

SOLUTION OF EQUILIBRIUM RADIATION DIFFUSION PROBLEMS USING IMPLICIT ADAPTIVE MESH REFINEMENT*

MICHAEL PERNICE[†] AND BOBBY PHILIP[†]

Abstract. Diffusion approximations to radiation transport feature a nonlinear conduction coefficient that leads to formation of a sharp thermal front, or Marshak wave, under suitable initial and boundary conditions. The thermal front can vary several orders of magnitude over a very short distance. Resolving the shape of the thermal front is essential, but using a global fine mesh can be prohibitively expensive. In such circumstances it is natural to consider using adaptive mesh refinement (AMR) to place a fine mesh only in the vicinity of the propagating front. In addition, to avoid any loss of accuracy due to linearization, implicit time integration should be used to solve the equilibrium radiation diffusion equation. Implicit time integration on AMR grids introduces a new challenge, as algorithmic complexity must be controlled to fully realize the performance benefits of AMR. A Newton-Krylov method together with a multigrid preconditioner addresses this latter issue on a uniform grid. A straightforward generalization is to use a multilevel preconditioner that is tuned to the structure of the AMR grid, such as the Fast Adaptive Composite grid (FAC) method. We describe the resulting Newton-Krylov-FAC method and demonstrate its performance on simple equilibrium radiation diffusion problems.

Key words. Adaptive mesh refinement, implicit time integration, Newton-Krylov methods, FAC, multilevel preconditioning, radiation diffusion

AMS subject classifications. 65H10, 65F10, 65N50, 65N55

1. Introduction. Radiation transport plays an important role in numerous fields of study, including astrophysics, laser fusion, and combustion applications such as modeling of coal-fired power generation systems and wildfire spread. A diffusion approximation provides a reasonably accurate description of penetration of radiation from a hot source to a cold medium. This approximation features a nonlinear conduction coefficient that leads to formation of a sharp thermal front, in which the solution can vary several orders of magnitude over a very short distance. The shape of the thermal front can be very complex as it interacts with different materials having different conduction properties. Resolving these localized features with a global fine mesh can be prohibitively expensive. It is natural to consider reducing the cost of accurately resolving these thermal fronts by using adaptive mesh refinement (AMR), which concentrates computational effort by increasing spatial resolution only locally.

Classical solution techniques for equilibrium radiation diffusion use a linearized conduction coefficient to avoid the expense of solving a system of nonlinear equations at each time step. This introduces a first order error in time, precluding effective use of higher order time integration methods, and small time steps must be used to control the size of this error. Analytic and computational results that demonstrate degradation in time accuracy associated with linearization in the presence of strong nonlinear coefficients can be found in [22] and [13]. Such effects can be avoided by using implicit time integration, which generally requires efficient nonlinear solution techniques to be competitive. Newton-Krylov methods, usually preconditioned by a multigrid method, have been instrumental in demonstrating that this approach is practical for equilibrium radiation diffusion problems [24].

*This work was supported under the auspices of the U.S. Department of Energy under DOE contract W-7405-ENG-36 at Los Alamos National Laboratory.

[†] Computer and Computational Sciences Division, Los Alamos National Laboratory, Los Alamos, NM 87545. {pernice,bphilip}@lanl.gov

We seek to combine the benefits of improved temporal accuracy of implicit time integration with the improved spatial accuracy that is made possible by AMR. Two important factors must be addressed in order to fully realize the potential efficiency gains. The first is the use of finer grids only in those regions where high resolution is required. This can be accomplished by determining regions where the spatial error is large or where localized solution features warrant enhanced resolution, and immediately translates into lower overall storage costs. The second is the use of algorithms whose arithmetic complexity scales linearly with problem size. This presents a significant challenge when using implicit time integration methods.

