UCRL-JC-133094
PREPRINT

A Block-Structured Adaptive Mesh Refinement Algorithm for Diffusion Radiation

L. H. Howell, J. A. Greenough

This paper was prepared for submittal to the
1998 Nuclear Explosives Development Conference
Las Vegas, NV
October 25-30, 1998

January 27, 1999

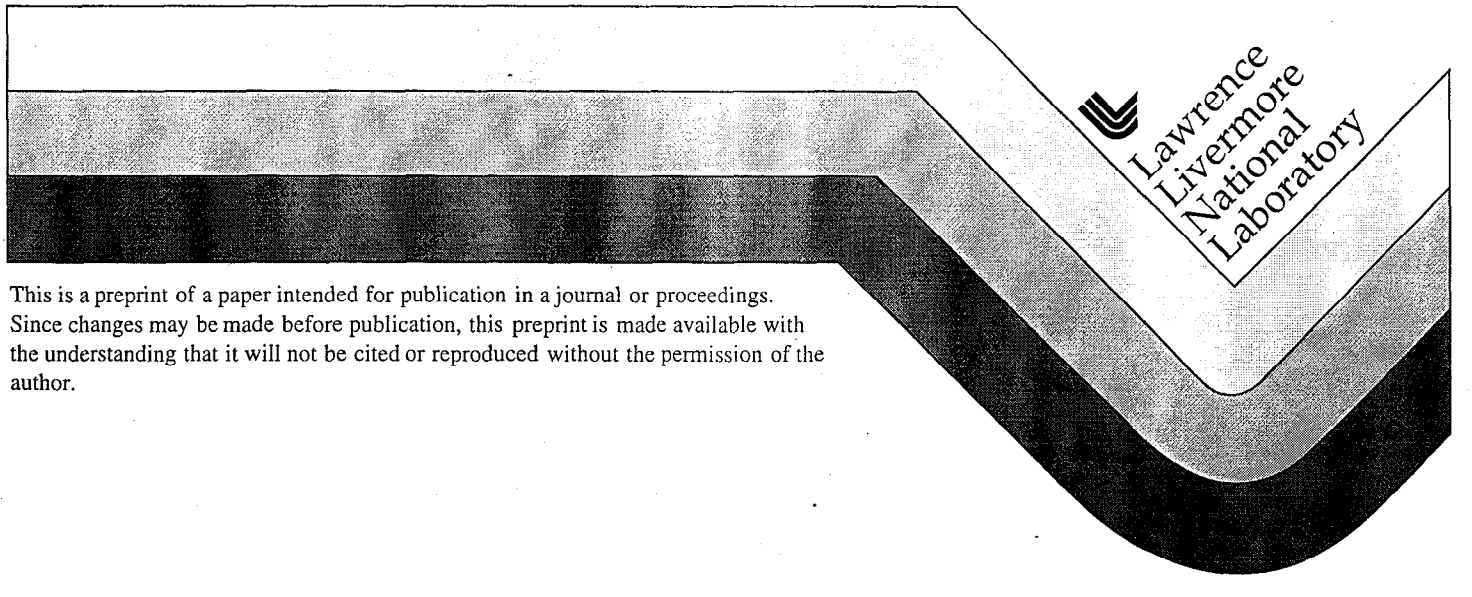

A Block-Structured Adaptive Mesh Refinement Algorithm for Diffusion Radiation

L. H. Howell, J. A. Greenough
Lawrence Livermore National Laboratory

Block-structured meshes provide the ability to concentrate grid points and computational effort in interesting regions of a flow field, without sacrificing the efficiency and low memory requirements of a regular grid. In unsteady calculations each level is advanced at a different timestep, both for greater efficiency and to maintain an optimal Courant condition. Conservation laws are enforced by using locally conservative difference schemes along with explicit synchronization operations between different levels of refinement. We describe a multigrid algorithm for diffusion radiation that we have added to an existing multi-fluid gasdynamics code. Particular attention is given to the appropriate coupling between the fluid energy and the radiation field, and to the form of synchronization between levels required for energy conservation in the diffusion process. (U)

Keywords: radiation transport, diffusion radiation, multi-fluid, multigrid, adaptive mesh

Introduction

The block-structured adaptive mesh refinement (AMR) algorithm, originally developed by Berger and Oliger (1984), was shown to be effective for 2D gasdynamics by Berger and Colella (1989) and was then extended to 3D by Bell et al. (1994). Coupling the explicit integrator for hyperbolic equations with an implicit multigrid solver permits the solution of systems which also involve parabolic and elliptic effects. For example, an algorithm for the incompressible Navier-Stokes equations, presented as a prototype in (Howell and Bell, 1997) and as a fully-developed code in (Almgren et al., 1998), uses multigrid solvers to implicitly model the elliptic incompressibility constraint and the parabolic terms for diffusion. Other examples of implicit processes coupled to explicit advection in AMR codes include self-gravitation in astrophysics (Truelove et al., 1997), and discrete ordinates for radiation in combustion applications (Howell et al., 1998).

In this paper we will present an Eulerian AMR implementation of single-group diffusion radiation, coupled to a Godunov integrator for explicit multi-fluid hydrodynamics. The advective portion of this algorithm is presented in this proceedings by Greenough et al. (1998), along with appropriate citations of previous work from which it was developed.

AMR Essentials

A complete description of the AMR timestepping scheme can be found in many of the references cited in the previous section, but for this paper we will confine ourselves to a very brief overview. The computational mesh consists of a hierarchy of refinement levels, each of which is locally uniform and organized as a union of rectangular patches. Cells on different levels are aligned, and refinement ratios between levels are typically 2 or 4. Edges of different levels may not be coincident except at the physical boundary—that is, level $\ell+2$ may not directly border level ℓ without an intervening region at refinement level $\ell+1$. For

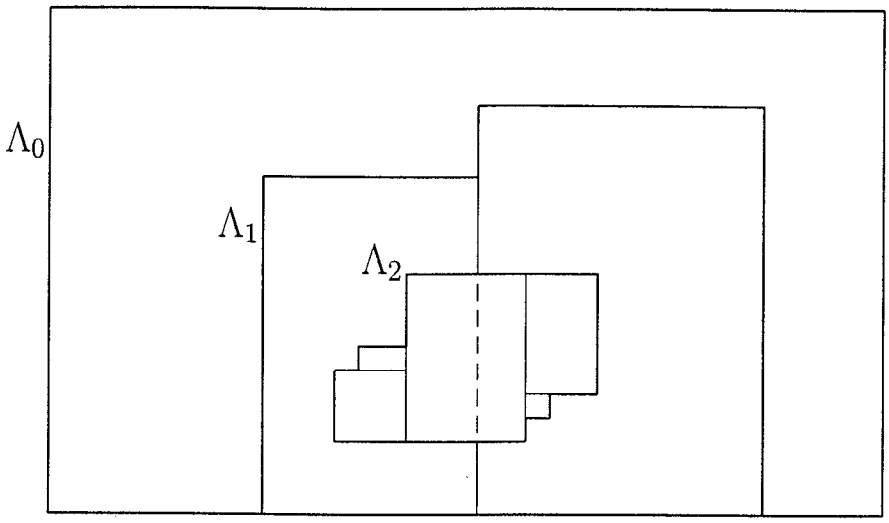

Figure 1: A properly nested hierarchy of grids

computational convenience coarse cells are defined under finer grids, but the finest solution in a given region takes precedence over coarser versions.

