions. Eq. (3) can then be written in flux conservative form as

$$\frac{1}{h_x} (F_{m,i+1/2,j} - F_{m,i-1/2,j}) + \frac{1}{h_y} (G_{m,i,j+1/2} - G_{m,i,j-1/2}) = S_{m,i,j}, \quad (4)$$

where the lumped source term on the right-hand side of Eq. (4) is given by

$$S_{m,i,j} = -(\kappa_{i,j} + \sigma_{i,j})I_{m,i,j} + \frac{\sigma_{i,j}}{4\pi} \sum_{k=1}^M w_k I_{k,i,j} + \kappa_{i,j} I_{b,i,j}. \quad (5)$$

From Eq. (4) we see that the choice of different control volume formulations is determined by the choice of the fluxes $F_{m,i-1/2,j}$ and $G_{m,i,j-1/2}$.

It is worthwhile to make a few observations about Eqs. (1) and (4):

(1) Notice that the scattering terms satisfy a principle of detailed balance. Thus the scattering terms do not contribute when one takes the integral of Eq. (1) over the full solid angle of 4π . In Eq. (1) the detailed balance is trivially satisfied. For moving media, see Mihalas and Klein [42], achieving such a detailed balance for the scattering terms requires preservation of additional moments. The same is true for anisotropic scattering even when the medium is stationary. The detailed balance in the continuous case implies that the analogous discrete ordinates equations, i.e. Eq. (4), should also satisfy these constraints as far as possible. That is the case for the ordinate sets given in the appendix and Fiveland [43].

(2) When the scattering terms are strong all ordinate directions are strongly coupled to each other. Thus in the limit of strong scattering it becomes increasingly unacceptable to relegate the scattering terms to a post-facto acceleration step.

(3) As the absorption or scattering opacities become stronger the importance of the advection terms in Eqs. (1) and (4) is decreased and the equations become more local. Thus an optimal solution strategy should converge very rapidly in that limit.

(4) Principal component analysis furnishes a well-formulated theory which can tell us which terms in a set of equations couple strongly with which other terms. All terms which have strong coupling with each other need to be relaxed simultaneously. When the principal component analysis of Yavneh [44] is applied to Eqs. (1) or (4) one realizes that the scattering terms can have strong coupling with the advection terms. Furthermore, as the mesh becomes coarser the importance of the absorption and scattering opacity terms increases. The absorption opacity is routinely included in the relaxation process so the fact that the terms involving absorption opacity dominate on coarser meshes does not damage the convergence. The scattering terms are routinely decoupled from the relaxation even on coarser meshes. The principal component analysis of Yavneh [44] shows that this is progressively more damaging to the convergence of the overall scheme as the mesh becomes coarser. The fact that the coarse mesh rebalance acceleration strategy in [33] loses its effectiveness when the scattering opacity increases can be attributed to this fact.

(5) The contributions from the absorption and scattering terms are at least formally linear in the radiation intensity variable. Thus a method that draws on this linearity should converge faster when the opacities assume larger values. In particular, the Newton method for solving a system of nonlinear equations assumes the existence of such a linearity in its solution strategy. Thus a method that builds in a Newton iteration step into its solution strategy should converge faster as the opacities increase. This is one of the reasons that motivates our choice of Newton–Krylov methods as smoothers for our numerical scheme.

(6) When the scattering opacity is zero the ordinates decouple. In that limit the pointwise smoothers devised in Sidilkover and Brandt [45] become useful. The pointwise smoothers might even retain their utility when the scattering terms are analytically differentiable, provided one is willing to carry out the analytical differentiation in every case of interest.

(7) Mihalas and Kline [42] pointed out that first-order discretizations of Eq. (1) are unacceptable. The reason they gave is that when a first-order discretization is used for the advection terms in Eq. (1) the terms that bear the resultant discretization error can contribute much like an additional radiation diffusion term. This makes it impossible to retrieve the diffusion limit at great optical depth. The asymptotic analysis of the RTE by Morel et al. [19] and Adams and Nowak [21] confirms the above-mentioned observation of Mihalas and Kline [42] and gives it a more formal

footing. For this reason, some of the first-order schemes in Jessee et al. [30] and Howell et al. [31] may be unsuitable for general purpose use.

(8) It is also important to be able to retrieve the free streaming limit when the opacities become very small. Thus an optimal scheme for radiative transfer is one which reduces the dissipation and dispersion as much as possible, retains positivity of the solution and converges very fast in the free streaming limit. Unfortunately, the first-order accurate discretizations that have been tried so far, see Fiveland [24] or Howell et al. [31], converge fast but are extremely diffusive and the second-order accurate discretizations that have been tried so far, see Jessee and Fiveland [28], have poor convergence properties.

(9) Putting the above points together we realize that the advection terms and the opacity terms interact most strongly with each other when the optical depth in a zone is of order unity. Thus we expect that a single scheme which is simultaneously optimal in the optically thick limit as well as in the free streaming limit should also be rather proficient in this intermediate regime.

(10) Thus in our quest for optimal schemes for numerical radiative transfer we seek a single scheme (i.e. without acceleration sub-iterations) which is simultaneously optimal in the optically thick and the free streaming limits.

In the next two sections we give explicit expressions for the fluxes in the dimensional TVD schemes and the genuinely multidimensional, nonlinearly positive schemes.

2.1. Dimensional TVD scheme

The TVD methods of vanLeer [26] and Harten [27] are applied on a dimension by dimension basis to obtain the fluxes. The x -flux is given by

$$\begin{aligned}
 F_{m,i-1/2,j} &= \Omega_{x,m}(I_{m,i-1,j} + 0.5 \text{ Limiter}[I_{m,i,j} - I_{m,i-1,j}, I_{m,i-1,j} - I_{m,i-2,j}]) \\
 &\quad \text{when } \Omega_{x,m} \geq 0 \\
 &= \Omega_{x,m}(I_{m,i,j} - 0.5 \text{ Limiter}[I_{m,i+1,j} - I_{m,i,j}, I_{m,i,j} - I_{m,i-1,j}]) \\
 &\quad \text{when } \Omega_{x,m} \leq 0,
 \end{aligned} \tag{6}$$

where $\text{Limiter}[x, y]$ is any symmetric slope-limiter function that keeps the slope within the TVD region, see Sweby [46]. Good examples of TVD limiters include the minmod, vanLeer or superbee limiters. It might even prove useful to use the vanAlbada limiter “ $\text{Limiter}[x, y] = (xy^2 + yx^2)/(x^2 + y^2)$ ” because of its differentiability even though it is not strictly speaking a TVD limiter. The y -flux can be obtained similarly. We have used both the minmod and vanAlbada limiters in all the simulations that will be presented in this paper.

Several useful points can be made about the flux in Eq. (6):

(1) Chakravarthy [47] and Harten [48] have shown that in one dimension such fluxes retain diagonal dominance. In multiple dimensions the diagonal dominance cannot be ensured.

(2) Mulder and vanLeer [49] have shown that Gauss–Seidel relaxation with FAS multigrid acceleration can iterate such problems to convergence. However, the convergence rates are slow. Sidilkover and Brandt [45] give some reasons for that fact. As a result, numerical schemes for solving the RTE that use the flux in Eq. (6) will not be optimal in the free streaming limit.

(3) Gropp et al. [50] have shown that Newton–Krylov methods can be used to iterate such dimensional TVD schemes to convergence. However, the convergence is again sub-optimal. This further reinforces the notion that dimensional TVD approaches yield sub-optimal numerical schemes for solving the RTE.

(4) The exploration of the convergence properties of the RTE when the vanAlbada limiter is used is very important for two specific reasons: (a) The vanAlbada limiter is a differentiable limiter and is, therefore, most beneficial to the differentiation that is carried out in the Newton step of the Newton–Krylov method used in this paper. (b) It can be shown that the lower-order weighted essentially non-oscillatory (WENO) schemes, see Liu et al. [51], Jiang and Shu [52] and Balsara and Shu [53], effectively have a limiter that is much like the vanAlbada limiter. These WENO schemes produce discretizations that are much better than second-order accurate when the solution is smooth. Thus, if it can be shown that the RTE converges well when the vanAlbada limiter is used, it would suggest that the RTE may converge well when WENO discretizations are used. The use of WENO discretizations for the RTE would in turn open the door for designing very high accuracy schemes for the RTE. We will explore this topic in a subsequent paper.

2.2. Genuinely multidimensional, nonlinearly positive scheme

Goodman and LeVeque [54] have shown that a literal extension of the TVD concept to multiple dimensions proves to be too restrictive in the design of genuinely multidimensional second-order accurate schemes. For that reason, several authors have focussed attention on schemes that retain positivity of the solution. Schemes that retain some concept of multidimensionality have been designed by numerous authors, Colella [55], Hirsch and van Ransbeek [56], LeVeque [57] to name but a few. However, very systematically designed strategies for constructing genuinely multidimensional schemes have mainly been offered in the work of Roe [58], Roe and Sidilkover [59], Sidilkover and Brandt [45], Sidilkover [60] and Sidilkover and Roe [38]. Those authors impose the twin requirements of positivity and multidimensionality which result in diagonally dominant discretizations and compact stencils. Sidilkover and Brandt [45] have also shown that the multidimensional methods are extremely well suited for use in multigrid methods. The genuinely multidimensional discretizations routinely produce convergence in substantially fewer multigrid cycles than the dimensional TVD discretizations. Roe and Sidilkover [59] show that the dissipation and dispersion characteristics of these methods are superior to those of the dimensional TVD discretizations. Sidilkover [60] showed that these methods can be extended to three dimensions. For all those reasons, we use the fluctuation splitting schemes given in Roe [58] and Sidilkover and Roe [38] in this work. The stencils that are used in these schemes depend on the characteristic directions. The fluxes, which retain second-order accuracy, are given by the following equations.

When $\Omega_{x,m} \geq 0$; $\Omega_{y,m} \geq 0$; $\Omega_{x,m}/h_x \leq \Omega_{y,m}/h_y$ we have

$$F_{m,i-1/2,j} = \Omega_{x,m} I_{m,i-1,j-1} + 0.5 \text{ Limiter} \left[\Omega_{x,m} (I_{m,i,j} - I_{m,i-1,j-1}) - h_x S_{m,i,j}, \right. \\ \left. \left(\Omega_{x,m} - \frac{\Omega_{y,m} h_x}{h_y} \right) (I_{m,i-1,j} - I_{m,i-1,j-1}) \right], \tag{7}$$

$$G_{m,i,j-1/2} = \Omega_{y,m} I_{m,i,j-1}. \tag{8}$$

When $\Omega_{x,m} \geq 0$; $\Omega_{y,m} \geq 0$; $\Omega_{x,m}/h_x \geq \Omega_{y,m}/h_y$ we have

$$F_{m,i-1/2,j} = \Omega_{x,m} I_{m,i-1,j}, \quad (9)$$

$$G_{m,i,j-1/2} = \Omega_{y,m} I_{m,i-1,j-1} + 0.5 \text{ Limiter} \left[\Omega_{y,m} (I_{m,i,j} - I_{m,i-1,j-1}) - h_y S_{m,i,j}, \right. \\ \left. \left(\Omega_{y,m} - \frac{\Omega_{x,m} h_y}{h_x} \right) (I_{m,i,j-1} - I_{m,i-1,j-1}) \right]. \quad (10)$$

The above expressions give us the fluxes when the characteristics point into the first quadrant of the xy -plane. When the characteristics point into other quadrants of the xy -plane, the fluxes can be obtained by applying obvious symmetries to Eqs. (7)–(10).

