Available online at www.sciencedirect.com
ScienceDirect
Journal of Computational Physics 226 (2007) 1181–1195
JOURNAL OF COMPUTATIONAL PHYSICS
www.elsevier.com/locate/jcp

Efficient solution of multi-dimensional flux-limited nonequilibrium radiation diffusion coupled to material conduction with second-order time discretization

Gordon L. Olson *

Computer and Computational Sciences Division (CCS-2), Los Alamos National Laboratory, 5 Foxglove Circle, Madison, WI 53717, USA

Received 2 January 2007; received in revised form 22 April 2007; accepted 14 May 2007
Available online 26 May 2007

Abstract

Many algorithms for the second-order time evolution of the coupled radiation diffusion and material conduction equations have been published. Most of them are cumbersome to implement and much slower computationally than their first-order equivalent algorithms. This paper presents a simpler approach that is both computationally efficient and easy to implement. Second-order behavior can be achieved even when the iteration at each time step is incompletely converged. The test problem uses multiple materials and nonlinear heat capacities. Unexpectedly, details in the discretization of the gradient in the flux limiter significantly affect the spatial and temporal convergence of the solution.
Published by Elsevier Inc.

Keywords: Radiation diffusion; Flux-limited diffusion; Crank–Nicolson; Second-order

1. Introduction

Many authors [1–8] have studied the numerical solution of the nonequilibrium radiation diffusion equation coupled to a material energy balance equation. Some authors include the conduction process in the material equation so that there are two coupled diffusion-reaction type equations. These equations are very nonlinear and tightly coupled; therefore, they offer an interesting challenge for their accurate solution. In order to obtain accurate solutions, it has been recognized that second-order time evolution is highly desirable. However, the equations are complex enough that each author has found a slightly different solution technique through different linearizations and operator splittings. Most researchers have used nonlinear Newton or Newton–Krylov solution techniques for these equations.

This paper takes a different approach. We will see how far we can push simple iteration schemes that are easier to code into a computer program and thus are also easier to verify for correctness. A potential problem with Newton-type solution methods is that they can fail if the initial guess is not close enough to the converged

* Tel.: +1 608 836 1779.
E-mail address: olson_g@tds.net

0021-9991/$ - see front matter Published by Elsevier Inc.
doi:10.1016/j.jcp.2007.05.015

solution. Simple iteration schemes can be more robust, but they are useful only if they converge fast enough. Of course, it is necessary to use a difficult two-dimensional test problem that will stress the numerical methods. We extend the multimaterial problem in [3,8] to also use a Saha ionization model [2,7] that makes the heat capacity nonlinear.

The following section will present the equations as they are solved here. Then Section 3 will show a second-order time discretization and iteration scheme. Section 4 discusses an unexpected sensitivity of the numerical results to details in the flux limiter that is applied to radiation diffusion. The main two-dimensional test results are given in Section 5. A comparison to a traditional Crank–Nicolson discretization method is presented in Section 6. A brief summary is in the last section.

2. Nonequilibrium diffusion equations

The equations for nonequilibrium flux-limited diffusion coupled to material conduction have been written using different notation by various authors (see [2,9,11]). We choose the following form for the two coupled partial differential equations:

$$\frac{1}{c}\frac{\partial E}{\partial t} - \nabla\cdot(D\nabla E) = \sigma_{\mathrm{a}}(T^4 - E), \tag{1}$$

$$\frac{c_V}{c}\frac{\partial T}{\partial t} - \nabla\cdot(K\nabla T) = \sigma_{\mathrm{a}}(E - T^4). \tag{2}$$

Here, E is the radiation energy density, t is time, c is the speed of light, D is the radiation diffusion coefficient, σ_{a} is the absorption opacity which must be evaluated as a Planck mean (see [11,12]) in a multi-frequency calculation, T is the material temperature, c_V is the heat capacity, and K is the material conduction coefficient. Here, we are using units such that $E = T_{\mathrm{r}}^4$, where T_{r} is the radiation temperature.

The form of the first term in Eq. (2) is the one that is usually chosen when the heat capacity is assumed to be constant, independent of temperature or time. Here, we will allow c_V to be variable and it will be evaluated at the old and new time step conditions just like the other material coefficients. By using this form, one can solve the equation as a linear system with nonlinear coefficients. If one uses an energy conservation form, the equation is less easily solved as a linear system. In Section 5, we will show that the error in energy conservation made by using this form is very small. Therefore, this is a useful approximation.

In order to keep the propagation velocity of a radiation wave front in a vacuum less than the speed of light, one must use a flux-limited diffusion coefficient. We use the form suggested by [9]:

$$D = \frac{1}{\sqrt[m]{(3\sigma_{\mathrm{t}})^m + \chi^m}}, \tag{3}$$

where

$$\chi = \frac{|\nabla E|}{E} \tag{4}$$

is the normalized energy gradient and σ_{t} is the total opacity including absorption and scattering (σ_{s}), which must be a Rosseland mean (see [11,12]) in a multi-frequency calculation. Traditional radiation flux limiters usually use $m = 1$ in Eq. (3). However, Morel [10] has shown that such a choice can distort the solution away from the correct transport solution. The traditional sum flux limiter ($m = 1$) can slow wave front propagation too much (see [9]). Therefore, in this work, $m = 2$ will be used in all the 2D test problems. The linear case will only be used for simple 1D examples in the Section 4.

The material properties must be specified through the coefficients in the above equations. Just one test problem will be examined in great detail in this work. We choose the following:

$$\sigma_{\mathrm{a}} = \sigma_{\mathrm{t}} = \frac{Z^3}{T^3}, \quad \sigma_{\mathrm{s}} = 0, \tag{5}$$

$$K = 0.01T^{2.5}, \tag{6}$$

where Z represents the material dependence. The two materials in the two-dimensional test problem have $Z = 1$ and 10. The conduction coefficient is the same for the two materials.

Most studies of radiation diffusion in the literature have assumed that the heat capacity in Eq. (2) is either constant, $c_V = 1$, or proportional to temperature cubed, $c_V = 4T^3$. The first assumption is quit simple, while the second allows for a transformation of variables to $B \equiv T^4$ that turns Eq. (2) into a linear equation for B when material conduction is ignored. Having a linear equation allows for some analytical analysis; however, this temperature dependence has no relation to an actual material. Recently, researchers [2,7] have introduced a scaled hydrogen-like Saha ionization model for a simplified, but physically plausible, heat capacity:

$$c_V = 1 + \alpha + (T + 0.3)\frac{\partial \alpha}{\partial T}, \tag{7}$$

$$\frac{\alpha^2}{1-\alpha} = \mathrm{e}^{-0.3/T} \quad \text{or} \quad \alpha = \frac{1}{2}\mathrm{e}^{-0.3/T}\left(\sqrt{1 + 4\mathrm{e}^{0.3/T}} - 1\right), \tag{8}$$

$$\frac{\partial \alpha}{\partial T} = \frac{0.3}{T^2}\left(\alpha - \frac{1}{\sqrt{1 + 4\mathrm{e}^{0.3/T}}}\right). \tag{9}$$

For a normalized problem where the material temperatures are of order unity, this heat capacity provides a strong nonlinearity for testing numerical methods. When using large time steps, this nonlinearity causes some numerical methods to fail [7].