The scheme typically refines in time as well as space, both for efficiency and to help maintain an optimal Courant condition on all levels. The refinement ratios in time are the same as those in space—these are specified by the user along with the specific criteria for refinement. Timestepping is a recursive process, where a level ℓ is advanced first without reference to any finer levels, then level $\ell + 1$ is advanced a number of times to catch up (this in turn requires advancing levels $\ell + 2$ and higher). After the finer levels have reached the time corresponding to the end of the coarse timestep, synchronization operations must be performed to correct any flux mismatches that have occurred at the coarse-fine interface. For hyperbolic terms in the equations this is just a local correction, which we refer to as refluxing. For elliptic and parabolic equations a separate call to the implicit solver is required.

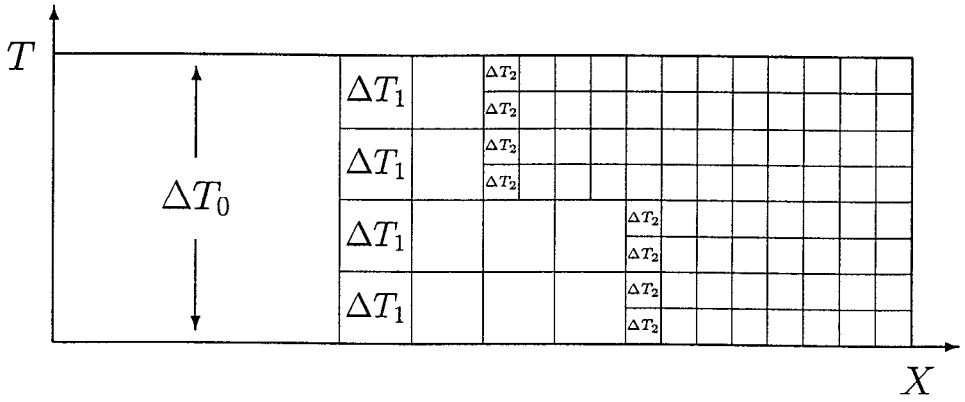

Figure 2: Fine levels advance at smaller timesteps than coarse levels.

Figure 1 shows a possible spatial layout for a grid hierarchy with two refined levels, and Figure 2 shows a space-time outline of the temporal refinement strategy. The latter figure shows the refinement ratios (4 and 2 in this case), and features a regridding event where the boundary between levels 1 and 2 was moved. Note that regridding of finer levels can occur even in the middle of a coarse level timestep.

Diffusion Radiation

Radiative transport is modeled using the diffusion approximation, which satisfies the following equation:

$$\frac{\partial E_R}{\partial t} = \nabla \cdot \left(\frac{c\lambda(E_R)}{\kappa_R} \nabla E_R \right) + \kappa_P \left(4\sigma T^4 - cE_R \right), \tag{1}$$

where E_R is the radiation energy density and κ_R and κ_P are the Rosseland and Planck mean opacities, respectively. The diffusion coefficient $c\lambda(E_R)/\kappa_R$ contains a nonlinear (E_R-dependent) limiter as described by Levermore and Pomraning (1981), the details of which are not a concern here. Note only that $\lambda \leq 1/3$, where $1/3$ is the value without limiting.

We have neglected terms relating to radiation pressure and to motion of the fluid in the Eulerian frame, which are small by a factor of U/c for fluid velocity U. (Future versions may include these terms.) The $\partial E_R/\partial t$ term also lacks a factor of c, but it is retained because without it the radiation field can transport but cannot contain energy. At high temperatures the energy content of the radiation field may be an important part of the total energy in the system.

The solution of (1) is the radiation energy density, defined as

$$E_R = \frac{1}{c} \int_0^\infty \int^{4\pi} I_\nu d\Omega d\nu. \tag{2}$$

This is the zeroth moment of the radiative intensity I. We will also have use for the first moment, the radiative flux vector

$$\mathbf{F} = \int_0^\infty \int^{4\pi} \Omega I_\nu d\Omega d\nu. \tag{3}$$

The diffusion approximation assumes that I is weakly dependent on angle and that therefore $\mathbf{F}$ can be represented as

$$\mathbf{F} = -\frac{c\lambda(E_R)}{\kappa_R} \nabla E_R. \tag{4}$$

This implies that E_R is slowly varying. There is no implication, however, that the radiation field and the fluid are locally coupled.

At vacuum boundaries we use the Marshak boundary condition

$$\hat{n} \cdot \mathbf{F} = \frac{c}{2} E_R - 2F_{\text{inc}}. \tag{5}$$

where F_{inc} is a specified incoming flux. This is a mixed boundary condition involving both E_R and its normal derivative. The simpler Dirichlet and Neumann boundary conditions on the radiation field are also supported.

Energy Update

The radiation field cannot be considered in isolation, as it continually exchanges energy with the fluid medium. The fluid itself obeys the equations of gasdynamics, which in our algorithm are the compressible Euler equations. We use a second-order Godunov method to explicitly advance the density, momentum, and energy of the fluid at each timestep. In addition, we maintain a multi-fluid representation of the state in order to prevent numerical mixing of different materials. The effect of this modification is that the volume fraction f^α, density fraction $f^\alpha\rho^\alpha$, and total energy fraction $f^\alpha\rho^\alpha E^\alpha$ for each fluid α become additional advected quantities. For a description of the multi-fluid algorithm see the paper by Greenough et al. (1998) in this proceedings.

We will start with the single-fluid form of the energy update, and for simplicity of notation will use this form for many of the derivations that follow. Most of the effects of the multi-fluid representation on the radiation algorithm are to require various sums and averages over the fluids in each cell, which must be specified for completeness but add little to the exposition.