Several useful points can be made about the fluxes in Eqs. (7)–(10):

(1) Sidilkover and Brandt [45] have shown that when these fluxes are used in a multigrid scheme the convergence is extremely rapid. The reasons for this rapid convergence stem from the compact stencil and the diagonal dominance that are mandated by the twin requirements of positivity and genuine multidimensionality. As a result, we use these fluxes in our scheme for solving the RTE because they allow us to obtain an optimal scheme in the free streaming limit.

(2) Use of the fluxes in Eqs. (7)–(10) yields an optimal scheme for another reason. Harten [61] showed that the cross-characteristic smearing produced by a numerical scheme is strongly determined by the dissipation structure of the scheme. This point has relevance for radiative transfer because in the free streaming limit the ability to accurately advect beams of radiation over several hundreds of zones is entirely determined by the cross-characteristic smearing. Roe and Sidilkover [59] and Sidilkover and Brandt [45] have shown that the cross-characteristic smearing of the present scheme is the least that is theoretically possible. Thus these schemes are optimal even from the point of view of having minimal dissipation.

(3) Sidilkover and Brandt [45] used a pointwise relaxation strategy as a smoother within a FAS multigrid scheme to rapidly iterate an advection equation to convergence. Pointwise smoothers may not be practicable when the zones have a bad aspect ratio or when complex scattering terms are involved. However, the diagonal dominance and compact stencils that result from this discretization would also help in the convergence of Newton–Krylov techniques. Newton–Krylov techniques remain practicable when the zones have a bad aspect ratio or when complex scattering terms are involved.

(4) Sidilkover and Roe [38] include source terms in the slope limiters in Eqs. (7) and (10) above in a slightly different fashion. Their way of including source terms does not result in properly centered representation of the source terms. For that reason, the form of the fluxes given here is to be preferred. It is also inadvisable to use the source terms in the limiters in Eqs. (7) and (10) when solving the problem on the coarser meshes in a multigrid hierarchy. This is so because, before the solution comes close to the converged solution, the source terms can take on large and unrealistic values. Even for a homogeneous equation the fine grid to coarse grid defect can make a large contribution to the source terms. We do not want those values to influence the fluxes. For that reason, we include the source terms in the limiters on the fine mesh that covers the entire computational domain but we exclude the source terms in the limiters on any of its coarsenings. We

have also explored the consequences of excluding the source terms from the limiters on all the meshes in the multigrid hierarchy and report on this topic later on in the paper.

(5) It is worthwhile pointing out that the original formulation of the fluctuation splitting scheme that was carried out by Roe [58] was done on unstructured meshes. While there is some ambiguity in the choice of triangles in that work the fact that it can be formulated on triangulated meshes holds out the hope that the methods for numerical radiative transfer devised here will see natural extensions to unstructured meshes.

(6) Raithby and Chui [62] have made an early effort to introduce some amount of multidimensionality into the discretization of the RTE. In the free streaming limit their method reduces to the skew upwind difference scheme of Raithby [63]. Roe and Sidilkover have shown that the method of Raithby [63] loses second-order accuracy for several ordinate directions. Thus the method of Raithby and Chui [62] is fundamentally a first-order method.

3. Relevant iterative techniques — a brief survey

In this section we briefly review Newton–Krylov methods and FAS multigrid. The discussion is meant to give the interested reader an entry point to the literature without repeating too many details in this paper. We also try wherever possible to provide insights that are specific to the solution of the RTE.

3.1. Restarted Newton–Krylov methods

Eq. (4) with the fluxes given in Sections 2.1 or 2.2 defines a nonlinear system of equations that has to be solved by iterating it to convergence. Denote the vector formed by the intensities in each zone and for every ordinate by U . Let $F(U)$ denote the vector of residuals associated with Eq. (4) in each zone and for every ordinate. The solution of the RTE is said to have converged when $|F(U)| \rightarrow 0$. Newton’s method provides a way for solving this nonlinear problem iteratively. Thus given the k th iterate U_k we wish to obtain the next iterate U^{k+1} via the solution of the linear system:

$$J^k \delta U^k = -F(U^k); \quad U^{k+1} = U^k + \delta U^k. \quad (11)$$

Here J^k is the Jacobian matrix for the current iterate. In general J^k is a large, sparse and non-symmetric matrix whose (i,j) th element in the k th iteration is given by

$$J_{i,j}^k = \frac{\partial F_i(U^k)}{\partial U_j^k}. \quad (12)$$

Such sparse matrices can be solved using the restarted generalized minimum residual method (GMRES) of Saad and Schultz [39]. The restarted GMRES(p) algorithm is an iterative method which seeks to find the solution of the linear system in Eq. (11) by projecting it onto the Krylov subspace $\{r_0, J^k r_0, (J^k)^2 r_0, (J^k)^3 r_0, \dots, (J^k)^{p-1} r_0\}$ where the vector of residuals r_0 , using an initial guess of zero, is given by $r_0 = -F(U^k)$. The projection is designed to minimize the residuals in the L_2 norm. After each complete projection the GMRES method is restarted. Here “ p ” denotes the dimension of the Krylov subspace used in each restarted GMRES(p) iteration. At the beginning of each restart of the GMRES method the Newton step from Eq. (11) is applied making this

a Newton–Krylov method. Brown [64] and Brown and Saad [40] pointed out that the Jacobian matrix J^k is never needed explicitly. All that is needed is the product of the Jacobian matrix with a Krylov vector V . This matrix–vector product can be obtained via an inexact Newton procedure for which we have the following approximation which holds true for small enough ε :

$$J^k V \approx [F(U^k + \varepsilon V) - F(U^k)]/\varepsilon. \quad (13)$$

Gropp et al. [50] catalogue some very useful strategies for choosing ε in a solution-dependent fashion.

Several useful points about the Newton–Krylov method that we have used in this work are given below:

(1) The GMRES method given by Saad and Schultz [39] is susceptible to numerical instabilities in the Gram–Schmidt orthogonalization procedure. A variation of GMRES that uses Householder transformations for the orthogonalization has been given by Walker [65]. In practical problems, Saad [66] suggests that the modified Gram–Schmidt method with reorthogonalization that is catalogued in the text of Parlett [67] is almost always adequate for rectifying this numerical instability. For all the radiative transfer problems presented here we have found Saad’s suggestion to be borne out.

(2) The dimension of the Krylov subspace, “ p ”, determines the amount of computer memory required by the method. Typically the memory required is approximately “ p ” times the memory required to store the entire solution vector. In general, larger values of “ p ” are known to produce better convergence. Thus the value of “ p ” is always a compromise between the computer’s available memory and the convergence rate that is desired. Based on several CFD calculations Gropp et al. [50] favor using “ $p \approx 30$ ”. For radiative transfer calculations we have found that using two restarted Newton-GMRES(15) iterations works just as well as one restarted Newton-GMRES(30) iteration. For that reason we prefer the smaller value of “ p ” with more frequent restarts. More frequent restarts can also be helpful in the initial phase of convergence because each restart allows us to start with a value of the solution that is closer to the final converged solution.

(3) Saad [68] has also designed preconditioning strategies for GMRES. We have not used those preconditioners in this work because the FAS multigrid method provides an equivalent strategy that is much better suited for the RTE.

(4) The experience of Gropp et al. [50] has shown that the convergence of Newton–Krylov methods is not restricted to two dimensional problems. Thus the Newton–Krylov methods converge well when applied to two-dimensional problems and also when applied to three-dimensional problems.

3.2. Full approximation scheme (FAS) multigrid

Radiation has the very interesting physical property that all parts of the computational domain can be radiatively coupled to all other parts of the computational domain. This is especially true when the problem is not very optically thick. Multigrid methods have the useful mathematical property that they represent the physical problem at all levels of coarsening and, therefore, allow all length scales to be coupled to each other. Because of that good match between the mathematical properties of multigrid and the physical property of the radiation field, we use multigrid methods in

this work. The full approximation scheme (FAS) multigrid and its utility for solving nonlinear problems has been extensively documented in Brandt [41]. For that reason we do not repeat any further details here.

Several useful points about the use of FAS multigrid in conjunction with a Newton–Krylov method are given below:

(1) When traditional FAS multigrid is applied to nonlinear equations it often uses a disjoint Newton step applied to the individual equations followed by relaxation sweeps applied to the resultant Jacobian matrix. The present method synthesizes the Newton and relaxation steps into one step and uses the two together in a much more powerful fashion. It also absolves the user of having to go through the inconvenience of explicitly constructing the Jacobian matrix. Avoiding the explicit construction of the Jacobian matrix also represents a very significant reduction in the amount of computer memory that is needed.

(2) Only a very small number of restarted Newton-GMRES iterations seem to be needed on each of the upward and downward passes in the multigrid cycling. Furthermore, using a Krylov subspace with a rather small dimension seems to be wholly adequate in all cases that we tested. It is popular among some users of Krylov methods to use the DQGMRES variant of the GMRES algorithm because it saves computational cost by doing an incomplete orthogonalization of the Krylov subspace. In view of the fact that a rather small dimension for the Krylov subspace seems to be adequate, it is not beneficial to use DQGMRES instead of GMRES for radiative transfer problems.

(3) The use of restarted GMRES as a smoother in a multigrid scheme seems more natural once one realizes that Kaczmarz [69] relaxation is equivalent to utilizing a Krylov subspace having a dimension of unity.

(4) When GMRES is used as a smoother in a multigrid hierarchy the coarsest mesh in the multigrid hierarchy does not have to be an extremely coarse mesh. Having just a few levels in the multigrid hierarchy has often been seen to suffice in several problems. This helps overcome one of the popular criticisms of multigrid methods that the number of zones in each mesh direction can only be a small multiple of a power of two.