Prior efforts in implicit AMR have not adequately addressed the algorithmic scalability issue. Early work by Winkler, Norman, and Mihalas [28] uses r -refinement to solve coupled radiation-hydrodynamics problems. They use a fully coupled formulation, solved by a classical Newton’s method, that includes a functional to determine new gridpoint locations for following features of the flow. Their approach is impractical for 2- or 3-dimensional problems, and even in one dimension difficulties are encountered with gridpoint movement. We use h -refinement, in particular structured AMR (SAMR), to avoid these problems. Both Trompert and Verwer [26, 27] and Li, Petzold, and Hyman [15] use incomplete factorization to precondition their linear solvers. Incomplete factorization can be expensive to set up, and, depending on the amount of fill-in allowed, can have high storage overhead. Further, incomplete factorization scales poorly with increasing resolution, and parallel implementation is problematic. In addition, this approach requires formation of the full Jacobian, through either analytic or automatic means. By using a Jacobian-free Newton-Krylov method, we eliminate the burden of supplying and storing the Jacobian.

We address the issue of algorithmic scalability by using the Fast Adaptive Composite grid (FAC) method of McCormick and Thomas [17, 18] to precondition the systems of linear equations that must be solved in every nonlinear iteration at every time step. FAC has low setup costs, low storage overhead, and converges at a rate independent of the number of refinement levels [18]. Similar methods for elasto-plasticity [4] and equilibrium radiation diffusion [25] use multilevel methods on unstructured grids to solve the linearized equations, but still require formation of the Jacobian. More recently, Howell and Greenough [12] solve a *linearized* radiation diffusion problem on a SAMR grid as part of a semi-implicit radiation-hydrodynamics algorithm, but performance of the multilevel linear solver is not their primary focus, and no details about its performance are provided. With minor modifications, our approach could be used to implement a fully nonlinear version of their reflux step.

We report on efforts to solve equilibrium radiation diffusion problems using structured AMR and the Newton-Krylov-FAC method. While structured AMR facilitates reuse of existing software written for logically rectangular grids, discretization at locations near changes in resolution must be treated carefully in order to avoid the creation of artificial sources.

This paper is organized as follows. The next section discusses the equilibrium radiation diffusion model. A discussion of structured adaptive mesh refinement follows in §3. Special considerations for spatial discretization on SAMR grids appear in §4. We describe our algorithmic components in §5. Results of computations appear in §6 and we summarize our conclusions in §7.

2. Equilibrium Radiation Diffusion. We provide a brief description of the equations that govern equilibrium radiation diffusion. See [19] for a more detailed discussion.

Propagation of a radiation field and its interaction with matter can be modeled by an integro-differential equation that accounts for transport, emission, absorption, and scattering of photons. When the radiation field is isotropic, detailed treatment of transport in angle is not needed. The dependence on angle can be averaged out to obtain a description based on spectral energy density and flux. In a static medium at local thermal equilibrium, absorption is independent of frequency, so dependence on frequency can also be averaged out. In this case, the radiative flux can be shown to be proportional to the gradient of the energy density, with the radiative conductivity proportional to the inverse of the opacity of the medium. Also, when the medium is in thermal equilibrium, emission equals absorption, and the total radiation energy density E is proportional to T^4 , where T is the temperature of the medium. These considerations lead to a parabolic partial differential equation

$$(2.1) \quad \frac{\partial E}{\partial t} = \nabla \cdot D(E) \nabla E \quad \text{in } \Omega \subset \mathbb{R}^d, \quad d \in \{1, 2, 3\}, \quad t > 0,$$

with initial conditions $E(t = 0) = E_0$. We use a model of radiative conductivity in which $D(E)$ scales as

$$(2.2) \quad D(E) = Z^\alpha E^\beta, \quad \alpha < 0, \quad \beta \in [0, 1],$$

where Z is the atomic number of the medium. The parameters α and β can be determined experimentally or through simulation [20]. In this work we take $\alpha = -3$ and $\beta = 0.75$.