The single-fluid energy equation with contributions from advection and radiant energy takes the form

$$\frac{\partial}{\partial t}(\rho E) + \nabla\cdot(\boldsymbol{u}\rho E + \boldsymbol{u}p) = -\kappa_P\left(4\sigma T^4 - cE_R\right). \tag{6}$$

Since our concern is primarily with the radiation terms on the right hand side of this equation, we define $(\rho E)^-$ as the fluid energy density updated to time $n+1$ for the advection process, but not for radiation. That is,

$$(\rho E)^- = (\rho E)^n - \Delta t\left[\nabla\cdot(\boldsymbol{u}\rho E + \boldsymbol{u}p)\right]^{n+1/2}. \tag{7}$$

In each timestep this quantity is computed by the Godunov module before the radiation calculation takes place.

In a single-grid code we could evaluate the radiation terms at time $n+1/2$ in order to do a time-centered update. For AMR this is not appropriate, since multilevel solves should be performed at integral times for which both coarse and fine level data exist. The ideal form of the radiation portion should therefore be

$$-\frac{1}{2}\left[\kappa_P^n\left(4\sigma(T^n)^4 - cE_R^n\right) + \kappa_P^{n+1}\left(4\sigma(T^{n+1})^4 - cE_R^{n+1}\right)\right]. \tag{8}$$

Since time n quantities are known, this leaves us the problem of computing T^{n+1} and E_R^{n+1}.

Implicit Radiation Update

In some regions of the flow field the fluid and radiation temperatures may be tightly coupled. It would therefore be inappropriate to advance T and E_R separately—an implicit form of the update is required which takes the coupling of these quantities into account where it exists.

Let us first define the integrated blackbody intensity $B = 4\sigma T^4$. The update for the fluid energy then takes the form

$$(\rho E)^{n+1} = (\rho E)^- - \Delta t\left\{(1-\theta)\kappa_P^n\left(B^n - cE_R^n\right) + \theta\kappa_P^{n+1}\left(B^{n+1} - cE_R^{n+1}\right)\right\}, \tag{9}$$

where we can set $\theta = 1/2$ for a time-centered update, or $\theta = 1$ for backwards-Euler.

Time centering, with the goal of second-order accuracy in time, seems appropriate for the fluid state updates so long as we can perform them without stability problems. In the diffusion radiation equation, though, the time difference is small (stiff) compared to the other terms by a factor of $1/c$. It therefore seems advisable to difference this term in backwards-Euler fashion. The equation to be satisfied is

$$\frac{E_R^{n+1} - E_R^n}{\Delta t} = \nabla \cdot \left(\frac{c\lambda(E_R^{n+1})}{\kappa_R^{n+1}} \nabla E_R^{n+1} \right) + \kappa_P^{n+1} \left(B^{n+1} - cE_R^{n+1} \right). \tag{10}$$

In the limit of large κ, this formulation has the desirable property of forcing emission and absorption—which are then both large—to balance each other. Their difference must be small enough to be comparable to the other terms in the equation. When the $\partial E_R/\partial t$ term is negligible the only remaining term is the flux, so the only possible solution is a quasi-steady balance among emission, absorption and flux at a single point in time.

There is a problem, however. The quantities $\kappa_P(B - cE_R)$ represent transfer of energy from the fluid to the radiation field. If we include these in the fluid update in Crank-Nicolson form, but in the radiation update in backwards-Euler form, the two will not balance and the system will not conserve energy. If we difference both in Crank-Nicolson form, then the radiation update (10) becomes

$$\frac{1}{2} v^{n+1} + \frac{1}{2} v^n = 0, \tag{11}$$

where $v = \kappa_P(B - cE_R) - \nabla \cdot \mathbf{F}$ and the time derivative is omitted. This supports an oscillatory solution $v^n = a(-1)^n$. A numerical imbalance among emission, absorption and flux at time n is merely inverted, rather than corrected, at time $n + 1$. In the limit of large κ we can have emission and absorption far out of sync, a nonphysical situation.

The alternative we have chosen is to modify the fluid energy update to be of the form

$$(\rho E)^{n+1} = (\rho E)^- - \Delta t \left\{ (1 - \theta)(\nabla \cdot \mathbf{F}^n - \nabla \cdot \mathbf{F}^{n+1}) + \kappa_P^{n+1} \left(B^{n+1} - cE_R^{n+1} \right) \right\}. \tag{12}$$

Since the emission and absorption appear as in a backwards-Euler difference, and the flux terms do not affect conservation on a uniform grid, the combined radiation and fluid energy update is conservative. On the other hand, this form reduces to the true Crank-Nicolson form (9) when $\partial E_R/\partial t$ is small. With the radiation equation permanently in backwards-Euler form this is the only case in which we would expect to see second-order accuracy.

We can perform an iteration to compute all of the appropriate quantities at time $n + 1$, but the parabolic solution itself is for E_R^{n+1}. To emphasize this, we rearrange terms:

$$-\nabla \cdot \left(\frac{c\lambda(E_R^{n+1})}{\kappa_R^{n+1}} \nabla E_R^{n+1} \right) + (\kappa_P^{n+1} + \tau^{n+1}) cE_R^{n+1} = \kappa_P^{n+1} B^{n+1} + \tau^{n+1} cE_R^n, \tag{13}$$

where $\tau^{n+1} = 1/(c\Delta t^{n+1})$ is expected to be small.

Multi-Fluid Modifications, Average Opacities

In the multi-fluid form of the algorithm, each cell may potentially be occupied by more than one material. Each fluid α in a cell takes up a volume fraction f^α, with $\sum_\alpha f^\alpha = 1$. Most quantities derived from the fluid state, including ρ, T and κ_P, are derived separately

$$\kappa_R = e^{\left\{\sum_\alpha \frac{f^\alpha \rho^\alpha}{\rho} \log \kappa_R^\alpha\right\}}$$

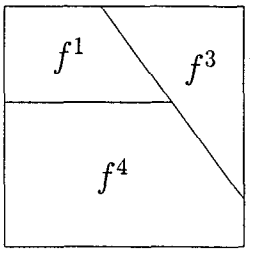

Figure 3: A weighted geometric mean of the Rosseland opacity avoids potential difficulties with anomalous values in small fractions.

for each fluid. The radiation energy density E_R, though, is given a single value over the entire cell. There is thus a separate update equation for each material energy, but a single equation for the radiation:

$$\begin{aligned}(f^\alpha \rho^\alpha E^\alpha)^{n+1} \;=\; & (f^\alpha \rho^\alpha E^\alpha)^- \\ & -\Delta t \left\{ (1-\theta) \frac{f^\alpha \kappa_P^{\alpha,n+1}}{\kappa_P^{n+1}} (\nabla \cdot \mathbf{F}^n - \nabla \cdot \mathbf{F}^{n+1}) \right. \\ & \left. + f^\alpha \kappa_P^{\alpha,n+1} \left(B^{\alpha,n+1} - c E_R^{n+1} \right) \right\},\end{aligned} \tag{14}$$

$$\frac{E_R^{n+1} - E_R^n}{\Delta t} = \nabla \cdot \left(\frac{c\lambda(E_R^{n+1})}{\kappa_R^{n+1}} \nabla E_R^{n+1} \right) + \sum_\alpha f^\alpha \kappa_P^{\alpha,n+1} \left(B^{\alpha,n+1} - c E_R^{n+1} \right). \tag{15}$$

Note that summing the equations (14) over all fluids recovers the single-fluid form of the energy update (12). The average Planck opacity κ_P appearing in (14) is simply $\sum_\alpha f^\alpha \kappa_P^\alpha$, and is chosen to make this sum come out correctly.