(5) When the mesh has zones with bad aspect ratios the GMRES algorithm (used as a smoother within FAS multigrid) is not as susceptible to convergence problems as the pointwise smoothers that are traditionally used in multigrid work. In particular, multigrid methods that use GMRES as a smoother do not seem to need directional coarsenings which increase the programming complexity of the algorithm.

(6) Problems with strong scattering terms are also easily handled. As the absorption *or scattering* opacities become large the convergence is improved.

(7) Putting the above two points together we claim that using the Newton–Krylov method as a smoother within FAS multigrid is a natural way of allowing the computer to carry out the principal component analysis. Thus, by permitting the maximal coupling between variables, the Newton–Krylov method naturally picks out the variables that strongly couple to each other in the relaxation process. This point might also be very useful for more complicated hyperbolic systems where the principal component analysis is difficult to carry out analytically. We remind the reader that the algebraic multigrid (AMG) method of Brandt et al. [70] also seeks to coarsen the mesh points based on evaluating strong and weak coupling between variables. However, it might prove difficult to incorporate the Newton step within AMG for general partial differential equations. The

extra logic that is required to evaluate couplings between variables on a computer has also impeded the use of AMG, Brandt [71]. Thus methods that use Newton–Krylov-based smoothers within FAS multigrid may well turn out to be general-purpose alternatives to the AMG method.

(8) When the genuinely multidimensional scheme in Section 2.2 is used in conjunction with the FAS multigrid scheme that uses the Newton–Krylov smoothers one gets a single scheme for numerical radiative transfer that is simultaneously optimal in the optically thick and free streaming limits. The resulting scheme is second-order accurate and has minimal dissipation and dispersion. It retains its good convergence rate when the zones have unit optical depth on the fine mesh. Indeed the use of a multigrid hierarchy ensures that there is always a mesh in the hierarchy where the optical depth in the zones becomes either small or large. Such meshes help accelerate the convergence of the overall radiative transfer problem on the fine mesh.

(9) The equations of radiative transfer are often solved in conjunction with the equations for the evolution of the irradiated matter. The equations for the matter can also draw on the existence of a multigrid hierarchy and the Newton–Krylov methods. For example, Gropp et al. [50] have shown that the equations of hydrodynamics can be solved using implicit methods that also use the restarted Newton–Krylov methodology.

(10) In this paper we do not explore the role of coarsening the frequency groups on coarser meshes. This would seem to be a useful thing to do because it helps produce better thermal coupling across frequency groups. This idea does not seem far-fetched when we realize that we only want the final solution to be available for all the frequency groups on the base level (i.e. finest) mesh in the multigrid hierarchy. The solution on all the coarsenings of the finest mesh only serve to accelerate the fine mesh solution to convergence. This idea also has the mathematical merit that it reduces the size of the vectors that form the Krylov subspace on the coarser meshes, making convergence easier to achieve on those meshes. We will explore this idea in a subsequent paper.

(11) The previous idea of coarsening the frequency groups on coarser meshes can also be extended to using a smaller ordinate set on coarser meshes. Morel and Manteuffel [72] have performed angular coarsening on the same spatial mesh and shown that convergence rates that are superior to diffusion synthetic acceleration can be obtained. We will explore this idea in a subsequent paper.

4. Model problems

In this section we focus on a few model problems. This is done with the intent of demonstrating the performance characteristics of our methods for numerical radiative transfer in an extremely simple setting. The simplicity permits us to pick out isolated effects. The model problems are based on solving the following equation:

$$a \frac{\partial u(x, y)}{\partial x} + b \frac{\partial u(x, y)}{\partial y} = -\kappa u(x, y) + s. \quad (14)$$

The above equation is solved with different values of the parameters “ $(a, b, \kappa$ and s ” which are analogous to the ordinates, opacity and thermal emissivity, respectively. The problem is solved either on the unit square $[0,1] \times [0,1]$ or on the domain $[0,1] \times [0,0.1]$ using a 256×256 zone mesh

in both situations. In the latter domain we have zones with a 10:1 aspect ratio allowing us to test the performance characteristics of the schemes when zones with very bad aspect ratios are present in the computation. In all cases we have used $a = 1$ and $b = 2$. As a result, inflow boundary conditions have to be specified on the lower and left boundaries. To make the problem nontrivial, a step function is specified on the lower y -boundary so that the boundary conditions are given by

$$\begin{aligned} u(x = 0, y) &= 1.0, \\ u(x, y = 0) &= 1.0 \quad \text{when } x \leq 0.2, \\ &= 0.1 \quad \text{when } x > 0.2. \end{aligned} \tag{15}$$

The model problems were always started with a value of 1.2 initialized on the fine grid. As a result the schemes we tested were required to converge from initial values that were more than 100% different from the converged values in all cases. Using 256×256 zones to solve some of the problems in this section might seem like we are using an excessive number of zones. However, the choice is justified because on grids with smaller number of zones the GMRES convergence (without multigrid) can show improvements that are not sustained on larger grids. In most cases we were interested in an inter-comparison of the convergence rates for the following solution strategies:

(1) Restarted Newton-GMRES(15) was used with the dimensional TVD scheme of Section 2.1. The minmod limiter was used. This represents in some sense the most conventional of technologies and, therefore, provides us with a baseline to compare against. We call this method “convergence strategy 0”.

(2) Restarted Newton-GMRES(15) was used as a smoother within the FAS multigrid and the dimensional TVD discretization from Section 2.1 was used for the fluxes. The minmod limiter was used. Two restarted Newton-GMRES(15) smoothing steps were carried out in the upward pass of a multigrid V-cycle and a similar number were used in the downward pass. We call this method “convergence strategy 1”. Twenty multigrid V-cycles were used.

(3) Restarted Newton-GMRES(15) was used as a smoother within the FAS multigrid and the genuinely multidimensional discretization from Section 2.2 was used for the fluxes. The minmod limiter was used and the source term was included in the limiter. Two restarted Newton-GMRES(15) smoothing steps were carried out in the upward pass of a multigrid V-cycle and a similar number were used in the downward pass. We call this method “convergence strategy 2”. Twenty multigrid V-cycles were used.

(4) Restarted Newton-GMRES(15) was used as a smoother within the FAS multigrid and the genuinely multidimensional discretization from Section 2.2 was used for the fluxes. The vanAlbada limiter was used and the source term was included in the limiter. Two restarted Newton-GMRES(15) smoothing steps were carried out in the upward pass of a multigrid V-cycle and a similar number were used in the downward pass. We call this method “convergence strategy 3”. Twenty multigrid V-cycles were used. It is interesting to make a comparison between the present point and the immediately preceding one because the vanAlbada limiter is a differentiable limiter and is, therefore, most beneficial to the differentiation that is carried out in the Newton step of the Newton–Krylov method.

We were primarily interested in evaluating the rates at which the residuals from the above four solution strategies converge. All four strategies are iterative. The solution strategies entail

evaluating the residual of Eq. (4) for each mesh point and each ordinate (and frequency, if multiple frequency groups are used). This is the computationally intensive step. Thus the cost of the scheme is determined by the number of times the residual of Eq. (4), viewed as a function, is evaluated for the whole spatial and angular mesh. An optimal solution strategy is one whose residual converges with the smallest number of function evaluations. Thus we use the number of function evaluations as the unit of work. In all cases we were interested in the convergence of the residuals evaluated in the L_∞ and L_2 norms with respect to the number of function evaluations. Several of these model problems have sharp singularities and large gradients in their solutions. The residual evaluated in the L_∞ norm is a good indicator of the error that builds up in singular regions. The residual evaluated in the L_2 norm is a better index of the error in the solution averaged over the whole computational domain. As a result, both are of interest here. Since we used two restarted Newton-GMRES(15) smoothing steps in each upward and downward pass of the multigrid V-cycle, each multigrid cycle resulted in sixty function evaluations on the finest grid. For a multigrid V-cycle, this is equivalent to eighty function evaluations for each traversal of the entire multigrid hierarchy.

4.1. The free streaming limit; $\kappa = s = 0$

The free streaming limit is one of the very important limits of the RTE. The reason for its importance is that when the opacity and emissivity of the matter are very small the radiation must be accurately propagated along rays over extremely large distances. Thus small advection errors in each zone of the discretized RTE can cumulatively pile up as the radiation propagates over a large number of zones. In this model problem, the step that is initialized on the lower y -boundary in Eq. (15) is propagated along the characteristics over the entire computational mesh. Fig. 1a shows the evolution of the residual evaluated in the L_∞ norm as a function of the number of function evaluations. The restarted Newton-GMRES(15) with the dimensional TVD scheme is shown with a legend of “0” and corresponds to convergence strategy 0. The FAS multigrid with the dimensional TVD discretization is shown with a legend of “1” and corresponds to convergence strategy 1. The FAS multigrid with the multidimensional discretization and minmod limiter is shown with a legend of “2” and corresponds to convergence strategy 2. The FAS multigrid with the multidimensional discretization and vanAlbada limiter is shown with a legend of “3” and corresponds to convergence strategy 3. Each legend of “1”, “2” or “3” gives the residual after one entire FAS multigrid cycle. Fig. 1b shows the evolution of the residual evaluated in the L_2 norm. Both Figs. 1a and b show that the convergence strategy 0 is very slow especially in the initial iterations. By comparison, convergence strategy 1 is much faster especially in the initial iterations. This is because the multigrid method provides an effective way of reducing the residuals in all wavelengths of the problem. The rapid initial convergence of the multigrid method is important because it is often considered meaningful to iterate the problem only till the residuals become smaller than the discretization error in the problem. When that is the case, strategy 1 handily beats strategy 0. We see, however, that once the residual has become small enough the restarted Newton-GMRES(15) (without multigrid) does indeed show the fast, quadratic convergence of the Newton method. Convergence strategies 2 and 3 are seen to converge at roughly comparable rates. It is also worth pointing out that convergence strategies 2 and 3 handily beat convergence strategies 0 and 1. This is because convergence strategies 2 and 3 have the dual advantages of using a compact, positivity




Fig. 1. (a) and (b) show the residual convergence histories in the L_∞ and L_2 norms for $\kappa = s = 0$ which corresponds to the model problem in the free streaming limit. The legend “0” indicates restarted Newton-GMRES(15) (without multigrid) with dimensional TVD discretization. The legend “1” indicates FAS multigrid using restarted Newton-GMRES(15) as a smoother for the dimensional TVD discretization. The legends “2” and “3” indicate FAS multigrid using restarted Newton-GMRES(15) as a smoother for the multidimensional discretization. Legend “2” corresponds to the use of a minmod limiter while legend “3” corresponds to the use of a vanAlbada limiter.

preserving discretization as well as using the convergence acceleration provided by multigrid cycling. Convergence strategies 2 and 3 could actually reduce the residual below the discretization error in about six multigrid cycles even though the initial guess was several 100% wrong. This underscores the great advantage of genuinely multidimensional discretizations over conventional, dimensional TVD discretizations. Not just do they converge faster but they are more faithful to the physics of the problem. The convergence of the dimensional TVD scheme is poor mainly because it tries to represent the advection as if it were taking place along mesh lines. But the advection is not oriented along the mesh lines in this problem causing the convergence of the dimensional TVD scheme to degrade. It is also interesting to notice that the L_∞ norm of the residual in convergence strategy 0 converges with a lot of fluctuations while the L_2 norm of the residual in convergence strategy 0 has a smoother convergence history. This is purely a consequence of the GMRES algorithm which is designed to minimize the error in the L_2 norm of the residual.