3. Time discretization

Using the traditional θ weighting of the time derivatives, the radiation energy density can be advanced in time using

$$E^{n+1} = E^n + \theta \Delta t \left.\frac{\partial E}{\partial t}\right|^{n+1} + (1-\theta)\Delta t \left.\frac{\partial E}{\partial t}\right|^{n}, \tag{10}$$

where Δt is the time step, $\Delta t = t^{n+1} - t^n$, and n indicates the time level. For $\theta = 1$, this reduces to the first order backward Euler (BE) method. For $\theta = 0.5$, this equation is second-order accurate and is usually referred to as the Crank–Nicolson (CN) method. While CN methods are relatively simple and very accurate, they have the unfortunate property of generating oscillating solutions when large time steps are used. The time averaged solution is correct, but the time and spatially resolved solutions show oscillations. In the two-dimensional problem to be presented in Section 5, this behavior shows up as circular ripples or waves near the origin. Of course, ripples in a purely diffusive problem are non-physical. In order to avoid these oscillations, authors [13,14] have discussed the following prescription for time advancement:

$$E^{n+1} = \frac{4}{3}E^n - \frac{1}{3}E^{n-1} + \frac{2\Delta t}{3}\left.\frac{\partial E}{\partial t}\right|^{n+1}. \tag{11}$$

This second-order temporal differencing uses three time levels and is fully implicit in the time derivative. The truncation error for this method with constant time step is $\frac{2(\Delta t)^3}{9}\frac{\partial^3 E}{\partial t^3}$, while for $\theta = 0.5$ in Eq. (10) the truncation error is $\frac{(\Delta t)^3}{12}\frac{\partial^3 E}{\partial t^3}$. So CN is theoretically more accurate by a factor of 8/3, but it has the disadvantage of oscillating when using large time steps. The biggest disadvantage of Eq. (11) is that one must keep around information from an additional time level. In the literature this three level method is sometimes referred to as BDF2, backward difference formula second-order. It has been used to successfully solve the Navier–Stokes and energy equations with phase change effects using a Jacobian-free Newton–Krylov method [15]. When using variable time steps, Eq. (11) must be generalized to

$$E^{n+1} = aE^n + (1-a)E^{n-1} + b\Delta t \left.\frac{\partial E}{\partial t}\right|^{n+1}, \tag{12a}$$

$$a = \frac{(1+\rho)^2}{1+2\rho}, \quad b = \frac{1+\rho}{1+2\rho}, \quad \rho = \frac{t^{n+1} - t^n}{t^n - t^{n-1}}. \tag{12b}$$

In practical calculations, one limits how quickly the time step is allowed to increase so that $0.1 \leqslant \rho \leqslant 1.25$. Therefore, the coefficients in Eq. (12) are very well behaved. Because of its robustness over a wide range of

time step sizes, the BDF2 method in Eq. (12) will be the preferred numerical algorithm in this paper. Later, in Section 6, a comparison will be made to CN results.

Inserting Eq. (1) into Eq. (12) and rearranging terms gives

$$(1 + bc\Delta t\sigma_a^{n+1})E^{n+1} - bc\Delta t\nabla\cdot(D^{n+1}\nabla E^{n+1}) = aE^n + (1-a)E^{n-1} + bc\Delta t\sigma_a^{n+1}(T^{n+1})^4. \tag{13}$$

The superscripts on the material coefficients indicate that they are evaluated at the conditions existing at the new time step. Doing the same for material temperature and Eq. (2) gives

$$\left[1 + bc\Delta t\frac{\sigma_a^{n+1}}{c_V^{n+1}}(T^{n+1})^3\right]T^{n+1} - \frac{bc\Delta t}{c_V^{n+1}}\nabla\cdot(K^{n+1}\nabla T^{n+1}) = aT^n + (1-a)T^{n-1} + bc\Delta t\frac{\sigma_a^{n+1}}{c_V^{n+1}}E^{n+1}. \tag{14}$$

These two equations are coupled and are nonlinear in T; therefore, to preserve the second-order accuracy, one must specify how to solve them.

The nonlinearity in temperature can be dealt with by the method introduced in [1]. Use the Taylor series expansion

$$(T^{n+1})^4 \approx (T^*)^4 + \Delta T\frac{\partial(T^*)^4}{\partial T} = (T^*)^4 + (T^{n+1} - T^*)4(T^*)^3 = (T^*)^3(4T^{n+1} - 3T^*), \tag{15}$$

where T^* is a reference temperature. When doing first-order calculations, one can choose either the old or new time step value for this reference temperature. However, in order to get second-order accuracy, one should use the previous iteration's value of T^{n+1} for T^*. As the iterations converge, the equation coefficients will be evaluated at the correct implicit temperature. Substituting Eq. (15) into Eq. (14) gives

$$\left[1 + 4bc\Delta t\frac{\sigma_a^{n+1}}{c_V^{n+1}}(T^*)^3\right]T^{n+1} - \frac{bc\Delta t}{c_V^{n+1}}\nabla\cdot(K^{n+1}\nabla T^{n+1}) = aT^n + (1-a)T^{n-1} + \frac{bc\Delta t\sigma_a^{n+1}}{c_V^{n+1}}[E^{n+1} + 3(T^*)^4]. \tag{16}$$

In one dimension, Eqs. (13) and (16) can be solved separately as tridiagonal systems of equations. In two dimensions, each is a system with five diagonals. If an energy conserving form had been used in Eq. (2), Eq. (16) would have a more complicated structure. The simplest solution scheme is to solve Eq. (13) for E^{n+1} with the current estimate for T^{n+1}, then solve Eq. (16) for T^{n+1} using the latest value for E^{n+1}. With an updated temperature the material coefficients must be recalculated and the process continued. These equations and their solution will be referred to in this paper as the non-split technique. Unfortunately, this iteration scheme converges quickly only if the time step is very small. An example will be discussed in Section 5.

The above technique converges slowly because the radiation and material energy balance equations are tightly coupled and the above algorithm lags this coupling. One improvement would be to solve the two equations simultaneously. This doubles the size of the matrix to be solved and can greatly increase the computational cost for each iteration. In large two- and three-dimensional problems, this may not be desirable. A comparison will be made at the end of Section 5.

Many authors [2–4,7,8] have explored solving the above equations with Newton's method for solving coupled nonlinear equations or have used Newton–Krylov techniques. These methods work, but can have a large computational cost. Most published second-order solutions are much slower computationally than backward Euler. Therefore, we will explore operator splitting techniques that can partially improve the nonlinear coupling without much added computational cost.