The average Rosseland opacity κ_R in the diffusion coefficient is not so trivial to define, however. The issue is more numerical than physical. In cases where a volume fraction for a particular fluid is small, the advected quantities f^α, $f^\alpha \rho^\alpha$ and $f^\alpha \rho^\alpha E^\alpha$ are well-behaved, but the primitive quantities ρ and T on which the opacity depends may become untrustworthy due to cancellation. Since opacity can depend on high powers of these primitive quantities, the value computed in a small fraction may be wildly inaccurate. The geometric mean shown in Figure 3 minimizes the influence of small fractions and thus keeps the cell average value within reasonable limits.

There is another context in which we must average Rosseland opacities, not because of the multi-fluid representation but because of the difference stencil for diffusion. The Rosseland mean appears in the diffusion coefficient, which falls on cell edges rather than cell centers in the standard 5-point discretization. The need to average values onto edges gives us an opportunity to address one of the awkward cases in the diffusion model—that of getting the correct flux at a sharp interface between optically thick and optically thin materials. In the formula shown in Figure 4 the value $4/(3\Delta x)$ is selected for this configuration, producing a reasonable radiative flux across an interface.

$$\kappa_R = \min\left[\frac{\kappa_R^l + \kappa_R^r}{2}, \max\left(\frac{2\kappa_R^l\kappa_R^r}{\kappa_R^l + \kappa_R^r}, \frac{4}{3\Delta x}\right)\right]$$

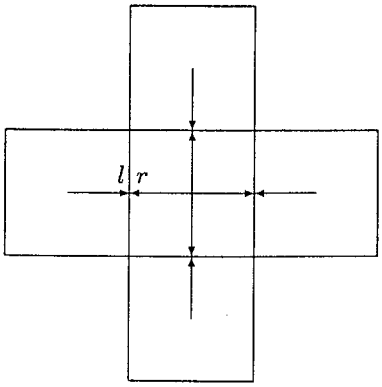

Figure 4: The difference stencil for diffusion uses opacities averaged to edges. The value $4/(3\Delta x)$ gives a reasonable flux across an interface where optically thick and optically thin materials meet.

Energy/Radiation Update Loop

We return now to single-fluid form for simplicity. Solution of equation (13) requires several quantities not yet available, but in particular it requires $B^{n+1} = B^{n+1}(T^{n+1})$. We can estimate this based on a temperature T^* from the partially-updated fluid energy $(\rho E)^-$, perhaps with some explicit guesswork about the expected effect of the radiation field, but whatever we do the estimate will not match the updated temperature to be obtained via equation (12). In general we will have to obtain T^{n+1} through an iterative process. To accelerate convergence we will use a semi-implicit representation for emission:

$$\kappa_P^{n+1}B^{n+1} \approx \kappa_P^* B^* + (T^{n+1} - T^*)\frac{\partial(\kappa_P B)}{\partial T}, \tag{16}$$

where the derivative is evaluated at T^*.

The kinetic portion of ρE is the same on both sides of (12), so we can subtract that off and then divide by ρ^{n+1} to obtain an update for the internal energy e. From ρ^{n+1} and any given e^* we can obtain T^*, c_v and other quantities via the LEOS table (Young and Corey, 1998). We substitute (16) into (12), replace $T^{n+1} - T^*$ with $(e^{n+1} - e^*)/c_v$, and solve for e^{n+1}. The only remaining time $n+1$ quantities are $\mathbf{F}^{n+1}$ and $\kappa_P^{n+1}E_R^{n+1}$, which we will have after we have solved the radiation equation (using κ_P^* computed from T^*). The result is a form for the energy update with the expected change in B built in:

$$e^{n+1} = \eta^* e^* + (1-\eta^*)\left\{e^- - \frac{\Delta t}{\rho^{n+1}}\left[(1-\theta)(\nabla\cdot\mathbf{F}^n - \nabla\cdot\mathbf{F}^{n+1}) + \kappa_P^*\left(B^* - cE_R^{n+1}\right)\right]\right\} \tag{17}$$

where

$$\eta^* = \frac{\dfrac{\partial(\kappa_P B)}{\partial T}}{\dfrac{\rho^{n+1}c_v}{\Delta t} + \dfrac{\partial(\kappa_P B)}{\partial T}}. \tag{18}$$

Now we similarly substitute (16) into (13), replace T again, use (17) to eliminate e^{n+1}, and group related terms. The result is

$$-\left[1-(1-\theta)\eta^*\right]\nabla\cdot\left(\frac{c\lambda(E_R^*)}{\kappa_R^*}\nabla E_R^{n+1}\right)+\left[(1-\eta^*)\kappa_P^*+\tau^{n+1}\right]cE_R^{n+1}=$$

$$(1-\eta^*)\kappa_P^*B^*+\tau^{n+1}cE_R^n+\eta^*\left\{\frac{\rho^{n+1}}{\Delta t}(e^- - e^*)-(1-\theta)\nabla\cdot\mathbf{F}^n\right\}. \quad (19)$$

We now have the pieces necessary to specify the iteration for the implicit update. We first select an initial guess to e^*, where the simplest choice is to take it directly from the advected energy $(\rho E)^-$. We solve (19) for E_R^{n+1}, then update the fluid energy using (17). If $e^{n+1}-e^*$ seems too large we replace e^* with e^{n+1} and do another iteration, continuing until the solution appears to have converged.