4.2. Low opacity; $\kappa = s = 1$

It is interesting to ask how the inclusion of opacity terms and source terms in the Newton-Krylov relaxation interacts with the discretization of the advection terms? We want to verify that our inclusion of opacity and source terms in the Newton-GMRES smoother does not have a negative impact on the convergence characteristics of the scheme. The present problem is designed to test the convergence when opacity and source terms are included. In principle, the

radiative intensity should reach an asymptotic value of s/κ after traversing several optical depths. The optical depth across this box is, however, of order unity so we do not expect it to reach that asymptotic value in this problem. In Figs. 2a and b we show the convergence histories of the residuals. We see that the convergence history does not look significantly different from the convergence history in Figs. 1a and b. This is consistent with our anticipation in the previous sections because the optical depth in each zone is much less than unity. Thus we expect the opacity term to be weakly coupled to the advection term which is the dominant term in this model problem. As a result, the convergence history in Fig. 2 looks much like that in Fig. 1, as expected. Convergence strategy 3 shows itself to be superior to convergence strategy 2 in this problem. This can be attributed to the fact that the vanAlbada limiter is differentiable limiter. Since the source terms are included in the limiter, the vanAlbada limiter is more responsive to the influence of the source terms, thereby producing better convergence. (When the source terms are excluded from the minmod limiter the results, which are not shown here, display a convergence rate that is comparable to convergence strategy 3.) Our inclusion of opacity and source terms in the Newton-GMRES smoother has not had a negative impact on the convergence characteristics of the scheme.

4.3. Intermediate opacity; $\kappa = 100$; $s = 10$

In Section 2.2 we had asserted that an optimal scheme for radiative transfer should converge rapidly in the free streaming limit. In Section 4.1 we showed that the multidimensional scheme indeed converges very well in that limit. Based on our discussion in Section 2 we also expect the Newton-GMRES smoother to cause the scheme to converge rapidly when the optical depth in a zone is larger than unity. In Section 2 we had, therefore, conjectured that a single scheme that is optimal in the free streaming limit and the optically thick limit should also have good convergence properties in the intermediate regime where the optical depth in each zone is of order unity. We test that conjecture here because the opacity in the present model problem is chosen so that the optical depth in each zone is of order unity. Figs. 3a and b show the convergence histories of the residuals. We see that convergence strategies 0 and 1 converge extremely fast and at roughly the same rate. This is to be expected because the convergence is driven by the fact that the high opacity makes this problem a local one in each of the zones. The Newton-GMRES smoother utilizes that fact to converge very rapidly. Thus the convergence acceleration from the FAS multigrid algorithm plays virtually no role in this intermediate opacity case. We also see that convergence strategies 2 and 3 are much better than convergence strategies 0 and 1. Thus the genuinely multidimensional scheme retains its advantages in the intermediate opacity case.

4.4. High opacity; $\kappa = 1000$; $s = 100$

In Figs. 4a and b we examine the convergence rates of the schemes in the limit where the zones are optically thick. For reasons that have already been given in Section 4.3 we verify that the convergence rate is excellent for all four convergence strategies. Comparing the convergence history of the residuals in Figs. 2–4 with that in Fig. 1 we see that the convergence in Figs. 2–4 is as good as or better than the convergence in Fig. 1. Thus, for the model problems catalogued so far, the convergence rate is *bounded from above* by the convergence rate in the free streaming limit! We see that convergence strategy 3 is optimally convergent in the free streaming limit as well as the




Fig. 2. (a) and (b) show the residual convergence histories in the L_∞ and L_2 norms for $\kappa = s = 1$ which corresponds to the low-opacity model problem. The legend “0” indicates restarted Newton-GMRES(15) (without multigrid) with dimensional TVD discretization. The legend “1” indicates FAS multigrid using restarted Newton-GMRES(15) as a smoother for the dimensional TVD discretization. The legends “2” and “3” indicate FAS multigrid using restarted Newton-GMRES(15) as a smoother for the multidimensional discretization. Legend “2” corresponds to the use of a minmod limiter while legend “3” corresponds to the use of a vanAlbada limiter.

optically thick limit. Furthermore, its convergence does not degrade when the optical depth of each zone is of order unity. (When the source terms are excluded from the minmod limiter the results, which are not shown here, display a convergence rate that is comparable to convergence strategy 3 in all the model problems analyzed thus far.) *As a result we identify FAS multigrid which uses the restarted Newton–Krylov method as a smoother for the genuinely multidimensional discretization of the RTE as our choice of an optimal scheme. We also notice that it is advantageous to use a differentiable limiter.*

4.5. Strong lasing; $\kappa = -10$; $s = 0$

In this section we examine the convergence properties of the schemes in a problem that mimics lasing phenomena. Here the opacity is set to a large negative value. Thus as the radiation propagates from the lower y -boundary to the upper y -boundary it undergoes an exponential amplification. The radiation undergoes a cumulative amplification by a factor of e^5 . Because of the strong exponential growth, this model problem cannot be done on a mesh that is substantially smaller than 256×256 zones. Figs. 5a and b present the convergence histories. We see that the Newton-GMRES method without multigrid acceleration (legend “0”) does not converge at all! Convergence strategy 1, which uses dimensional TVD discretization, converges exceedingly slowly. This is because the TVD limiter, which is applied on a dimension-by-dimension basis, does not




Fig. 3. (a) and (b) show the residual convergence histories in the L_∞ and L_2 norms for $\kappa = 100$; $s = 10$ which corresponds to the intermediate opacity model problem. The legend “0” indicates restarted Newton-GMRES(15) (without multigrid) with dimensional TVD discretization. The legend “1” indicates FAS multigrid using restarted Newton-GMRES(15) as a smoother for the dimensional TVD discretization. The legends “2” and “3” indicate FAS multigrid using restarted Newton-GMRES(15) as a smoother for the multidimensional discretization. Legend “2” corresponds to the use of a minmod limiter while legend “3” corresponds to the use of a vanAlbada limiter.

detect the fact that the radiation intensity is undergoing exponential amplification in a direction that is not aligned with the mesh. Thus the limiter tries to damp out exponential growth even though such growth is physically justified in this model problem. Convergence strategy 2 shows a modest improvement over convergence strategy 1. Convergence strategy 3, which uses the genuinely multidimensional scheme and vanAlbada limiter in conjunction with the multigrid acceleration and the Newton-GMRES smoother, is the only scheme that has a healthy convergence rate. (When the source terms are excluded from the minmod limiter in this strong lasing problem the results, which are not shown here, display a convergence rate that is not as good as convergence strategy 3.) Note that in this model problem the convergence rate is not bounded from above by the convergence rate shown in Fig. 1 for the free streaming limit. However, a problem that has strong lasing cannot be thought of as being in the same category as problems with a positive opacity because lasing problems permit exponential growth in the radiation intensity.

4.6. Zones with bad aspect ratio; $\kappa = s = 0$ and $\kappa = 100$; $s = 10$

In this section we examine the convergence characteristics when zones with a bad aspect ratio (i.e. 10:1) are used. Figs. 6a and b present the convergence histories in the free streaming limit. We see again that the Newton-GMRES method without multigrid acceleration (legend “0”) does not converge too well. Convergence strategies 1–3 converge rather well with convergence strategy




Fig. 4. (a) and (b) show the residual convergence histories in the L_∞ and L_2 norms for $\kappa = 1000$; $s = 100$ which corresponds to the high-opacity model problem. The legend “0” indicates restarted Newton-GMRES(15) (without multigrid) with dimensional TVD discretization. The legend “1” indicates FAS multigrid using restarted Newton-GMRES(15) as a smoother for the dimensional TVD discretization. The legends “2” and “3” indicate FAS multigrid using restarted Newton-GMRES(15) as a smoother for the multidimensional discretization. Legend “2” corresponds to the use of a minmod limiter while legend “3” corresponds to the use of a vanAlbada limiter.

2 having a distinct advantage. This is the only test problem we have tried where convergence strategy 3 produced convergence that was inferior to that of convergence strategy 2. Since 20 multigrid cycles were used in all model problems it is worthwhile making a direct comparison between Figs. 6 and 1. After 20 multigrid cycles the residuals in the free streaming problem are reduced by almost fifteen orders of magnitude on a uniform mesh when convergence strategy 2 is used. On a mesh that has zones with bad aspect ratio, the same twenty multigrid cycles with convergence strategy 2 reduce the residuals by almost eight orders of magnitude. Thus we see that there is some degradation in the convergence rate when the mesh has zones with bad aspect ratios. However, we also see that this degradation in the convergence rate is not too extreme.

Figs. 7a and b present the convergence histories when an opacity of $\kappa = 100$ is used. This is an intermediate opacity case. As a result, all the general trends that we observed in Section 4.3 are repeated here. Thus we see that all four methods display rather good convergence rates even when the mesh has zones with bad aspect ratios.

5. Test problems

In this section we focus on a few test problems with pure absorption opacities. We also focus on several test problems with a pure scattering opacity.




Fig. 5. (a) and (b) show the residual convergence histories in the L_∞ and L_2 norms for $\kappa = -10$; $s = 0$ which corresponds to the strong lasing model problem. The legend “0” indicates restarted Newton-GMRES(15) (without multigrid) with dimensional TVD discretization. The legend “1” indicates FAS multigrid using restarted Newton-GMRES(15) as a smoother for the dimensional TVD discretization. The legends “2” and “3” indicate FAS multigrid using restarted Newton-GMRES(15) as a smoother for the multidimensional discretization. Legend “2” corresponds to the use of a minmod limiter while legend “3” corresponds to the use of a vanAlbada limiter.