There are many operator splits that have been proposed and tested [1–8] for both first- and second-order accuracy. Some of the second-order algorithms use several steps and can be difficult to implement. Here, we will use a relatively simple two-step process. In the first phase, the full radiation and material equations will be solved with all but the nonlocal part of the material equation. Then in the second phase, the nonlocal part of the material equation will be solved along with a local update of the radiation energy density. How this differs from previous work is in the important details presented next.

The first phase of the operator split will advance the variables from time level n to an intermediate time labeled with primes. The radiation equation will include all the terms in Eq. (13) while the material equation will include all terms in Eq. (16) except for the implicit conduction term. It will be lagged in time and moved to the right-hand side. The material equation then becomes

$$\left[1+4bc\Delta t\frac{\sigma_a^{n+1}}{c_V^{n+1}}(T^*)^3\right]T'=aT^n+(1-a)T^{n-1}+bc\Delta t\frac{\sigma_a^{n+1}}{c_V^{n+1}}[E'+3(T^*)^4]+\frac{bc\Delta t}{c_V^n}\nabla\cdot(K^n\nabla T^n). \tag{17}$$

The last term in Eq. (17) may look like a nonlocal term, but this divergence at the old time level can be calculated once at the beginning of each time step and then it is just an additive constant term during the iteration process. In the next phase, this term will be subtracted from the right-hand side while added implicitly at the advanced time on the left-hand side. By including this term in this step, as much information as possible is being used to advance the intermediate temperature. Inserting Eq. (17) into Eq. (13) and rearranging gives the following radiation equation:

$$\left[1+\frac{bc\Delta t\sigma_a^{n+1}}{1+4bc\Delta t\frac{\sigma_a^{n+1}}{c_V^{n+1}}(T^*)^3}\right]E'-bc\Delta t\nabla\cdot(D^{n+1}\nabla E')$$
$$=aE^n+(1-a)E^{n-1}+\frac{4bc\Delta t\sigma_a^{n+1}(T^*)^3}{1+4bc\Delta t\frac{\sigma_a^{n+1}}{c_V^{n+1}}(T^*)^3}\left[aT^n+(1-a)T^{n-1}-\frac{3}{4}T^*+\frac{bc\Delta t}{c_V^n}\nabla\cdot(K^n\nabla T^n)\right]. \tag{18}$$

This equation includes all the information known at this point of the iteration process about both the radiation and conduction physics. The current estimated value for T^{n+1} is used for T^*. This system of equations for E' has more complicated diagonal and right-hand side terms but is no more difficult to solve than Eq. (13). In multi-dimensions where iterative techniques are used for the matrix solution, Eq. (18) may require more iterations to solve because it is less diagonally dominant than Eq. (13). After solving for E', Eq. (17) is easily evaluated for T'.

The second phase of the operator split is to do the nonlocal conduction for the material and update the radiation energy density for the new temperature:

$$T^{n+1}-\frac{bc\Delta t}{c_V^{n+1}}\nabla\cdot(K^{n+1}\nabla T^{n+1})=T'+bc\Delta t\frac{\sigma_a^{n+1}}{c_V^{n+1}}[E^{n+1}-(T^{n+1})^4-E'+(T')^4]-\frac{bc\Delta t}{c_V^n}\nabla\cdot(K^n\nabla T^n), \tag{19}$$

$$E^{n+1}=E'+bc\Delta t\sigma_a^{n+1}[(T^{n+1})^4-E^{n+1}-(T')^4+E']. \tag{20}$$

In Eq. (19), note that the conduction at the old time step that was added in Eq. (17) now is being subtracted from the right-hand side. The right-hand side of Eq. (20) and the first part of (19) are just adding the energy exchange between the radiation and material at the new time level and subtracting the energy exchange at the intermediate prime level. Eq. (20) can trivially be solved for the radiation energy density and then substituted into a temperature linearized version of Eq. (19) to get

$$E^{n+1}=E'+bc\Delta t\sigma_a^{n+1}[(T^{n+1})^4-(T')^4]/(1+bc\Delta t\sigma_a^{n+1}), \tag{21}$$

$$\left[1+\frac{4bc\Delta t\sigma_a^{n+1}(T^*)^3/c_V^{n+1}}{1+bc\Delta t\sigma_a^{n+1}}\right]T^{n+1}-\frac{bc\Delta t}{c_V^{n+1}}\nabla\cdot(K^{n+1}\nabla T^{n+1})$$
$$=T'+\frac{bc\Delta t\sigma_a^{n+1}/c_V^{n+1}}{1+bc\Delta t\sigma_a^{n+1}}[(T')^4+3(T^*)^4]-\frac{bc\Delta t}{c_V^n}\nabla\cdot(K^n\nabla T^n). \tag{22}$$

Solving Eq. (22) and then Eq. (21) completes the process of advancing to the next time level. However, the material coefficients must be evaluated at the new temperature and the process iterated to convergence. The solution method represented by Eqs. (17), (18), (21) and (22) will be referred to in this paper as the split, or operator split technique.

If the above process were repeated for the backward Euler equation, one would find that terms like $aE^n+(1-a)E^{n-1}$ reduce to just E^n and one would set $b=1$. Since b is less than unity for the three time level method, the implicit terms in all the above equations that must be iterated to convergence are smaller in the BDF2 method than for the BE method. The terms are identical with the same temperature and time dependences, only the coefficient is smaller for the BDF2 method. Therefore, in general, this second-order solution should converge in fewer iterations than a fully implicit first-order solution. It is surprising to this author that all the published comparisons between first- and second-order methods (known to this author) show that sec-

ond-order methods are computationally slower, often much slower, than fully implicit first-order methods. That fact was a major motivation in publishing this current work.

In one dimension, the nonlocal systems equations that arise from Eqs. (18) and (22) are tridiagonal equations and are thus trivial to solve. In two dimensions, the corresponding systems have five diagonals with two diagonals adjacent to the main diagonal and the other two spaced out by the number of zones in one of the two dimensions. This sparse system is solved here using the "linbcg" routine given in [16]. This may not be the fastest algorithm for this purpose; however, it is reliable and readily available. We will use it with its trivial default of diagonal preconditioning. The choice of algorithm for the sparse system solver does not affect how quickly the temperature iterations converge; it just affects the total computational time. The exact same linear solver will be used for both first- and second-order methods. A difference choice for solver and preconditioner might be faster, but it would affect the above first- and second-order methods uniformly. It would not affect the conclusions presented here.

One must specify the criteria for deciding when one of the above iterations methods is adequately converged. Traditionally, the main criterion is that the maximum relative change in E or T in any cell is less than a small quantity, typically $1.0e-7$. It is often argued that a very tight convergence is required for second-order time evolution. We will test this assumption in Section 5. In difficult nonlinear problems there are some time levels when the solution can be very tough to converge; therefore, a maximum number of iterations is also specified. Here, we will use 20 iterations as the default maximum but will study other values.