In the above analysis the only nonlinear effect that is linearized is emission. The other nonlinear factors—Rosseland mean opacity and flux limiter in the diffusion coefficient, Planck mean opacity in the absorption term—may be updated at each iteration in an attempt to reach a fixed point, but the convergence rate may suffer as a result. On the other hand, holding some of these quantities fixed at their original values may improve the convergence rate, but at the cost of a reduction in time accuracy or even stability due to the lagged coefficients. It does seem advisable to use the same Planck mean for both emission and absorption, so either both should be updated, or both should be lagged (in which case κ_P should be moved out of the derivatives in (18).)

Update Loop in Matrix Form

Since the goal of the previous section gets a bit lost in the algebra, it may be helpful to present the linearization in a somewhat different form. For simplicity we will let $\theta=1$ (backwards-Euler form) and hold the opacities constant. We will work with the multi-fluid form of the equations, though, to show that there are no major complications introduced by doing so. The independent variables for fluid internal energy will therefore be $Q_\alpha=f^\alpha\rho^\alpha e^\alpha$.

We express the implicit update in the form of Newton's method. The coupled equations are written as dependent variables that should all be zero when the system is converged:

$$F_\alpha = (f^\alpha\rho^\alpha e^\alpha)^{n+1}-(f^\alpha\rho^\alpha e^\alpha)^- +\Delta t f^\alpha\kappa_P^\alpha\left(B^{\alpha,n+1}-cE_R^{n+1}\right), \quad (20)$$

$$F_R = E_R^{n+1}-E_R^n-\Delta t\nabla\cdot(D\nabla E_R^{n+1})-\Delta t\sum_\alpha f^\alpha\kappa_P^\alpha\left(B^{\alpha,n+1}-cE_R^{n+1}\right). \quad (21)$$

Partial derivatives of these dependent variables give us the form of the Jacobian matrix. We then have a linear system to solve for updates to the radiation and fluid energies. Note that for each cell there are multiple fluid energies Q_α and multiple fluid energy equations represented by the F_α, one for each fluid.

$$\frac{\partial F_\alpha}{\partial Q_\alpha} = 1+\Delta t\frac{f^\alpha\kappa_P^\alpha}{f^\alpha\rho^\alpha c_v^\alpha}\frac{\partial B^\alpha}{\partial T^\alpha} \quad (22)$$

$$\frac{\partial F_\alpha}{\partial E_R} = -\Delta t f^\alpha\kappa_P^\alpha c \quad (23)$$

$$\frac{\partial F_R}{\partial Q_\alpha} = -\Delta t \frac{f^\alpha \kappa_P^\alpha}{f^\alpha \rho^\alpha c_v^\alpha} \frac{\partial B^\alpha}{\partial T^\alpha} \tag{24}$$

$$\frac{\partial F_R}{\partial E_R} = 1 - \Delta t \nabla \cdot (D \nabla *) + \Delta t \kappa_P c \tag{25}$$

$$\begin{bmatrix} \frac{\partial F_\alpha}{\partial Q_\alpha} & \frac{\partial F_\alpha}{\partial E_R} \\ \frac{\partial F_R}{\partial Q_\alpha} & \frac{\partial F_R}{\partial E_R} \end{bmatrix} \begin{bmatrix} \delta Q_\alpha \\ \delta E_R \end{bmatrix} = \begin{bmatrix} -F_\alpha \\ -F_R \end{bmatrix} \tag{26}$$

The upper-left block of the Jacobian matrix is diagonal and therefore trivial to invert. We can therefore solve the system by forming the Schur complement:

$$\begin{bmatrix} \frac{\partial F_\alpha}{\partial Q_\alpha} & \frac{\partial F_\alpha}{\partial E_R} \\ 0 & \frac{\partial F_R}{\partial E_R} - \frac{\partial F_R}{\partial Q_\alpha} \left(\frac{\partial F_\alpha}{\partial Q_\alpha} \right)^{-1} \frac{\partial F_\alpha}{\partial E_R} \end{bmatrix} \begin{bmatrix} \delta Q_\alpha \\ \delta E_R \end{bmatrix} = \begin{bmatrix} -F_\alpha \\ -F_R + \frac{\partial F_R}{\partial Q_\alpha} \left(\frac{\partial F_\alpha}{\partial Q_\alpha} \right)^{-1} F_\alpha \end{bmatrix} \tag{27}$$

We emphasize that the two forms of the implicit update presented here and in the previous section are essentially equivalent, given only that the same nonlinear quantities are linearized in each. In both cases the right hand side and the diagonal of the diffusion operator are modified by an expression involving

$$\eta^* = -\frac{\partial F_R}{\partial Q_\alpha} \left(\frac{\partial F_\alpha}{\partial Q_\alpha} \right)^{-1}. \tag{28}$$

The matrix form, though, makes it more obvious how to extend the linearization to the flux limiter and to other nonlinear quantities—just include the appropriate partial derivatives in the Jacobian matrix (sometimes easier said than done). It also provides a more easily-extensible framework to which we can add additional physics: electron heat conduction, for example, or multigroup radiation.

Stability of the Update

There are two separate questions we may ask about the update iteration (17–19): "Does the iteration itself converge?" and "Given convergence at each timestep, what behavior do we expect over multiple timesteps?" To address the first, we note first that since the iteration is essentially Newton's method, we can expect the usual behavior of Newton iterations. The method should converge rapidly near the solution, but this convergence may degrade if the Jacobian is inaccurate (e.g., has missing terms). If started far from the solution, in the sense that nonlinearities become important, the method may diverge. This suggests a timestep limitation based on a maximum permissible change in temperature.

A limiting configuration provides insight into both questions. If we neglect τ, so that the radiation field acts as a transport mechanism but not as an energy reservoir, we can solve (19) for $(1-\eta^*)\kappa_P^*(B^* - cE_R^{n+1})$. Substituting the result into (17) and simplifying then gives

$$e^{n+1} = e^- - \frac{\Delta t}{\rho^{n+1}} \left[(1-\theta) \nabla \cdot \mathbf{F}^n + \theta \nabla \cdot \mathbf{F}^{n+1} \right]. \tag{29}$$

Note that all of the starred quantities representing the previous iteration have disappeared. This is encouraging, since it means that in this limit the accuracy of the energy update depends only on the accuracy of the current estimate to the radiative flux, not on some complex combination of B^*, η^*, etc.

In the limit of close coupling between the radiation and the fluid, this update becomes a simple Crank-Nicolson update for radiation heat conduction:

$$\rho^{n+1}c_v^{n+1}\frac{T^{n+1}-T^-}{\Delta t} = (1-\theta)\nabla\cdot\left(\frac{16\sigma T^3}{3\kappa_R}\nabla T^n\right) + \theta\nabla\cdot\left(\frac{16\sigma T^3}{3\kappa_R}\nabla T^{n+1}\right). \tag{30}$$

The usual stability considerations for a Crank-Nicolson update then apply. When the timestep is too large this discretization can be vulnerable to oscillations, but these can be avoided (with some loss of accuracy) by running the scheme in backwards-Euler mode.