5.1. Pure absorption opacity

When the problem has a pure absorption opacity, i.e. $\sigma = 0$ in Eq. (1), the ordinates decouple from each other. Thus the convergence properties of problems with pure absorption opacity are no different from the convergence properties of some of the model problems in Section 4 where a single ordinate was examined. For that reason we do not show too many examples of the pure absorption opacity case. Figs. 8a and b show the mean, angle-averaged intensity when $\kappa = 1$ and 5, respectively. In these test problems the absorption opacity κ was assigned a constant value throughout the medium. The thermal emissivity was initialized so that $\kappa I_b(\mathbf{r}) = 1/\pi$ in the medium. The problems were solved on the unit square using a 48×48 zone mesh. The S_6 LSH ordinate set from the appendix was used. The walls of the square had unit surface emissivity and were held at zero degrees Kelvin. Similar test problems have been described in Fiveland and Jessee [23]. In the optically thick limit the mean intensity should reach a value of $1/(\kappa\pi)$ in the central part of the domain. Fig. 8a does not correspond to an optically thick problem. As a result we see that the mean intensity in the center of Fig. 8a is substantially smaller than $1/\pi$. Fig. 8b does correspond to an optically thick problem. Thus for Fig. 8b we see that in the center of the domain the mean intensity reaches a value that is quite close to the expected value of $1/(5\pi)$. Closer to the boundaries the mean intensity diminishes in value over a distance that is comparable to a unit optical depth. As a result, we see that the mean intensity diminishes rather gradually towards the boundaries in Fig. 8a. For




Fig. 6. (a) and (b) show the residual convergence histories in the L_∞ and L_2 norms for $\kappa = s = 0$ which corresponds to the model problem in the free streaming limit with zones having a bad aspect ratio of 10:1. The legend “0” indicates restarted Newton-GMRES(15) (without multigrid) with dimensional TVD discretization. The legend “1” indicates FAS multigrid using restarted Newton-GMRES(15) as a smoother for the dimensional TVD discretization. The legends “2” and “3” indicate FAS multigrid using restarted Newton-GMRES(15) as a smoother for the multidimensional discretization. Legend “2” corresponds to the use of a minmod limiter while legend “3” corresponds to the use of a vanAlbada limiter.

Fig. 8b a unit optical depth corresponds to a rather short distance. As a result, we see that the mean intensity diminishes rather rapidly towards the boundaries in Fig. 8b. Other test problems with $\kappa = 10$ and 25 were also tried and were seen to yield results that were consistent with expectations.

5.2. Pure scattering opacity

In this set of test problems we set $\kappa = 0$; $I_b(\mathbf{r}) = 0$ in Eq. (1). The scattering opacity was assigned values of $\sigma = 1, 5$ and 25 permitting us to test a substantial range of scattering opacities. The problems were solved on the unit square using resolutions that ranged from 12×12 to 192×192 mesh points. The S_6 and S_8 LSH ordinate sets from the appendix were used. The walls of the square had surface emissivities of $\epsilon = 0.3, 0.5$ and 1.0. The upper and lower y -boundaries as well as the upper x -boundary were held at zero degrees Kelvin. The lower x -boundary had $I_b(\mathbf{r}) = 1/\pi$. Figs. 9a and b show the mean, angle-averaged intensity for the cases where $\epsilon = 0.3$ and $\sigma = 5$ and 25, respectively. 192×192 mesh points along with the S_8 LSH ordinate set were used. We see that for the lower scattering opacity the radiation permeates much of the computational domain. For the higher scattering opacity the radiation has a significant mean intensity only till about half-way across the unit square.

The convergence rates of the discrete ordinates methods presented here in the pure scattering case are also of interest. This is because other practitioners, see Fiveland and Jessee [33] and




Fig. 7. (a) and (b) show the residual convergence histories in the L_∞ and L_2 norms for $\kappa = 100$; $s = 10$ which corresponds to the intermediate opacity model problem with zones having a bad aspect ratio of 10:1. The legend “0” indicates restarted Newton-GMRES(15) (without multigrid) with dimensional TVD discretization. The legend “1” indicates FAS multigrid using restarted Newton-GMRES(15) as a smoother for the dimensional TVD discretization. The legends “2” and “3” indicate FAS multigrid using restarted Newton-GMRES(15) as a smoother for the multidimensional discretization. Legend “2” corresponds to the use of a minmod limiter while legend “3” corresponds to the use of a vanAlbada limiter.

Ramone, Adams and Nowak [34], have found that the pure scattering limit is especially problematic for discrete ordinates methods. Figs. 10a and b show the residual convergence histories in the L_∞ and L_2 norms for several problems that have a pure scattering opacity. The walls of the domain have an emissivity ε of 0.3. FAS multigrid with V-cycling was used. Two restarted Newton-GMRES(15) smoothing steps were used on each of the upward and downward passes of the multigrid V-cycles. The legend “0” corresponds to dimensional TVD discretization and $\sigma = 5$. The legend “1” corresponds to the multidimensional discretization with $\sigma = 1$. The legend “2” corresponds to the multidimensional discretization with $\sigma = 5$. The legend “3” corresponds to the multidimensional discretization with $\sigma = 25$. In Figs. 10a and b the minmod limiter was used and the source terms were excluded from the limiter. 96×96 mesh points along with the S_6 LSH ordinate set were used. We see that the dimensional TVD discretization has the slowest rate of convergence. We also point out the trend that is obvious from Fig. 10 that increasing scattering opacities produce better convergence rates. Thus the multidimensional discretization with $\sigma = 25$ has the best convergence rate followed by the multidimensional discretizations with $\sigma = 5$ and lastly $\sigma = 1$. This improving convergence rate with increasing scattering opacity mirrors the similar trend with increasing absorption opacity that we found in Section 4. This is because an increasing scattering opacity, just like an increasing absorption opacity, makes the problem more local. The Newton-GMRES(15) smoothing draws on this fact and provides enhanced convergence with increasing scattering opacity. Fig. 11 shows a convergence study for the same physical problems as




Fig. 8. (a) and (b) show the mean, angle-averaged intensity for the pure absorption problems when $\kappa = 1$ and 5, respectively. The problems were solved on the unit square using a 48×48 zone mesh. The S_6 LSH ordinate set was used.




Fig. 9. (a) and (b) show the mean, angle-averaged intensity for the pure scattering problems where $\varepsilon = 0.3$ and $\sigma = 5$ and 25, respectively. 96×96 mesh points along with the S_6 LSH ordinate set were used.

were shown in Fig. 10. The only difference between them is that the vanAlbada limiter was used and the source terms were included in the limiter. We see that the dimensional TVD scheme in Fig. 11 converges faster than the dimensional TVD scheme in Fig. 10. This shows that the use of a differentiable limiter, such as the vanAlbada limiter, can also improve the convergence of




Fig. 10. (a) and (b) show the residual convergence histories in the L_∞ and L_2 norms for several problems that have a pure scattering opacity. The walls of the domain have an emissivity ε of 0.3. FAS multigrid with Newton-GMRES(15) smoothing was used in all four cases. The legend “0” corresponds to dimensional TVD discretization and $\sigma = 5$. The legend “1” corresponds to the multidimensional discretization with $\sigma = 1$. The legend “2” corresponds to the multidimensional discretization with $\sigma = 5$. The legend “3” corresponds to the multidimensional discretization with $\sigma = 25$. 192×192 mesh points along with the S_6 LSH ordinate set were used. The minmod limiter was used and the source terms were excluded from the limiter.

dimensional TVD schemes. As pointed out in Section 2.1, this point has important implications for WENO discretizations of the RTE. Legends “2” and “3” in Figs. 10 and 11 show comparable convergence rates. Legend “1” in Fig. 10 converges faster than legend “1” in Fig. 11. This shows that the exclusion of the source terms from the limiter in the genuinely multidimensional schemes can on some occasions help improve the convergence rate. It must be pointed out that the scattering terms couple strongly with the advection terms in the regime of intermediate scattering opacities. Thus it is impossible (i.e. mathematically unjustified) to use operator splitting to treat the scattering terms separately in the relaxation process and yet expect good convergence rates. Unfortunately, previous practitioners had always made this operator splitting, thereby making it very difficult to treat problems with strong scattering terms. In Section 2 we had identified optimal schemes as ones which have a good convergence rate in both the free streaming and optically thick regimes. This includes situations where the high optical depth is purely a consequence of a high scattering opacity. *The results in Figs. 10 and 11 further reinforce our assertion that the FAS multigrid which uses the restarted Newton–Krylov method as a smoother for the genuinely multidimensional discretization of the RTE is such an optimal scheme. The importance of using differentiable limiters is also clearly demonstrated.*

It is also worthwhile to comment on the role that the surface emissivity ε plays in determining the convergence rates. As ε decreases, the surface reflectivity ρ increases in Eq. (2). Thus the value of the radiation intensity at the wall increasingly depends on the scattering terms, which themselves




Fig. 11. (a) and (b) show the residual convergence histories in the L_∞ and L_2 norms for several problems that have a pure scattering opacity. The walls of the domain have an emissivity ε of 0.3. FAS multigrid with Newton-GMRES(15) smoothing was used in all four cases. The legend “0” corresponds to dimensional TVD discretization and $\sigma = 5$. The legend “1” corresponds to the multidimensional discretization with $\sigma = 1$. The legend “2” corresponds to the multidimensional discretization with $\sigma = 5$. The legend “3” corresponds to the multidimensional discretization with $\sigma = 25$. 192×192 mesh points along with the S_6 LSH ordinate set were used. The vanAlbada limiter was used and the source terms were included in the limiter. The difference in the limiters and the inclusion of the source terms in the limiters was the only difference between Figs. 10 and 11.

depend on the interior solution and are not fixed in the same way that $I_b(\mathbf{r})$ is fixed in Eq. (2). Thus as the surface emissivity increases the boundary conditions that are applied in the relaxation of the RTE undergo increasing changes with each iteration. This impedes convergence. Convergence is further impeded if there are two highly reflective walls that are parallel to each other in the problem because the radiation can undergo multiple reflections between the walls before leaving the domain. For that reason, both Fiveland and Jessee [33] and Ramone, Adams and Nowak [34] note that problems with highly reflective walls that are parallel to each other present the maximum difficulty as far as convergence is concerned. The methods presented here also experience difficulty in sustaining rapid convergence rates when the problem has parallel walls with surface emissivity ε much less than 0.3. One way of viewing this problem is to say that the convergence rate is no longer dependent on the convergence rate of the basic algorithm in the interior of the computational domain. Instead, the convergence rate becomes sensitively dependent on the rate at which the solution converges at the boundaries. In that situation, Ramone et al. [34] suggest that one should use some form of relaxation at the boundaries. They, however, do not give any further details. We have made movies of the evolution of the mean intensity and the residual when parallel walls with surface emissivity ε much less than 0.3 are present in the problem. We can see clearly that the radiation intensity undergoes many small fluctuations and the residual oscillates back and forth

between the parallel highly reflective walls for several iterations. This is consistent with the physics of the problem where the radiation indeed undergoes multiple reflections when trapped between highly reflecting parallel walls. We do not analyze this problem further in this paper. It just suffices to mention that for problems with parallel walls having surface emissivity ε greater than or equal to 0.3 the algorithms presented here seem to converge with convergence rates that are consistent with the interior convergence rates.