The time steps will be adjusted in size to give the maximum relative change in E or T that is desired. Define the time step tolerance for change as

$$\eta = \max\left[\max_i\left(\frac{2\left(E_i^{n+1}-E_i^n\right)}{E_i^{n+1}+E_i^n}\right), \max_i\left(\frac{2(T_i^{n+1}-T_i^n)}{T_i^{n+1}+T_i^n}\right)\right], \tag{23}$$

where the maximum over i is symbolically over all cells, whether the problem is 1D or 2D. The properties of the numerical solutions will be studied in the range of $0.0125 \leqslant \eta \leqslant 1.25$. These two orders of magnitudes cover the practical applications of the techniques presented here. This definition is similar in philosophy, but different in detail, with the scaling of dynamical times scales to determine time steps as used by other authors.

4. Flux limiter discretization

In order to discuss flux limiting more clearly, it is useful to look at the radiation flux, which is the physical quantity inside the divergence in Eq. (1):

$$F = -D\nabla E, \tag{24a}$$

or, in component form:

$$F_x = -D\frac{\partial E}{\partial x} \quad \text{and} \quad F_y = -D\frac{\partial E}{\partial y}. \tag{24b}$$

Radiation flux is the rate of radiation energy flowing through a unit area in a given direction. This diffusion equation is derived by dropping a time derivative term from the transport equation. This is an appropriate and accurate approximation in optically thick problems where the radiation mean free path is small compared to the dimensions of the medium. However, in a low opacity region, mean free paths are large and the diffusion equation has too large of a propagation velocity. In order to apply the diffusion approximation everywhere, one often uses a flux limiter. In the units used here, saying that the propagation velocity must be less than the speed of light is equivalent to saying that the magnitude of the flux must be less than the energy density: $|F| \leqslant E$. Although the components of the flux are across the faces of a 2D cell, the opacity in the diffusion coefficient is a cell-centered quantity. Therefore, the normalized gradient in Eqs. (3) and (4) is often evaluated at the centers of the cells. Here, we will examine both cell-centered and face-centered flux limiters.

It is important to remember that flux limiting is a crude fix to a complicated transport problem. The full radiation transport equations are dependent on angle. The diffusion approximation assumes a linear function in angle in addition to dropping the time derivative of the flux. See [9,12] for more information on this topic.

In order to conserve energy locally, it is necessary that the radiation that flows out of one cell must be identical to the energy flowing into the next cell. To demonstrate this concept in 1D, use a mesh with the energy densities at the cell centers labeled by i. Assume that the control volume coincides with the cell and that temperatures and diffusion coefficients are constant within a cell. Introduce a temporary face value of $E_{i+1/2}$ and assume that the energy density is a piecewise linear continuous function. Then equate the flux evaluated from the left side of the face to the flux evaluated from the right side:

$$F_{i+1/2} = -D_i \frac{E_{i+1/2} - E_i}{\Delta x/2} = -D_{i+1} \frac{E_{i+1} - E_{i+1/2}}{\Delta x/2}, \tag{25}$$

where the two cells are assumed to be of equal size and the factor of one half in the denominator is because the difference is only over half of each cell. Solve this for $E_{i+1/2}$ and substitute it back into Eq. (25) to get

$$E_{i+1/2} = \frac{D_i E_i + D_{i+1} E_{i+1}}{D_i + D_{i+1}}, \tag{26}$$

$$F_{i+1/2} = -\frac{2D_i D_{i+1}}{D_i + D_{i+1}} \frac{E_{i+1} - E_i}{\Delta x} \equiv -\overline{D}_{i+1/2} \frac{E_{i+1} - E_i}{\Delta x}. \tag{27}$$

Therefore, the correct diffusion coefficient across a cell face is the harmonic average of the two cell-centered coefficients. In two dimensions, each of the four faces of a quadrilateral cell will have a different harmonically averaged diffusion coefficient. Only this harmonic mean coefficient will exactly conserve energy locally between two cells. Also, when cells are optically thick, this harmonic mean gives a much more accurate solution to the diffusion equation. Simple arithmetic means work well only on highly resolved meshes where the cells are less than a mean free path thick.

With this discussion as background, it seems reasonable to use the results of Eq. (26) to form a normalized gradient as follows:

$$\chi_i = \left(\frac{1}{E} \left| \frac{\partial E}{\partial x} \right| \right)_i = \frac{|E_{i+1/2} - E_{i-1/2}|}{\Delta x E_i} = \frac{1}{\Delta x E_i} \left| \frac{D_i E_i + D_{i+1} E_{i+1}}{D_i + D_{i+1}} - \frac{D_i E_i + D_{i-1} E_{i-1}}{D_i + D_{i-1}} \right|. \tag{28}$$

However, the final equality shows that this gradient is dependent on the diffusion coefficients that are dependent on the normalized gradient. This circular definition creates a nonlinear dependence that can generate unreliable diffusion coefficients. Therefore, one should use a gradient based directly on the radiation energy densities without the diffusion coefficient weighting. We will examine three possible choices, illustrating them in 1D calculations. One could use a geometric average of forward and backward differences:

$$\text{(a) } \chi_i = \sqrt{\left| \frac{4(E_{i+1} - E_i)(E_i - E_{i-1})}{\Delta x^2 (E_{i+1} + E_i)(E_i + E_{i-1})} \right|}. \tag{29a}$$

Or one could use a weighted algebraic average of forward and backward differences:

$$\text{(b) } \chi_i = \frac{1}{2\Delta x} \left| \frac{2(E_{i+1} - E_i)}{(E_{i+1} + E_i)} + \frac{2(E_i - E_{i-1})}{(E_i + E_{i-1})} \right| = \frac{2E_i |E_{i+1} - E_{i-1}|}{\Delta x (E_{i+1} + E_i)(E_i + E_{i-1})}. \tag{29b}$$

Or one could use a simple central difference:

$$\text{(c) } \chi_i = \frac{|E_{i+1} - E_{i-1}|}{2\Delta x E_i}. \tag{29c}$$

Case (c) is probably the most common choice. In principal, one could expect that all three should give nearly identical results. In practice, on poorly resolved meshes where there are steep gradients in just a few cells, it turns out that they give significantly different results and these differences impact the accuracy of the time evolution.

As a test problem, we will use the 1D blast wave used by other authors [7]. The initial condition is a peak of radiation near the origin:

$$E(x) = 0.001 + 100 \exp[-100x^2]. \tag{30}$$

Reflecting boundary conditions are used at the origin and at $x = 3$. The opacity inversely proportional to the cube of the temperature is used with $Z = 1$ and the heat capacity is constant, $c_V = 1$. Initially the material temperature is set equal to the radiation temperature, $T(x) = T_r(x)$. Then the equations are evolved in time to $ct = 3$.

Some comparisons are shown in Figs. 1 and 2. For the case of no flux limiter, there is only a slight difference using these two different mesh resolutions. Also, when using the geometrical gradient in Fig. 1, the two meshes give very similar results. As expected, using $m = 1$ in the flux limiter slows the wave front more than $m = 2$ does. The surprise is in Fig. 2 where using the central difference for the gradient shows a large mesh sensitivity. As the mesh is refined further, the central difference solutions do approach those using the geometric gradient. Case (b) with the weighted algebraic average gives results between the other two; however, they are much closer to the results using (a). Since cases (a) and (c) appear to be the extremes of these three cases, results from (b) will not be shown. One concludes from these two figures that using a flux limiter with the simple central difference gradient creates a mesh sensitivity that does not exist in the underlying equations.