AMR Time Step

Up to this point we have been discussing operations confined to a single level of refinement. In the AMR timestepping scheme each level advances at its own local timestep, obtaining boundary data from the next coarser level or from the physical boundary as appropriate. The pattern of rectangular grids within a level should make no difference to the computed solution, since a different division of a refined region into patches does not alter the discretization.

The real issue comes at the end of a coarse timestep, when coarse and fine levels have advanced to equivalent points in time. Fluxes across the coarse-fine interface have been computed separately by the two levels and therefore do not match exactly. In the explicit advection scheme this mismatch can be remedied by a local update to the coarse cells adjacent to the interface. In the parabolic radiation update, however, information from a localized region can potentially affect the entire domain, so an additional call to the linear solver is required.

Unlike the equations solved in advancing a level, the synchronization operation directly impacts at least two levels simultaneously. We therefore need a multilevel solver which computes a solution on two or more levels at a time, in the style of Almgren et al. (1994) or McCormick (1989). Fortunately, though, the corrections are generally small, so we require only a single linear solution, not an iterative nonlinear solution as we do for the level advance operations.

The form of synchronization required for radiative energy conservation in an AMR time step is similar to that developed by Howell et al. (1998) for the AMR discrete ordinates code, with the added complication that in the present algorithm the radiation field stores energy and is not merely a transport mechanism. The algorithm in the discrete ordinates paper also included correction terms designed to improve the accuracy of the coarse grid solution by anticipating the effects of the finer levels. We have experimented with these corrections in the present diffusion algorithm, but they are not included in the following description because we have not found them to be beneficial.

The AMR time step is defined recursively, in terms of operations on a level ℓ and its interactions with coarser and finer levels. We are advancing level ℓ from time index n to $n+1$, corresponding to time values $t^{\mathrm{old},\ell}$ and $t^{\mathrm{new},\ell}$, respectively. The region covered by

level ℓ is denoted Λ^ℓ, its border is $\partial\Lambda^\ell$, and the border of the next higher level, projected onto level ℓ, is $\mathbf{P}(\partial\Lambda^{\ell+1})$.

If $(\ell < \ell_{\max})$ or $t^{\text{old},\ell} = t^{\text{old},\ell-1}$) **then**

- Store $\mathbf{F}^{\ell,n}$ in flux registers on $\partial\Lambda^\ell$ and $\mathbf{P}(\partial\Lambda^{\ell+1})$, and compute $\nabla\cdot\mathbf{F}^{\ell,n}$ on Λ^ℓ.

Endif

Level Time Step, level ℓ:

Multi-Fluid Advection (not described here)

Implicit Energy Update

A check indicates quantities that can be updated during the iteration:

- $\dfrac{\check{E}_R^{\ell,n+1} - E_R^{\ell,n}}{\Delta t^\ell} = \nabla\cdot\left(\dfrac{c\lambda(\check{E}_R^{\ell,n+1})}{\check{\kappa}_R^{\ell,n+1}}\nabla\check{E}_R^{\ell,n+1}\right) + \check{\kappa}_P^{\ell,n+1}\left(\check{B}^{\ell,n+1} - c\check{E}_R^{\ell,n+1}\right)$
- $(\rho\check{E})^{\ell,n+1} = (\rho E)^{\ell,-} - \Delta t^\ell\left\{(1-\theta)(\nabla\cdot\mathbf{F}^{\ell,n} - \nabla\cdot\check{\mathbf{F}}^{\ell,n+1}) + \check{\kappa}_P^{\ell,n+1}\left(\check{B}^{\ell,n+1} - c\check{E}_R^{\ell,n+1}\right)\right\}$

End

- $\bar{\mathbf{F}}^{\ell,n+\theta} := (1-\theta)\mathbf{F}^{\ell,n} + \theta\mathbf{F}^{\ell,n+1}$
- Advance levels $\ell+1,\ldots,\ell_{\max}$
- $\delta\bar{\mathbf{F}}^{\ell+1,n+\theta} := \langle\langle\bar{\mathbf{F}}^{\ell+1}\rangle\rangle - \bar{\mathbf{F}}^{\ell,n+\theta}$ on $\mathbf{P}(\partial\Lambda^{\ell+1})$, $\ell < \ell_{\max}$

End

If $(\ell < \ell_{\max})$ **then**

Synchronization/refluxing between levels ℓ and $\ell+1$:

Multilevel Solve: $\ell' \in \{\ell,\ldots,\ell_{\max}\}$

Solve is for $\tilde{E}_R^{\ell',n+1}$ and $\tilde{\mathbf{F}}^{\ell',n+1}$

- $\dfrac{\tilde{E}_R^{\ell',n+1}}{\Delta t^\ell} = \nabla\cdot\left(\dfrac{c\lambda(E_R)}{\kappa_R^{\ell',n+1}}\nabla\tilde{E}_R^{\ell',n+1}\right) - \kappa_P^{\ell',n+1}c\tilde{E}_R^{\ell',n+1} - D_R(\delta\bar{\mathbf{F}}^{\ell'+1,n+\theta})$ on $\Lambda^{\ell'} - \mathbf{P}(\Lambda^{\ell'+1})$
- $\tilde{\mathbf{F}}^{\ell',n+1} = \langle\tilde{\mathbf{F}}^{\ell'+1,n+1}\rangle$ on $\mathbf{P}(\partial\Lambda^{\ell'+1})$, $\ell' < \ell_{\max}$

End

For $\ell' \in \{\ell_{\max},\ldots,\ell\}$ **do**

- $(\rho E)^{\ell',n+1} := (\rho E)^{\ell',n+1} + \Delta t^\ell\kappa_P^{\ell',n+1}c\tilde{E}_R^{\ell',n+1}$ on $\Lambda^{\ell'} - \mathbf{P}(\Lambda^{\ell'+1})$
- $E_R^{\ell',n+1} := E_R^{\ell',n+1} + \tilde{E}_R^{\ell',n+1}$ on $\Lambda^{\ell'} - \mathbf{P}(\Lambda^{\ell'+1})$
- $(\rho E)^{\ell',n+1} := \langle(\rho E)^{\ell'+1,n+1}\rangle$, $E_R^{\ell',n+1} := \langle E_R^{\ell'+1,n+1}\rangle$ on $\mathbf{P}(\Lambda^{\ell'+1})$, $\ell' < \ell_{\max}$