We now turn our attention to demonstrating the order of accuracy of the genuinely multidimensional schemes presented here. Jessee and Fiveland [28] have presented an accuracy analysis of discrete ordinate methods for the RTE when dimensional TVD discretizations and several different choices of limiters were used. We focus here on the genuinely multidimensional discretizations of the RTE. Because the role of the limiters on the accuracy of the multidimensional schemes has not been catalogued in the literature, we catalogue a few interesting ones here. The accuracy analysis presented here uses the S_6 ordinate set. The unit square is covered with a range of meshes having as few as 24×24 zones to as many as 192×192 zones. Large calculations that used the S_8 ordinate set on 256×256 zone meshes were used as reference calculations. Several accuracy analyses were carried out where the scattering opacity assumed values of $\sigma = 5$ and 25, respectively. In these problems we set $\kappa = 0$; $I_b(\mathbf{r}) = 0$ in Eq. (1). The walls of the square had a surface emissivity of $\varepsilon = 0.3$. The upper and lower y -boundaries as well as the upper x -boundary were held at zero degrees Kelvin. The lower x -boundary had $I_b(\mathbf{r}) = 1/\pi$. The accuracy analyses are tabulated in Tables 1–6. Table 1 corresponds to $\sigma = 5$ when the minmod limiter was used and the source terms were excluded from the limiter. Table 2 corresponds to $\sigma = 5$ when the minmod limiter was used and the source terms were included in the limiter. Table 3 corresponds to $\sigma = 5$ when the vanAlbada limiter was used and the source terms were included in the limiter. Table 4 corresponds to $\sigma = 25$ when the minmod limiter was used and the source terms were excluded from the limiter. Table 5 corresponds to $\sigma = 25$ when the minmod limiter was used and the source terms were included in the limiter. Table 6 corresponds to $\sigma = 25$ when the vanAlbada limiter was used and the source terms were included in the limiter. Tables 1–3 correspond to the simulation shown in Fig. 9a. Tables 4–6 correspond to the simulation shown in Fig. 9b. We are interested in the error in the angle-averaged mean intensity evaluated in the L_∞ and L_2 norms. In each case we show the error as a percentage of the mean intensity averaged over the entire computational domain. The order of accuracy of the solution is also evaluated in each of those two norms. We see that the methods do not display their higher-order accuracy on meshes with poor resolutions. However, as

Table 1
Accuracy analysis for $\sigma = 5$; $\varepsilon = 0.3$. Source terms excluded from minmod limiter

No. of zones	% error, L_∞ norm	Order, L_∞ norm	% error, L_2 norm	Order, L_2 norm
24×24	11.56		8.882	
48×48	5.953	0.957	4.570	0.959
96×96	2.355	1.338	1.773	1.366
144×144	1.360	1.353	0.726	2.201
192×192	0.862	1.585	0.441	1.733

Table 2
Accuracy analysis for $\sigma = 5$; $\varepsilon = 0.3$. Source terms included in minmod limiter

No. of zones	% error, L_∞ norm	Order, L_∞ norm	% error, L_2 norm	Order, L_2 norm
24 × 24	9.371		5.888	
48 × 48	5.338	0.812	3.464	0.765
96 × 96	2.393	1.158	1.552	1.159
144 × 144	1.423	1.280	0.765	1.742
192 × 192	0.988	1.268	0.339	2.827

Table 3
Accuracy analysis for $\sigma = 5$; $\varepsilon = 0.3$. Source terms included in vanAlbada limiter

No. of zones	% error, L_∞ norm	Order, L_∞ norm	% error, L_2 norm	Order, L_2 norm
24 × 24	8.716		5.309	
48 × 48	4.881	0.837	3.007	0.820
96 × 96	2.241	1.123	1.333	1.174
144 × 144	1.374	1.206	6.579	1.742
192 × 192	0.970	1.209	0.292	2.818

Table 4
Accuracy analysis for $\sigma = 25$; $\varepsilon = 0.3$. Source terms excluded from minmod limiter

No. of zones	% error, L_∞ norm	Order, L_∞ norm	% error, L_2 norm	Order, L_2 norm
24 × 24	24.47		16.38	
48 × 48	13.92	0.814	9.830	0.737
96 × 96	6.213	1.164	4.605	1.094
144 × 144	3.075	1.735	2.308	1.703
192 × 192	1.536	2.410	1.096	2.588

the mesh resolution improves, the higher-order nature of the genuinely multidimensional schemes becomes apparent. It is also significant that the order of accuracy (as measured in the L_2 norm) of the genuinely multidimensional schemes on meshes with sufficient resolution seems to saturate to values that are rather close to the theoretically predicted second-order accuracy. This is very significant because Jessee and Fiveland [28] have shown that the order of accuracy of schemes that use the dimensional TVD discretizations saturates to values that are far short of their predicted second-order accuracy. We point out that the test problems used in our accuracy analysis as well as in the accuracy analysis of Jessee and Fiveland [28] do not have sharp discontinuities in them. As a result, one may justifiably expect them to display accuracies that are close to second-order accuracy. We also point out that the problem studied in Tables 4–6 requires larger meshes to reach

Table 5

Accuracy analysis for $\sigma = 25$; $\varepsilon = 0.3$. Source terms included in minmod limiter

No. of zones	% error, L_∞ norm	Order, L_∞ norm	% error, L_2 norm	Order, L_2 norm
24 × 24	22.48		12.91	
48 × 48	11.32	0.989	7.662	0.752
96 × 96	5.301	1.095	3.720	1.042
144 × 144	2.675	1.687	1.927	1.622
192 × 192	1.271	2.586	0.877	2.734

Table 6

Accuracy analysis for $\sigma = 25$; $\varepsilon = 0.3$. Source terms included in vanAlbada limiter

No. of zones	% error, L_∞ norm	Order, L_∞ norm	% error, L_2 norm	Order, L_2 norm
24 × 24	20.31		13.83	
48 × 48	12.17	0.739	8.324	0.733
96 × 96	5.585	1.124	3.981	1.064
144 × 144	2.787	1.714	2.036	1.654
192 × 192	1.374	2.458	0.919	2.764

its design accuracy than the problem studied in Tables 1–3. The source of this difference becomes apparent when one compares Fig. 9b to a. We see clearly that Fig. 9b has steeper gradients in its solution. This explains why the problem studied in Tables 4–6 requires larger meshes to reach its design accuracy. From Tables 1–3 we see that the order of accuracy as measured in the L_∞ norm does not track the order of accuracy as measured in the L_2 norm. This is known to happen even in CFD problems when second-order schemes are used. It happens because a few extremal points can reduce the order of accuracy as measured in the L_∞ norm. From Tables 4–6 we see that the order of accuracy as measured in the L_∞ norm tracks the order of accuracy as measured in the L_2 norm. It is also significant that the exclusion of the source terms from the minmod limiter in Tables 1 and 4 does not cause a significant degradation in accuracy. The present accuracy analysis also makes it possible to inter-compare the role of the minmod limiter in Tables 2 and 5 with the vanAlbada limiter in Tables 3 and 6. We see that the minmod and vanAlbada limiters produce results with comparable accuracies. We have also carried out the same accuracy analysis with the vanLeer limiter and found that it too produces results that are as accurate as those produced by the minmod and vanAlbada limiters. The fact that the accuracies are comparable prompts us to conclude that the specific choice of limiter plays a small role in determining the order of accuracy of the genuinely multidimensional schemes. Thus we conclude that the genuinely multidimensional discretizations of the RTE outperform the dimensional TVD discretizations of the RTE not because of the specific choice of limiter but rather because of the superior discretization that closely emulates the physics of the problem.

The fact that the genuinely multidimensional schemes show accuracies that are close to second-order has two further implications:

(1) The requirement of Mihalas and Klein [42] that the scheme should be second-order accurate in order to correctly retrieve the optically thick limit is more closely met by the present genuinely multidimensional schemes than it is met by the dimensional TVD schemes.

(2) Jessee et al. [30] used rather low-quality adaptation criteria for their adaptive solution of the RTE. Their adaptation criteria were, however, well-matched with the fact that their scheme was only first-order accurate. The availability of second-order schemes for the solution of the RTE implies that better adaptation criteria from Berger and Olinger [73] or Zienkiewicz and Zhu [74] can now be used for the adaptive solution of the RTE.

6. Conclusions

In this paper we have designed several useful strategies for numerically solving the RTE using discrete ordinates methods. We have also analyzed the reasons for the poor convergence and accuracy of previous methods and demonstrated that our methods circumvent many of the pitfalls that plagued previous methods. Based on the work presented here we draw the following conclusions:

(1) It is possible to discretize the RTE using dimensional TVD discretizations as well as genuinely multidimensional discretizations.

(2) The genuinely multidimensional discretizations have several advantages over the dimensional TVD discretizations. Because the former take account of the direction of characteristic propagation, they represent the physics with greater fidelity.

(3) The genuinely multidimensional discretizations also have minimal cross-stream dissipation and dispersion errors making them very useful in accurately representing the free streaming limit of the RTE.

(4) Furthermore, using the genuinely multidimensional discretizations results in a compact stencil. This endows the scheme with superior convergence properties.

(5) Newton–Krylov methods are shown to provide us with a natural way of treating the advection and scattering terms in the RTE without resort to operator splitting. This permits maximal coupling between the terms. It is demonstrated that this coupling enhances convergence even in the strong scattering limit where previous methods have been known to perform poorly.

(6) The convergence of the Newton–Krylov methods has been shown not to degrade too much even when zones with very bad aspect ratios are present. This enhances their utility in general purpose codes.

(7) When the Newton–Krylov method is used as a smoother within FAS multigrid a strategy for numerical radiative transfer with superior convergence is seen to result. This is especially true when the genuinely multidimensional discretizations of the RTE are used.

(8) The resulting scheme is optimal in the sense that it has good convergence properties in the free streaming limit as well as in the limit where the optical depth in each zone is very large. This is true even when a strong scattering opacity is present in the problem.

(9) A variety of illustrative model problems and stringent test problems have been presented. Several of the test problems involved pure scattering with strong scattering terms. The methods presented here have been shown to perform very well on those problems. Based on our test problems we made the very useful and interesting discovery that increasing the strength of the scattering terms actually results in an improvement of the convergence rates for the methods presented here.

(10) Accuracy analysis for the genuinely multidimensional discretizations of the RTE shows that they are close to second-order accurate. This stands in marked contrast to the dimensional TVD discretizations of the RTE which fall short of their designed second-order accuracy.