An alternative to the above cell-centered flux limiter is to use a face-centered limiter with a slightly different formulation:

$$F = -\frac{\lambda}{3\sigma_t}\nabla E \equiv \lambda f, \quad \lambda \equiv \frac{1}{\sqrt[m]{1+\xi^m}}, \tag{31}$$

$$\xi \equiv \frac{|\nabla E|}{3\sigma_t E} = \frac{|f|}{E}, \tag{32}$$

where f is the scaled flux. If one repeats the derivation leading up to Eqs. (26) and (27) with this definition of the flux, one arrives at

$$E_{i+1/2} = \frac{\Delta x_{i+1}\sigma_{t,i+1}E_i + \Delta x_i \sigma_{t,i}E_{i+1}}{\Delta x_{i+1}\sigma_{t,i+1} + \Delta x_i \sigma_{t,i}}, \tag{33}$$

$$F_{i+1/2} = -\frac{2\lambda_{i+1/2}(E_{i+1} - E_i)}{3(\Delta x_{i+1}\sigma_{t,i+1} + \Delta x_i \sigma_{t,i})}. \tag{34}$$

In 1D, the normalized gradient that appears in the flux limiter can be expressed in terms of the energy gradient or the scaled flux, whichever is more convenient:

$$\xi_{i+1/2} = \frac{2|E_{i+1} - E_i|}{3(\Delta x_{i+1}\sigma_{t,i+1}E_i + \Delta x_i \sigma_{t,i}E_{i+1})} = \frac{|f_{i+1/2}|}{E_{i+1/2}}. \tag{35}$$

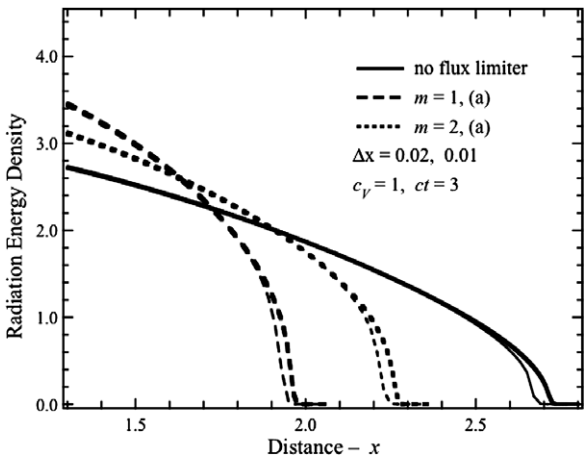

Fig. 1. The radiation energy density at time $ct = 3$ is shown with and without flux limiters on two different spatial grids. For each pair of lines the thinner one has a spatial resolution of 0.02 while the thicker line has a resolution of 0.01. The flux limiter uses the geometric differenced gradient.

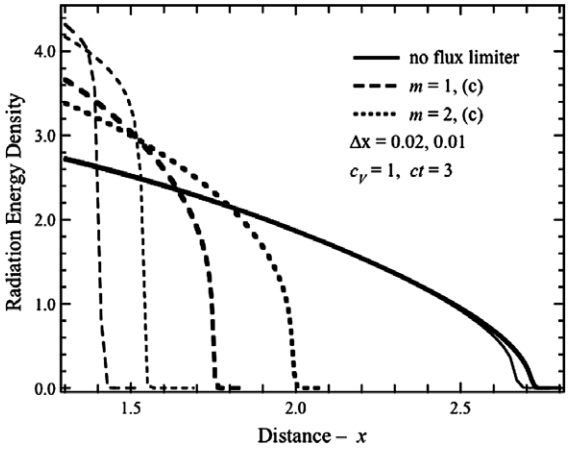

Fig. 2. The radiation energy density at time $ct = 3$ is shown with and without flux limiters on two different spatial grids. For each pair of lines the thinner one has a spatial resolution of 0.02 while the thicker line has a resolution of 0.01. The flux limiter uses the central differenced gradient.

The advantage of using this form of flux limiting on cell faces is that, in 1D, it is straightforward to make its discretization consistent with the flux. However, in multi-dimensions the components of the flux are not evaluated on the same face; therefore, one must average either energy gradients or flux values in order to calculate normalized gradient for the flux limiter.

In 1D, this face flux limiter gives results very similar to the results in Fig. 1, except that the separation between the curves for the two spatial resolutions is smaller. If one defines the position of the wave front as the point where $E = 0.1$, this position can be determined very accurately because the radiation energy density is known both at cell centers and cell faces. For the case of using no flux limiter, the wave front moves forward by 0.046 when the mesh resolution is reduced from 0.02 to 0.01. For the cell-centered geometric gradient flux limiter and $m = 2$, this position difference is 0.029. For the face-centered flux limiter with $m = 2$, the wave front moves backward by only 0.010 with this resolution change. Therefore, this face-centered flux limiter appears to be less sensitive to mesh resolution than non-flux-limited diffusion or any of the cell-centered flux limiters examined here.

The next step is to look at the convergence behavior of the above 1D problem as the time step is refined. As a convergence metric we use the root mean square difference of the radiation energy density from a reference calculation:

$$\text{RMS Error} \equiv \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left[E(x_i, ct) - E_i^{\text{ref}}\right]^2}. \tag{36}$$

The reference solution is done with $\eta = 0.00125$, i.e., 10 times smaller than the smallest time step tolerance shown in the figures. Some authors calculate the RMS error of the radiation temperature rather than the energy density. Since the equations are being solved for the energy density and the radiation temperature is a transformation of this variable, we prefer to use E. In practice it makes little difference, but this measure of error seems to be closer to the differential equations.

Fig. 3 shows the convergence with time step and mesh resolution for the backward Euler (BE) and three time level (BDF2) methods using cell-centered flux limiters. The most significant feature in this figure is that the errors for the central difference gradient are more than an order of magnitude larger than for the geometric gradient. Secondly, the slopes for the geometric gradient calculations are very close to the expected values of 1 and 2. As the RMS Error approaches unity, the curves deviate from the expected straight lines. At this point it must be clearly stated that the reference calculations used the same spatial mesh and the same flux limiter as the calculations being tested. Only the time step was smaller. The more accurate BDF2 results are used as the reference for the BE calculations. The poor spatial convergence of the central difference in Fig. 2 shows up in Fig. 3 as large errors in the temporal convergence.