Enddo

- $\bar{\mathbf{F}}^{\ell,n+\theta} := \bar{\mathbf{F}}^{\ell,n+\theta} + \tilde{\mathbf{F}}^{\ell,n+1}$ on $\mathbf{P}(\partial\Lambda^\ell)$, $\ell > 0$

Endif

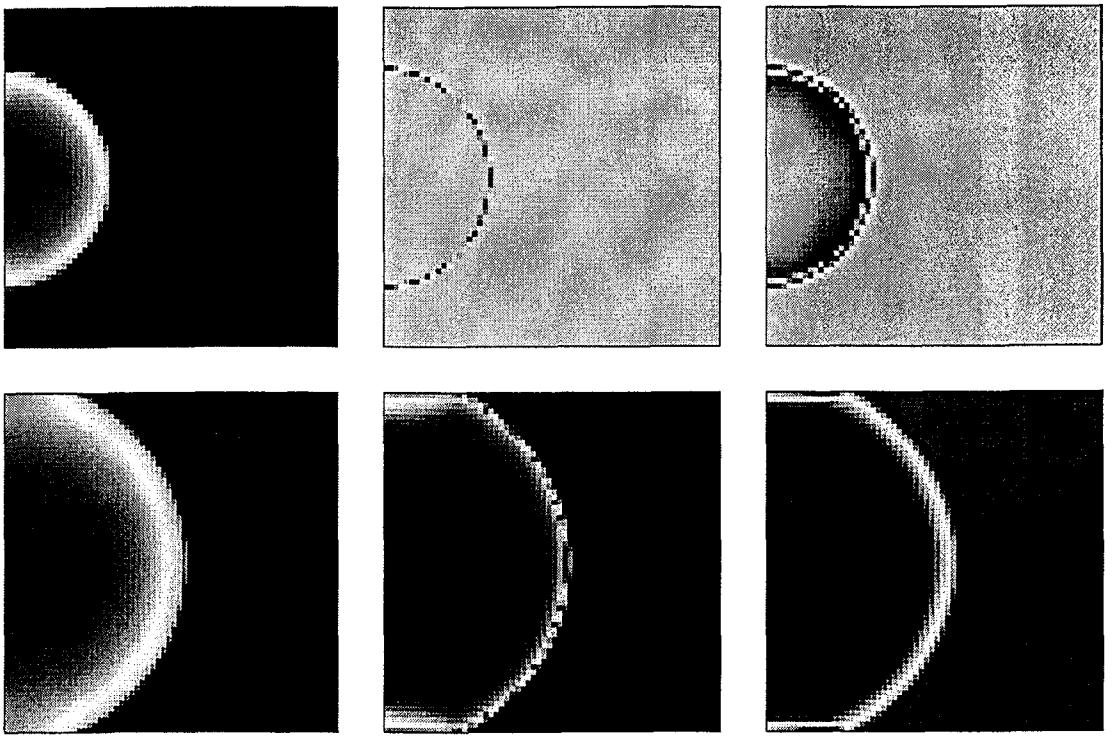

Figure 5: Nonlinear thermal wave expanding from a point heat source in RZ coordinates, shown at two different times. Left: Temperature. Center: Difference between computed temperature and analytic solution. Right: Density. At the early time the density range is (0.999–1.002), indicating that the thermal wave is dominant. At the later time the density range is (0.883–1.687), as hydrodynamic effects are becoming important.

The notation $\langle \cdot \rangle$ indicates an average in space, while $\langle\langle \cdot \rangle\rangle$ denotes an average in both space and time. The function D_R is the reflux divergence, which is nonzero only in coarse cells bordering an interface with a finer level.

Numerical Results

We first show results involving nonlinear thermal waves, since these demonstrate the behavior of the method in a radiation-dominated regime. The sharp front of a nonlinear wave provides a convenient feature to use as a refinement criterion, but we will start with results computed on a single grid.

Zel'dovich and Raizer (1967) give analytic solutions for a pure thermal wave. In our calculations we also model hydrodynamic effects, which become important at later times as the thermal wave weakens and transforms into a shock wave. The pure thermal solutions include the cases of expansion from an initial plane source and from a point source. In our code we have computed the first of these using 2D Cartesian coordinates, and the second using both cylindrical and spherical coordinates. All solutions agree well with the analytic result, but for the sake of brevity we will show only the cylindrical (RZ) case here.

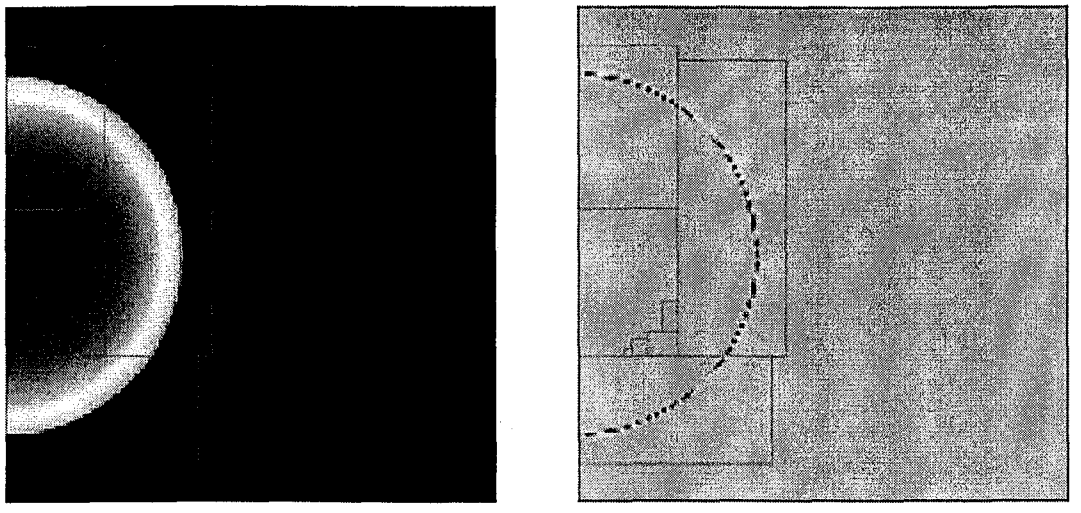

Figure 6: Nonlinear thermal wave expanding from a point heat source in RZ coordinates, with adaptive mesh refinement. Left: Temperature. Right: Difference between computed temperature and analytic solution.

Figure 5 shows the wave expansion at two different times. At early time the thermal wave is much stronger than any hydrodynamic effects. Because of this the density variations are still negligible, and the difference with the analytic solution consists primarily of discretization effects at the interface. At later time the thermal wave has weakened and a shock is forming. Note the broader front in the difference with the analytic solution, which does not include hydrodynamic effects. (Also, the wave has hit the walls of the domain, which are not included in the analytic model either.)