(11) Different limiters are used in this work and an inter-comparison is made of their role in producing good convergence and accuracy. It is shown that the use of differentiable limiters in the genuinely multidimensional schemes produces better convergence rates when used in conjunction with the Newton–Krylov method. The accuracy that one obtains seems to be close to the designed second-order accuracy of the multidimensional schemes and does not depend on the specific choice of limiter.

(12) It is shown that the exclusion of the source terms from the limiters in the multidimensional schemes does not cause a significant degradation in the accuracy of the solution in all the cases that we have tested. This point is important because when the source terms are excluded from the limiters the genuinely multidimensional discretizations of the RTE also display slightly faster convergence in some instances.

Appendix A

One problem that we faced in an early phase of this work was the unavailability of accurate values for the discrete ordinates and their weights. Values of the discrete ordinates that retain seven or less digits of accuracy have been published. However, due to the availability of 64-bit computers, we wanted values of the ordinates that were at least accurate up to 14 or more digits. The theory for constructing these ordinates has been documented by Fiveland [43] and references therein. Fiveland [43] also showed that by preserving a mix of even and odd moments one could derive level symmetric hybrid (LSH) ordinates which were in some cases much superior to the level symmetric even (LSE) ordinates that were proposed previously by Lathrop and Carlson [17]. For that reason, we catalogue LSH ordinate sets below whenever they are available. The S_8 LSH ordinate set that we catalogue below is a little different from Fiveland's choice but yields superior preservation of moments. It satisfies the first eight moment equations to within 0.47% accuracy and the first 15 moment equations to within 7.0% accuracy. The S_{10} LSH ordinate set given here is based on Fiveland's choice. It satisfies the first 15 moment equations to within 0.37% accuracy and the first 20 moment equations to within 2.4% accuracy. For S_{12} we catalogue a newly designed LSH ordinate set that yields positive weights and satisfies the first 15 moment equations to within 0.15% accuracy! It also satisfies the first 20 moment equations to within 1.1% accuracy. In all cases, the ordinate sets given below satisfy the required moment equations to fourteen digits of accuracy. Only the ordinates in the first octant are given. The remaining ordinates can be obtained by using symmetry arguments. Since the ordinates are level symmetric, each ordinate corresponds to a triplet of indices which add up to " $2 + n/2$ ". These indices may be useful to some researchers. As a result, we give the indices that correspond to the ordinates below each table of ordinates.

S₂ LSE

$\Omega_{x,m}$	$\Omega_{y,m}$	$\Omega_{z,m}$	w_m
0.5773502691896257	0.5773502691896257	0.5773502691896257	1.5707963267948966

The above ordinate corresponds to indices (1,1,1).

S₄ LSE

$\Omega_{x,m}$	$\Omega_{y,m}$	$\Omega_{z,m}$	w_m
0.8688903007222013	0.3500211745815406	0.3500211745815406	0.52359877559829893
0.3500211745815406	0.8688903007222013	0.3500211745815406	0.52359877559829893
0.3500211745815406	0.3500211745815406	0.8688903007222013	0.52359877559829893

The above ordinates correspond to indices (2,1,1), (1,2,1) and (1,1,2), respectively.

S₆ LSH

$\Omega_{x,m}$	$\Omega_{y,m}$	$\Omega_{z,m}$	w_m
0.9626351140347762	0.1914858183094162	0.1914858183094162	0.1780146904741121
0.6940220390543709	0.6940220390543709	0.1914858183094162	0.3455840851241866
0.6940220390543709	0.1914858183094162	0.6940220390543709	0.3455840851241866
0.1914858183094162	0.9626351140347762	0.1914858183094162	0.1780146904741121
0.1914858183094162	0.6940220390543709	0.6940220390543709	0.3455840851241866
0.1914858183094162	0.1914858183094162	0.9626351140347762	0.1780146904741121

The above ordinates correspond to indices (3,1,1), (2,2,1), (2,1,2), (1,3,1), (1,2,2) and (1,1,3), respectively. It can be derived by satisfying the moment equations for moment numbers 0,1,2 and 4.

S₈ LSH

$\Omega_{x,m}$	$\Omega_{y,m}$	$\Omega_{z,m}$	w_m
0.9709745907812476	0.1691276796641945	0.1691276796641945	0.1461389389986996
0.7987881412728111	0.5773502691896257	0.1691276796641945	0.1598388991393510
0.7987881412728111	0.1691276796641945	0.5773502691896257	0.1598388991393510
0.5773502691896257	0.7987881412728111	0.1691276796641945	0.1598388991393510
0.5773502691896257	0.5773502691896257	0.5773502691896257	0.1733461149626915
0.5773502691896257	0.1691276796641945	0.7987881412728111	0.1598388991393510
0.1691276796641945	0.9709745907812476	0.1691276796641945	0.1461389389986996
0.1691276796641945	0.7987881412728111	0.5773502691896257	0.1598388991393510
0.1691276796641945	0.5773502691896257	0.7987881412728111	0.1598388991393510
0.1691276796641945	0.1691276796641945	0.9709745907812476	0.1461389389986996

The above ordinates correspond to indices (4,1,1), (3,2,1), (3,1,2), (2,3,1), (2,2,2), (2,1,3), (1,4,1), (1,3,2), (1,2,3) and (1,1,4), respectively. It can be derived by satisfying the moment equations for moment numbers 0,1,2,5 and 8.

S₁₀ LSH

$\Omega_{x,m}$	$\Omega_{y,m}$	$\Omega_{z,m}$	w_m
0.9809754496167162	0.1372719331277908	0.1372719331277908	0.0944411600027882
0.8523177344565812	0.5046889100289010	0.1372719331277908	0.1483951164707604
0.8523177344565812	0.1372719331277908	0.5046889100289010	0.1483951164707604
0.7004128840817248	0.7004128840817248	0.1372719331277908	0.0173702170292469
0.7004128840817248	0.5046889100289010	0.5046889100289010	0.1149971656247430
0.7004128840817248	0.1372719331277908	0.7004128840817248	0.0173702170292469
0.5046889100289010	0.8523177344565812	0.1372719331277908	0.1483951164707604
0.5046889100289010	0.7004128840817248	0.5046889100289010	0.1149971656247430
0.5046889100289010	0.5046889100289010	0.7004128840817248	0.1149971656247430
0.5046889100289010	0.1372719331277908	0.8523177344565812	0.1483951164707604
0.1372719331277908	0.9809754496167162	0.1372719331277908	0.0944411600027882
0.1372719331277908	0.8523177344565812	0.5046889100289010	0.1483951164707604
0.1372719331277908	0.7004128840817248	0.7004128840817248	0.0173702170292469
0.1372719331277908	0.5046889100289010	0.8523177344565812	0.1483951164707604
0.1372719331277908	0.1372719331277908	0.9809754496167162	0.0944411600027882

The above ordinates correspond to indices (5,1,1), (4,2,1), (4,1,2), (3,3,1), (3,2,2), (3,1,3), (2,4,1), (2,3,2), (2,2,3), (2,1,4), (1,5,1), (1,4,2), (1,3,3), (1,2,4) and (1,1,5), respectively. It can be derived by satisfying the moment equations for moment numbers 0,1,2,5,8 and 10.

S₁₂ LSH

$\Omega_{x,m}$	$\Omega_{y,m}$	$\Omega_{z,m}$	w_m
0.9834365126237636	0.1281651778729502	0.1281651778729502	0.0802616917313290
0.8814778284523741	0.4545003026706165	0.1281651778729502	0.1082299400202983
0.8814778284523741	0.1281651778729502	0.4545003026706165	0.1082299400202983
0.7660671966248366	0.6298529490572852	0.1281651778729502	0.0451194191848942
0.7660671966248366	0.4545003026706165	0.4545003026706165	0.0713859132472878
0.7660671966248366	0.1281651778729502	0.6298529490572852	0.0451194191848942
0.6298529490572852	0.7660671966248366	0.1281651778729502	0.0451194191848942
0.6298529490572852	0.6298529490572852	0.4545003026706165	0.0652524522092968
0.6298529490572852	0.4545003026706165	0.6298529490572852	0.0652524522092968
0.6298529490572852	0.1281651778729502	0.7660671966248366	0.0451194191848942
0.4545003026706165	0.8814778284523741	0.1281651778729502	0.1082299400202983
0.4545003026706165	0.7660671966248366	0.4545003026706165	0.0713859132472878
0.4545003026706165	0.6298529490572852	0.6298529490572852	0.0652524522092968
0.4545003026706165	0.4545003026706165	0.7660671966248366	0.0713859132472878
0.4545003026706165	0.1281651778729502	0.8814778284523741	0.1082299400202983
0.1281651778729502	0.9834365126237636	0.1281651778729502	0.0802616917313290
0.1281651778729502	0.8814778284523741	0.4545003026706165	0.1082299400202983
0.1281651778729502	0.7660671966248366	0.6298529490572852	0.0451194191848942
0.1281651778729502	0.6298529490572852	0.7660671966248366	0.0451194191848942
0.1281651778729502	0.4545003026706165	0.8814778284523741	0.1082299400202983
0.1281651778729502	0.1281651778729502	0.9834365126237636	0.0802616917313290

The above ordinates correspond to indices (6,1,1), (5,2,1), (5,1,2), (4,3,1), (4,2,2), (4,1,3), (3,4,1), (3,3,2), (3,2,3), (3,1,4), (2,5,1), (2,4,2), (2,3,3), (2,2,4), (2,1,5), (1,6,1), (1,5,2), (1,4,3), (1,3,4), (1,2,5) and

(1,1,6), respectively. It can be derived by satisfying the moment equations for moment numbers 0,1,2,5,8,10 and 13.