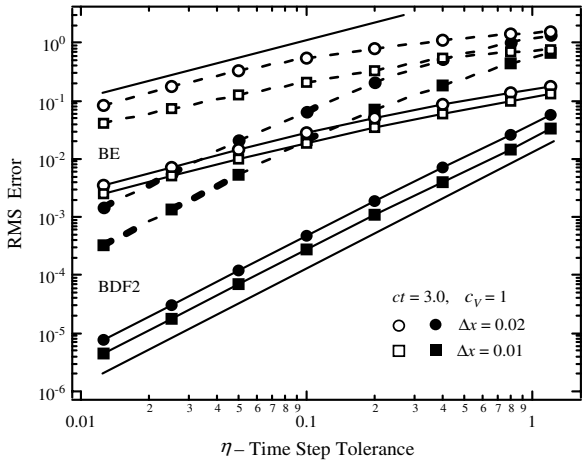

Fig. 3. Dashed lines show the errors when using central differences in the gradient. Solid lines are for geometric differences. Filled symbols are used for the BDF2 method, while open symbols are used for backward Euler. The shape of the symbol designates the mesh resolution. The thin lines show slopes of 1 and 2.

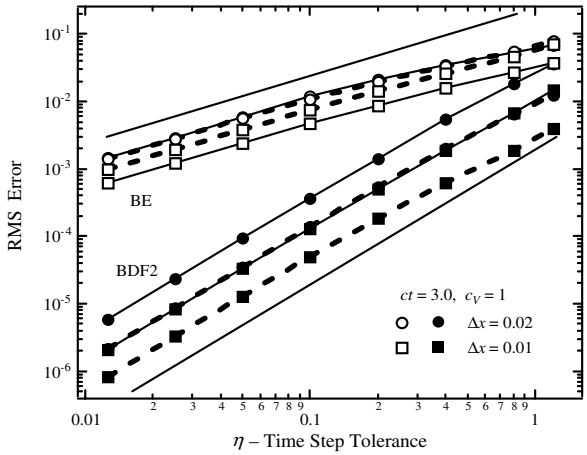

Fig. 4. Solid lines show the errors for no flux limiting. Dashed lines are for the cell face flux limiter and $m = 2$. Filled symbols are used for the BDF2 method, while open symbols are used for the BE method. The thin lines show slopes of 1 and 2.

Fig. 4 show the time step convergence with no flux limiter ($\chi = 0$ or $\xi = 0$) and when using the face-centered flux limiter given in Eqs. (31)–(35). The errors here are somewhat smaller than those for the cell-centered geometric gradient and much less than those for the cell-centered central differenced gradient. It appears that rapid convergence with spatial resolution is correlated with temporal accuracy. As a side note, if one solves the non-split Eqs. (13) and (16) simultaneously as a large system of equations, the error plot looks very similar to Figs. 3 and 4. Therefore, the operator splitting does not appear to affect the temporal convergence.

5. Two-dimensional results

Mousseau and Knoll [3,8] presented an interesting two-dimensional multimaterial test problem that has a central blast wave moving out around two square obstacles. On a unit square Cartesian region with uniform zoning the equations are solved here on either a 64^2 or 128^2 spatial mesh. The background material uses $Z = 1$ for the opacity in Eq. (5), while the obstacles use $Z = 10$. These obstacles are located at

$$x_1 = 0.3125 \pm 0.125, \quad y_1 = 0.6875 \pm 0.125,$$
$$x_2 = 0.6875 \pm 0.125, \quad y_2 = 0.3125 \pm 0.125.$$

The edges of these obstacles correspond exactly to cell edges on these two meshes; therefore, there are no mixed material cells. The initial energy distribution has a Gaussian peak near the origin:

$$E(x,y) = 0.001 + 100\exp[-100(x^2 + y^2)]. \tag{37}$$

The material temperature is initialized to be equal to the radiation temperature. All edges of the region use reflection boundary conditions. So the initial radiation spreads out and flows around the obstacles. When using a constant heat capacity, $c_V = 1$, the solutions will be compared at a time of $ct = 3$. When using the Saha heat capacity, the radiation flows more slowly so the comparison time will be at $ct = 6$. Mousseau and Knoll [3,8] analyzed this test problem for the case of $c_V = 1$, $m = 1$, $ct = 3$, three mesh sizes, and the central difference gradient. Here, we use two different heat capacities, $m = 2$, two mesh sizes, and the geometric difference for the normalized gradient in the cell-centered flux limiter. The different choice for the flux limiter does not give a big difference in a visual display of the solution; therefore, the $c_V = 1$ case looks very similar to Mousseau and Knoll's figures [3,8]. Fig. 5 shows contours of the radiation energy distribution at $ct = 6$ when using the Saha heat capacity and the 128^2 mesh. The radiation has flowed around the opaque squares and is flowing into them from all sides. The peak energy density at the origin has a value of 0.32; therefore, the peak temperature is 0.75. Since the temperature drops from this value down to $0.18 = (0.001)^{0.25}$, the nonlinearities of the Saha ionization model strongly affect the solution.

For the two different heat capacities, Fig. 6 shows the error convergence as the time step tolerance is changed. Because this test problem has two materials and the material interfaces are aligned with the mesh, a finer mesh has steeper spatial gradients at the material interfaces. Therefore, the mean error per cell decreases only slightly as the mesh is refined. The simple trends with mesh resolution seen in 1D in Fig. 3 are repeated in 2D in Fig. 6. For backward Euler, the errors very accurately lie on straight lines with slopes of unity for small time steps, $\eta \leqslant 0.2$. For the three time level method, the errors closely follow straight lines with the expected slope of 2 over the whole range of time steps.

If the cell-face flux limiter had been used instead of the cell-center flux limiter, Fig. 6 would change only slightly. The main difference is that the calculations at different mesh resolutions would lie closer together. Just as in 1D, the face flux limiter appears to be less sensitive to mesh resolution than the cell center limiter.

The next step is to look at the relative computational efficiency of the numerical methods presented here. One important efficiency detail is that at the end of a time step one should always do a linear extrapolation to get the initial guess for the next time step:

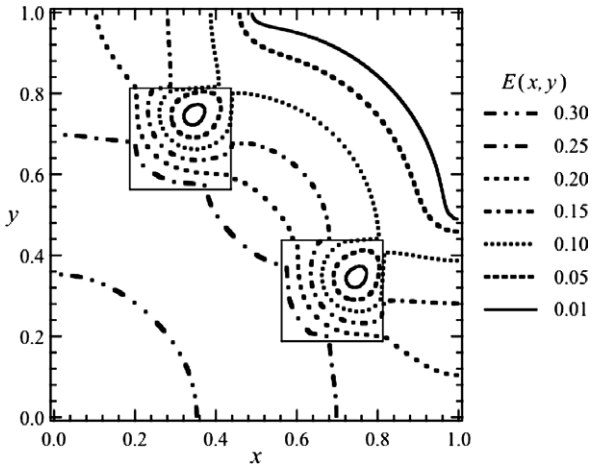

Fig. 5. Contours of the radiation energy density at $ct = 6$ for the case with a Saha heat capacity. The squares are the more opaque obstacles.