Next, Figure 6 shows a similar calculation using adaptive mesh refinement. The sharp gradient at the wave front is used as the criterion for adding grids a factor of 2 finer than the base grid. As the wave has expanded the grids have changed, so that the front is now in a region which was not initially refined. Still, the difference with the analytic solution is confined to the discretization of the sharp front itself. Note the absence of errors associated with the interfaces between refinement levels.

Finally, we show a sample calculation using real materials, with equation of state and opacity information obtained from the LEOS table (Young and Corey, 1998). The configuration is based on a NOVA experiment designed by D. Ward and T. Peyser, in which two half-hohlraums generate shocks moving in opposing directions. One of the shocks accelerates a sinusoidally perturbed interface. The perturbations grow due to a Richtmeyer-Meshkov instability, then interact with the second shock in a way which should lead to a rapid transition to turbulence.

We performed this calculation in RZ geometry using a 56×96 base grid and two levels of refinement by factors of 2. The five materials were carbon foam (inside tube), brominated polystyrene (ends of tube), beryllium (tube wall), air (outside) and silver (mounting), in the configuration shown on the left in Figure 7. Refinement was based on the locations of the material interfaces. The drive was in the form of a 2 ns square pulse, with a peak

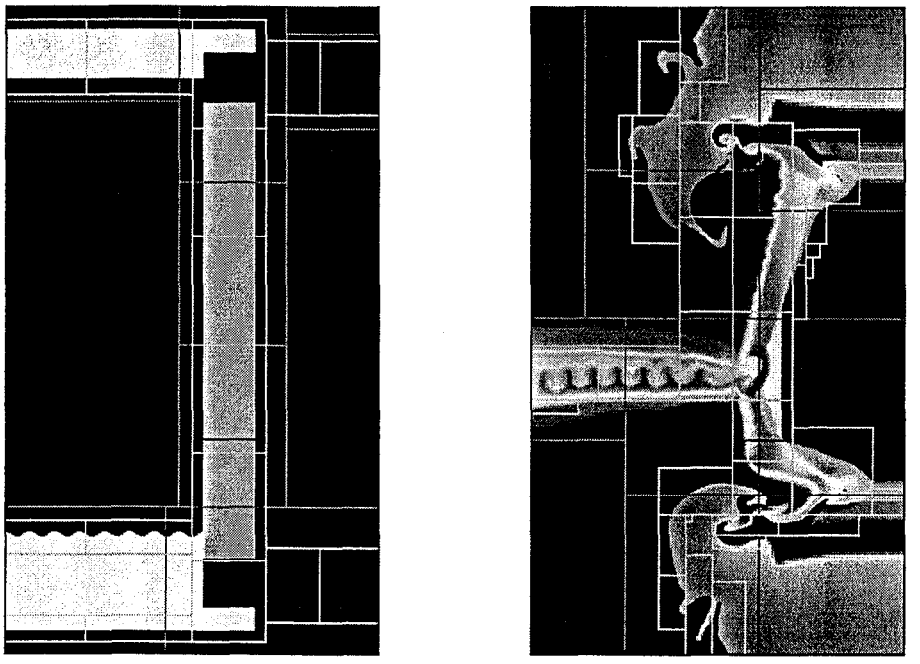

Figure 7: Left: Initial configuration in RZ coordinates with axis at left. Right: Radiation drive has created shocks from above and below, which meet in the middle of the tube as shown here. Both images are based on material density.

temperature of 200 eV. The right side of Figure 7 shows the configuration at 27 ns, when the two shocks have driven together in the center of the tube.

Acknowledgment

This work was performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under contract No. W-7405-ENG-48.

References

Almgren, A. S., Bell, J. B., Colella, P., Howell, L. H., and Welcome, M. L., "A Conservative Adaptive Projection Method for the Variable Density Incompressible Navier-Stokes Equations," *J. Comput. Phys.*, **142**, 1–46 (1998).

Almgren, A. S., Buttke, T., and Colella, P., "A Fast Adaptive Vortex Method in Three Dimensions," *J. Comput. Phys.*, **113**, 177–200 (1994).

Bell, J., Berger, M., Saltzman, J., and Welcome, M., "Three-Dimensional Adaptive Mesh Refinement for Hyperbolic Conservation Laws," *SIAM J. Sci. Comput.*, **15**, 127–138 (1994).

Berger, M. J., and Colella, P., "Local Adaptive Mesh Refinement for Shock Hydrodynamics," *J. Comput. Phys.*, **82**, 64–84 (1989).

Berger, M. J., and Oliger, J., "Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations," *J. Comput. Phys.*, **53**, 484–512 (1984).

Greenough, J. A., Graves, D. T., and Colella, P., "A Second-Order Method for Interface Reconstruction in Orthogonal Coordinate Systems," *Proceedings of the Nuclear Explosives Code Development Conference,* (Las Vegas, October 25–30, 1998).

Howell, L. H., and Bell, J. B., "An Adaptive Mesh Projection Method for Viscous Incompressible Flow," *SIAM J. Sci. Comput.,* **18**, 996–1013 (1997).

Howell, L. H., Pember, R. B., Colella, P., Jessee, J. P., and Fiveland, W. A., "A Conservative Adaptive Mesh Algorithm for Unsteady, Combined-Mode Heat Transfer Using the Discrete Ordinates Method," Lawrence Berkeley National Laboratory, Berkeley, CA, LBNL-41311 (1998).

Levermore, C. D., and Pomraning, G. C., "A Flux-Limited Diffusion Theory," *Astrophys. J.,* **248**, 321–334 (1981).

McCormick, S. F., *Multilevel Adaptive Methods for Partial Differential Equations,* (SIAM, Philadelphia, 1989).

Truelove, J. K., Klein, R. I., McKee, C. F., Holliman, J. H. II, Howell, L. H., and Greenough, J. A., "The Jeans Condition: A New Constraint on Spatial Resolution in Simulations of Isothermal Self-Gravitational Hydrodynamics," *Astrophys. J.,* **489**, L179–L183 (1997).

Young, D., and Corey, E., "Release of New LEOS Equation of State Library," Lawrence Livermore National Laboratory, Livermore, CA, unpublished memo (April 1, 1998).

Zel'dovich, Ya. B., and Raizer, Yu. P., *Physics of Shock Waves and High-Temperature Hydrodynamic Phenomena, Vol. 2,* (Academic Press, New York, 1967), pp. 664–672.