References

- [1] Castor JI. Radiative transfer in spherically symmetric flows. *Astrophys J* 1972;178:779.
- [2] Buchler JR. Radiation hydrodynamics in the fluid frame. *JQSRT* 1979;22:293.
- [3] Buchler JR. Radiation transfer in the fluid frame. *JQSRT* 1982;30:395.
- [4] Balsara D. The Eigenstructure of the equations of radiation hydrodynamics. *JQSRT* 1999;61(5):637.
- [5] Balsara D. Linearized formulation of the Riemann problem for radiation hydrodynamics. *JQSRT* 1999;61(5):629.
- [6] Balsara D. Exact Jacobians of Roe-type flux difference splitting of the equations of radiation hydrodynamics (and Euler equations) for use in time-implicit higher order Godunov schemes. *JQSRT* 1999;62:255.
- [7] Balsara D. An analysis of the hyperbolic nature of the equations of radiation magnetohydrodynamics. *JQSRT* 1999;62(5):617.
- [8] Balsara D. Linearized formulation of the Riemann problem for radiation magnetohydrodynamics. *JQSRT* 1999;62:167–89.
- [9] Krook M. On the solutions of the equations of transfer I. *Astrophys J* 1955;488–95.
- [10] Cheng P. Two dimensional radiating gas flow by a moment method. *AIAA J* 1964;2:1662–4.
- [11] Gelbard EM. Simplified spherical harmonics equations and their use in shielding problems. WAPD-T-1182 (Rev. 1), February 1961.
- [12] Liu YH, Gelbard EM. Accuracy of Nodal transport and simplified P_3 fluxes in benchmark tests. *Trans Am Nucl Soc* 1986;52:430.
- [13] Morel JE. The SP_n equations for radiative transfer with material motion corrections. Proceedings of the Joint International Conference on Mathematical Methods and Supercomputing in Nuclear Applications, vol. 1, Kongresszentrum, Karlsruhe, Germany, April 1993. p. 718.
- [14] Larsen EW, Morel JE, McGhee JM. Asymptotic derivation of the simplified P_n equations. Proceedings of the Joint International Conference on Mathematical Methods and Supercomputing in Nuclear Applications, vol. 1, Kongresszentrum, Karlsruhe, Germany, April 1993. p. 718.
- [15] Ratzell AC, Howell JR. Two dimensional radiation in absorbing-emitting media using the P-N approximation. *J Heat Transfer* 1983;105:333.
- [16] Chandrasekhar S. Radiative transfer. Oxford: Clarendon Press, 1950.
- [17] Lathrop KD, Carlson BG. Discrete ordinates angular quadrature of the neutron transport equation. Los Alamos Scientific Laboratory Report 3186, 1965.
- [18] Larsen EW, Morel JE. *J Comput Phys* 1989;83:212.
- [19] Morel JE, Wareing TA, Smith K. A linear-discontinuous spatial differencing scheme for S_n radiative transfer calculations. *J Comput Phys* 1996;128:445.
- [20] Adams ML. Subcell balance methods for radiative transfer on arbitrary grids. *Transp Theory Stat Phys* 1997;26:385.
- [21] Adams ML, Novak PF. Asymptotic analysis of a computational method for time and frequency-dependent radiative transfer. *J Comput Phys* 1998;146:366.
- [22] Fiveland WA, Jessee JP. Finite element formulation of the discrete ordinates method in multidimensional geometries. *J Thermophys Heat Transfer* 1994;8(3):426.
- [23] Fiveland WA, Jessee JP. Comparison of discrete ordinate formulations for radiative heat transfer in multidimensional geometries. *J Thermophys Heat Transfer* 1995;9(1):47.
- [24] Fiveland WA. Discrete ordinates solutions of the radiative transport equation for rectangular enclosures. *Trans ASME, J Heat Transfer* 1984;106:699–706.
- [25] Haferman JL, Smith TF, Krajewski WF. ASME, HTD-vol. 325, National Heat Transfer Conference, 1996. p. 125–32.
- [26] VanLeer B. Towards the ultimate conservation difference scheme V, A second order sequel to Godunov's method. *J Comput Phys* 1979;23:101–36.

- [27] Harten A. High resolution schemes for hyperbolic conservation laws. *J Comput Phys* 1983;49(3):357–93.
- [28] Jessee JP, Fiveland WA. Bounded, high resolution differencing schemes applied to the discrete ordinates method. *ASME, HTD-vol. 325, Proceedings of the 31st National Heat Transfer Conference, vol. 3, 1996. p. 133.*
- [29] Fiveland WA. Three-dimensional radiative heat transfer solutions by the discrete ordinates method. *J Thermophys Heat Transfer* 1988;2(4):309–16.
- [30] Jessee JP, Fiveland WA, Howell LH, Colella P, Pember RB. An adaptive mesh refinement algorithm for the radiative transport equation. *J Comput Phys* 1998;139:380–98.
- [31] Howell LH, Pember RB, Colella P, Jessee JP, Fiveland WA. A conservative adaptive mesh algorithm for unsteady, combined mode heat transfer using the discrete ordinates method. *Numer Heat Transfer, Part B* 1999;35:407–30.
- [32] Balsara D, Norton C. Highly parallel structured adaptive mesh refinement using parallel language-based approaches. *J Parallel Comput* (2000; to appear).
- [33] Fiveland WA, Jessee JP. Acceleration schemes for the discrete ordinates method. *ASME, HTD-vol. 315, Proceedings of the 30th National Heat Transfer Conference, vol. 13, 1995. p. 445.*
- [34] Ramone GL, Adams ML, Nowak PF. A transport synthetic acceleration method for transport iterations. *Nucl Sci Eng* 1997;125:257.
- [35] Larsen EW. Unconditionally stable diffusion-synthetic acceleration methods for the slab geometry discrete ordinates equations, Part I: Theory. *Nucl Sci Eng* 1982;82:47.
- [36] Burns SP, Christon MA. Spatial domain-based parallelism in large scale, participating-media, radiative transport applications. Sandia Report SAND96-2485.
- [37] Morel JE, Larsen EW, Matzen MK. A synthetic acceleration scheme for radiative diffusion calculations. *JQSRT* 1985;34:243.
- [38] Sidilkover D, Roe PL. Unification of some advection schemes in two dimensions. *Math Comput* (2000), submitted for publication.
- [39] Saad Y, Schultz MH. GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J Sci Stat Comput* 1986;7(1):860.
- [40] Brown PN, Saad Y. Hybrid Krylov methods for nonlinear systems of equations. *SIAM J Sci Stat Comput* 1990;11(3):450.
- [41] Brandt A. Guide to multigrid development.. In: Hackbusch W, Trottenberg U, editors. *Multigrid methods. Lecture Notes in Mathematics, vol. 960. Berlin: Springer, 1982.*
- [42] Mihalas D, Klein RI. On the solution of the time-dependent inertial-frame equation of radiative transfer in moving media to $O(v/c)$. *J Comput Phys* 1982;46:97.
- [43] Fiveland WA. The selection of discrete ordinate quadrature sets for anisotropic scattering. *ASME, HTD-vol. 160, Fundamentals of radiation heat transfer, 1991. p. 89.*
- [44] Yavneh I. A method for devising efficient multigrid smoothers for complicated PDE systems. *SIAM J Sci Comput* 1993;14:1437–63.
- [45] Sidilkover D, Brandt A. Multigrid solution to steady-state 2d conservation laws. *SIAM J Numer Anal* 1993;30:249–74.
- [46] Sweby P. High resolution schemes using flux limiters for hyperbolic conservation laws. *SIAM J Numer Anal* 1984;21:995–1011.
- [47] Chakravarthy SR. AIAA Paper #84-0165, 1984.
- [48] Harten A. *SIAM J Numer Anal* 1984;21:1.
- [49] Mulder WA, vanLeer B. *J Comput Phys* 1985;59:232.
- [50] Gropp WD, Keyes DE, McInnes LC, Tidiri MD. Globalized Newton-Krylov-Schwarz algorithms and software for parallel implicit CFD. ICASE Report No. 98-24, 1998.
- [51] Liu X-D, Osher S, Chan T. Weighted essentially nonoscillatory schemes. *J Comput Phys* 1994;115:200–12.
- [52] Jiang G-S, Shu C-W. Efficient implementation of weighted ENO schemes. *J Comput Phys* 1996;126:202–28.
- [53] Balsara D, Shu C-W. Monotonicity preserving weighted essentially non-oscillatory schemes with increasingly high order of accuracy. *J Comput Phys* 2000, to appear.
- [54] Goodman JB, LeVeque RJ. On the accuracy of stable schemes for two dimensional conservation laws. *Math Comp* 1985;45:15–21.
- [55] Colella P. Multidimensional upwind methods for hyperbolic conservation laws. *J Comput Phys* 1990;87:171.

- [56] Hirsch Ch, van Ransbeek P. Cell centered multidimensional upwind algorithm and structured meshes. Proceedings of ECOMAS 1st European CFD Conference, vol. 1. Amsterdam: Elsevier, 1992. p. 53–60.
- [57] LeVeque RJ. High resolution finite volume methods on arbitrary grids via wave propagation. *J Comput Phys* 1988;78:36.
- [58] Roe PL. Optimum upwind advection on a triangular mesh. ICASE Report #90-75, 1990.
- [59] Roe PL, Sidilkover D. Optimum positive linear schemes for advection in two and three dimensions. *SIAM J Numer Anal* 1992;29:1542–68.
- [60] Sidilkover D. A genuinely multidimensional upwind scheme and an efficient multigrid for compressible Euler equations. ICASE Report 94-84, 1994.
- [61] Harten A. The artificial compression method for computation of shocks and contact discontinuities: III Self-adjusting hybrid schemes. *Math Comp* 1978;32:363–89.
- [62] Raithby GD, Chui EH. A finite-volume method for predicting radiant heat transfer in enclosures with participating media. *ASME Trans J Heat Transfer* 1990;112:415.
- [63] Raithby GD. Skew upwind difference schemes for problems involving fluid flow. *Comput Methods Appl Mech Eng* 1976;9:153–64.
- [64] Brown PN. A local convergence theory for combined inexact-Newton/finite-difference projection methods. *SIAM J Sci Stat Comput* 1987;24(2):407.
- [65] Walker HF. Implementation of the GMRES method using householder transformations. *SIAM J Sci Stat Comput* 1988;9(1):152.
- [66] Saad Y. Iterative methods for sparse linear systems. Kenf: PWS Publishing Company, a division of International Thomson Publishing Inc., 1996.
- [67] Parlett BN. The symmetric Eigenvalue problem. Englewood Cliffs: Prentice-Hall, 1980.
- [68] Saad Y. A flexible inner-outer preconditioned GMRES algorithm. *SIAM J Sci Stat Comput* 1993;14(2):461.
- [69] Kaczmarz S. Angenaherte Auflosung von Systemen Linearer Gleichungen. *Bull Int Acad Polon Sci Lett* 1939;6:355–7.
- [70] Brandt A, McCormick S, Ruge J. Algebraic multigrid (AMG) for automatic multigrid solution with applications to geodetic computations. Institute for Computational Studies, POB 1852, Fort Collins, Colorado, 1982.
- [71] Brandt A. General highly accurate algebraic coarsening schemes. preprint received, 1999.
- [72] Morel JE, Manteuffel TA. An angular multigrid acceleration technique for S_n equations with highly forward-peaked scattering. *Nucl Sci Eng* 1991;107:330–42.
- [73] Berger M, Oliger J. Adaptive mesh refinement for hyperbolic partial differential equations. *J Comput Phys* 1984;53:64–84.
- [74] Zienkiewicz OC, Zhu JZ. A simple error estimator and adaptive procedure for practical engineering analysis. *Int J Numer Methods Eng* 1987;24:337–57.