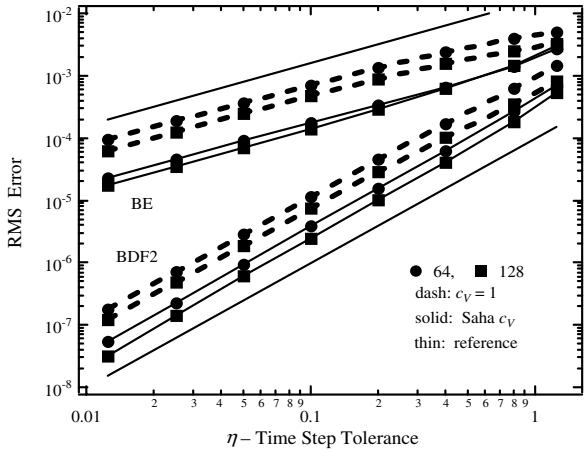

Fig. 6. The root mean square error relative to a reference solution is shown for the two numerical methods two mesh sizes, and two different heat capacity models. The quadratic flux limiter is used with the geometric gradient.

$$E^{n+2} = E^{n+1} + \frac{t^{n+2} - t^{n+1}}{t^{n+1} - t^{n}}(E^{n+1} - E^{n}). \tag{38}$$

This same type of extrapolation must also be done for the material temperature. Then the opacities, diffusion coefficients, conduction coefficients, and heat capacities must be evaluated at these extrapolated temperatures before beginning the iteration process for the new time step. Sometimes this extrapolation saves just one iteration, other times it can save several iterations. Using the current time level as an initial guess for the next time level does work, but not as well. One could use a forward differenced Euler equation that explicitly predicts the values at the new time step; however, it does not consistently give better results than the simple extrapolation in Eq. (38). One of the problems with solution techniques that depend on Newton or Newton–Krylov methods is that they can be sensitive to initial guesses. If the initial trial value is not close enough to the final solution, the technique may not converge. Here, with simpler iterations methods, the solutions always converge. The question is how to best accelerate the convergence. For some problems a better initial guess is obtained by only using 0.9–0.95 of the extrapolation in Eq. (38). In the present work the full extrapolation will be used.

Fig. 7 shows how the errors decrease with increasing computer time. Obviously, the finer spatial mesh takes more computer time than the coarse mesh. For the four pairs of curves, the one on the right, the more com-

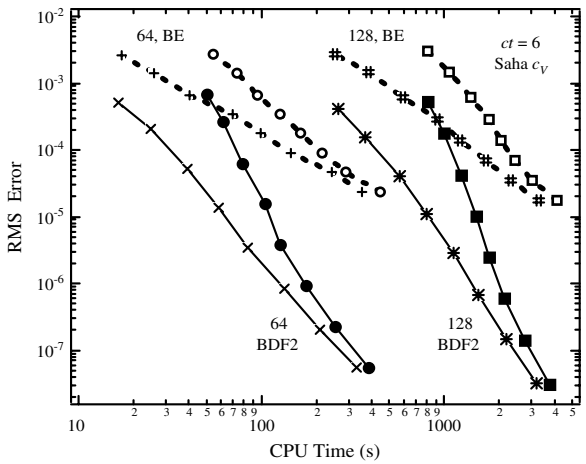

Fig. 7. The error is shown as a function of computer time in seconds. These curves are for the Saha heat capacity and the quadratic flux limiter with the geometric gradient. For each of the four pairs of curves, the left one uses a maximum of three iterations, while the right one may use up to 20 iterations.

putationally expensive curve, uses the tight convergence tolerance discussed in Section 3, a relative change of $1.0e-7$, with a maximum of 20 iterations. The left curve in each pair stops with only three iterations, no matter how well the solution is converged. The surprising aspect of these figures is that the error is roughly the same, independent of whether or not tight convergence was required. It appears that the error at the final time has not accumulated over the time steps. The solution in most cells changes very little from one time to the next. In those cells the solution converges immediately. Only the cells at the wave front are changing rapidly in time and are slow to converge with iteration. It seems that this error at the wave front caused by incomplete convergence does not affect the RMS error at the final time. In other test problems with other error criteria, it may be necessary to converge to high accuracy all the nonlinearities at all time levels in order to get an accurate final solution. If the wave front position is critically sensitive to the exact solution in the wave front cells, then stopping at three iterations would not be appropriate. The results for a constant heat capacity are very similar to Fig. 7. The results when using the face-centered flux limiter are just shifted downward slightly.

Note that the computational time required by the three time level solutions is comparable to or less than the backward Euler solutions. Greatly improved accuracy over first-order methods can be achieved in this problem without an increase in computational resources. On the finer mesh with large time steps, BDF2 is roughly a factor of six more accurate than BE. When using the smallest time step control, BDF2 is more than a factor of 500 times more accurate the BE. Whether one prefers cell center or face flux limiters does not significantly affect the computer time.

Since this test problem has reflecting boundary conditions, the total energy in the problem should be conserved. At $ct = 6$ with the Saha heat capacity and the 64^2 mesh, for BDF2 the error in the energy for $\eta = 1.25$ is just 0.065% and varies roughly as η^2. For BE and $\eta = 1.25$, the error is 0.4% and it decreases linearly with η. On the finer mesh, the energy errors are only slightly smaller. The size of these energy errors does not increase when the maximum number of iterations is dropped to three. In problems where energy conservation is crucial, other accuracy constraints will probably keep the time step change tolerance much less than unity. Therefore, these conservation numbers seem to be quite manageable for most real world problems. As an alternative, one could reformulate the energy equation to be in a conservative form. This was not done here because it added to the iteration complexity.

After showing that simple iterations of the operator split equations can be a viable effective solution technique, it is interesting to go back to test the temperature linearized but non-split Eqs. (13) and (16). Unfortunately, without the operator splitting, these equations converge much too slowly. For the BDF2 method and $\eta > 0.05$, 20 iterations are not enough for convergence. Therefore the errors build up and the solution drifts away from the correct solution. The errors in the backward Euler solution become objectionable for $\eta > 0.2$.

Another solution strategy is to solve the non-split Eqs. (13) and (16) simultaneously as a coupled system of 2×64^2 or 2×128^2 sparse equations. For each mesh point the radiation energy density and temperature are placed next to each in the matrix with their spatial neighbors generating both near and widely spaced band structure. Unfortunately, the iterative solution of these larger systems of equations takes more computational time than solving the two smaller systems of equations separately. When solving this coupled system, fewer overall iterations are required to converge the material coefficients, but each matrix solve is more expensive. While the accuracy is similar, for the case shown in Fig. 7, limited to just three iterations, the increase in cpu time varies from a factor of 1.3–3.8. For BDF2 and small values of η, the increase in computational cost is modest. However, for BE and large values of η, the added cost is more than a factor of three on the 64^2 mesh. On the finer mesh, the additional time is slightly less, probably because the time steps are smaller and one has a better initial guess. These results all use the *linbcg* linear solver [16] with the trivial default of diagonal preconditioning. Using a more sophisticated solver or preconditioner would affect all the timings and may change the comparison between solving one large system or two smaller systems.

6. Comparison to Crank–Nicolson time discretization

One can use $\theta = 0.5$ in Eq. (10) and follow the derivation in Section 3 to arrive at a CN equivalent scheme for solving the coupled radiation material equations. The equations will not be presented here because they are just a tedious variation of the earlier equations and have no inherent value.

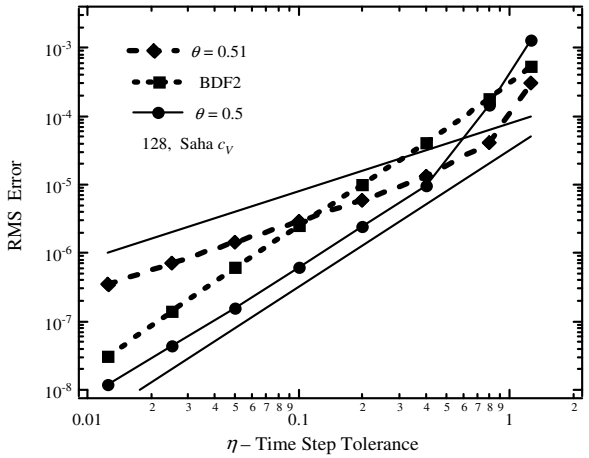

Fig. 8. The root mean square errors are shown for three methods. The thin lines have slopes of 1 and 2.

However, Fig. 8 shows how the BDF2 method compares to using $\theta = 0.5$ and $\theta = 0.51$. By using a value for θ that is slightly greater than one half, the oscillations present when using large time steps in the Crank–Nicolson method are greatly reduced. Unfortunately, it also then becomes a first-order method with a small truncation error rather than a true second-order accurate method. For small time steps the three time level method is almost exactly a factor of $^8/_3$ above CN, just as predicted in Section 3. Its error curve is a straight line with no surprises and has a slope of 2.12 over it whole length. Using $\theta = 0.5$ works fine for $\eta \leqslant 0.4$, but in this problem it gives oscillatory solutions for $\eta = 0.8$ and 1.25. The computational times and energy conservation for the CN and BDF2 methods are very similar for small time steps. Therefore, a figure equivalent to Fig. 7 using CN would just show the curves moved down slightly for most values of η. However, for $\eta = 1.25$, the error is smaller for the three time level method than for Crank–Nicolson.

For the class of problems studied in this research, the Crank–Nicolson method is very accurate and reliable only for small and moderate time steps. When using large time steps, the nonphysical oscillations yield unacceptably large errors. While the three time level method is slightly less accurate when using small time steps and it requires keeping information from a previous time step, it is reliably accurate over the whole range of time steps tested. Therefore, we highly recommend the BDF2 method.

7. Summary

This paper has demonstrated that a simple iteration scheme can yield very accurate second-order solutions to the coupled radiation diffusion and material conduction problem. Also, it is not necessary to completely converge the nonlinearities in every time step to obtain these accurate solutions. Simpler algorithms are easier to convert into computer programs that are less likely to contain mistakes. The operator splitting presented here can be used with either the three-time level (BDF2) or the Crank–Nicolson (CN) discretization method; however, CN is unreliable for large time steps. The greatly improved accuracy is achieved without taking more computer time than a simple first-order backward Euler (BE) calculation. It is quite feasible to modify existing BE codes to use the algorithms presented here.

One surprise uncovered by this work is that details of how the gradient in the flux limiter are differenced can greatly influence the spatial and temporal convergence. The most obvious choice of simple cell centering with central differencing gives the slowest convergence with spatial resolution. This choice also strongly affects the temporal convergence. Weighted algebraic and geometric differencing give much faster spatial and temporal convergence. The discretization errors in the flux limiter can mask the convergence trends of the underlying system of equations. In the problems tested here, it appears that placing the flux limiter on cell faces rather than at cell centers gives results that are less sensitive to the spatial mesh resolution.

The two-dimensional test problem used here was developed in [3,8]. We extended it by also using the Saha ionization model of [2,7] that creates a nonlinear heat capacity. One possible next step would be to test the technique presented here in a full 2D radiation-hydrodynamics calculation.

References

[1] G.L. Olson, J.E. Morel, Solution of the Radiation Diffusion Equation on an AMR Eulerian Mesh with Material Interfaces, Tech. Rep. LA-UR-99-2949, Los Alamos National Laboratory, 1999.
[2] D.A. Knoll, W. Rider, G.L. Olson, Nonlinear convergence, accuracy, and time step control in nonequilibrium radiation diffusion, JQSRT 70 (2001) 25–36.
[3] V.A. Mousseau, D.A. Knoll, New physics-based preconditioning of implicit methods for non-equilibrium radiation diffusion, J. Comp. Phys. 190 (2003) 42–51.
[4] D.A. Knoll, D. Keyes, Jacobian-free Newton–Krylov methods: a survey of approaches and applications, J. Comp. Phys. 193 (2004) 357–397.
[5] R.B. Lowrie, A comparison of implicit time integration methods for nonlinear relaxation and diffusion, J. Comp. Phys. 196 (2004) 566–590.
[6] P.N. Brown, D.E. Shumaker, C.S. Woodward, Fully implicit solution of large-scale non-equilibrium radiation diffusion with high order time integration, J. Comp. Phys. 204 (2005) 760–783.
[7] R.M. Rauenzahn, V.A. Mousseau, D.A. Knoll, Temporal accuracy of the nonequilibrium radiation diffusion equations employing a Saha ionization model, Comp. Phys. Commun. 172 (2005) 109–119.
[8] V.A. Mousseau, D.A. Knoll, Temporal accuracy of the nonequilibrium radiation diffusion equations applied to two-dimensional multimaterial simulations, Nuc. Sci. Eng. 154 (2006) 174–189.
[9] G.L. Olson, L.H. Auer, M.L. Hall, Diffusion, P_1, and other approximate forms of radiation transport, JQSRT 64 (2000) 619–634.
[10] J.E. Morel, Diffusion-limit asymptotics of the transport equation, the $P_{1/3}$ equations, and two flux-limited diffusion theories, JQSRT 65 (2000) 769–778.
[11] D. Mihalas, Stellar Atmospheres, second ed., W.H. Freeman and Co., San Francisco, 1978.
[12] D. Mihalas, B.W. Mihalas, Foundations of Radiation Hydrodynamics, Oxford University Press, New York, 1984.
[13] J.H. Ferziger, M. Peric, Computational Methods for Fluid Dynamics, Springer-Verlag, Berlin Heidelberg, 1996.
[14] C.A.J. Fletcher, Computational Techniques for Fluid Dynamics, vol. I, Springer, Berlin, 1991.
[15] K.J. Evans, D.A. Knoll, M. Pernice, Development of a 2-D algorithm to simulate convection and phase transition efficiently, J. Comp. Phys. 219 (2006) 404–417.
[16] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, Numerical Recipes in Fortran, second ed., Cambridge University Press, New York, 1992.