Because (2.1) is derived from asymptotic arguments, its solution can exhibit non-physical behavior, such as propagation of radiative energy at a rate faster than the speed of light. In particular, diffusion theory overestimates the amount of energy deposited into a cold medium from a hot source. This can be handled by flux limiting, which is an ad hoc adjustment of $D(E)$ that produces the correct asymptotic behavior. We use Wilson's form for flux limiting [19, 29]

$$(2.3) \quad D_L = \left(\frac{1}{D(E)} + \frac{\|\nabla E\|}{|E|} \right)^{-1}.$$

For the remainder of this paper we will drop the subscript and always use flux-limited diffusion defined by the combination of (2.2) and (2.3).

Finally, boundary conditions for (2.1) must be prescribed. Penetration of heat into a cold medium at rest from a hot source can be modeled by assuming a constant imposed radiation field on a portion of the physical boundary $\partial\Omega_{\mathcal{R}} \subset \partial\Omega$. We also assume that the remainder of the physical boundary $\partial\Omega_{\mathcal{N}} = \partial\Omega - \partial\Omega_{\mathcal{R}}$ is perfectly insulating. Under the same assumptions that led to (2.1), the boundary conditions may be expressed as

$$\begin{aligned} \mathbf{n} \cdot D(E) \nabla E + E/2 &= R \quad \text{on } \partial\Omega_{\mathcal{R}}, \quad t > 0, \\ \mathbf{n} \cdot D(E) \nabla E &= 0 \quad \text{on } \partial\Omega_{\mathcal{N}}, \quad t > 0, \end{aligned}$$

where \mathbf{n} is the unit outward normal to $\partial\Omega$. These conditions lead to solutions of (2.1) known as Marshak waves, which were first described in [16].

3. Structured Adaptive Mesh Refinement. Localized sharp propagation fronts and material discontinuities make the radiation diffusion problem an excellent candidate for AMR. h -refinement AMR techniques provide local mesh resolution by refining the computational mesh locally. Structured adaptive mesh refinement (SAMR)




FIG. 3.1. A composite grid Ω^c of two levels and its component grids Ω_1^h and Ω_2^h .

is h -refinement with local fine grid patches placed over coarser grid patches providing increased local resolution. SAMR techniques enjoy several advantages. Regular array access patterns improve cache performance, simple data structures for bookkeeping minimize the overhead for computations, uniform stencil discretizations provide increased accuracy, and reuse of software developed for single grids is possible.

A SAMR grid consists of a collection of grid patches at different grid resolutions which together cover the computational domain. Grids with the same mesh resolution are grouped together and form a refinement level. Grids at a refinement level are disjoint or touch along cell boundaries, but do not overlap. A fine grid patch overlying a coarse grid patch is called a child grid with the underlying coarse grid being the parent grid. A child grid can have several parent grids and vice versa. The boundaries of child and parent grids are assumed not to align except possibly at physical domain boundaries. This leads to a natural hierarchical structure for SAMR grids that is exploited while designing multilevel algorithms. Operations on the composite grid are decomposed into operations on individual refinement levels which in turn further decompose into operations on individual patches. Valid degrees of freedom at level k are defined in cells not covered by cells belonging to level $k + 1$. Cells covered by grids at finer levels are 'slave' cells with their values being derived from fine cells at the next finer level.

Notation is now introduced for a SAMR grid. Ω_J^c represents a composite SAMR grid with J refinement levels. The subscript is dropped when there is no ambiguity. Refinement level k , $k = 1, 2, \dots, J$, denoted by $\Omega_k^{h_k}$ or more simply as Ω_k^h , consists of a collection $\{\mathcal{P}\}$ of grid patches at the same grid resolution h_k . The refinement levels are ordered by increasing grid resolution with $h_{k+1} = r * h_k$, $k = 1, 2, \dots, J - 1$, where $r \in \mathbb{N}^d$ is the refinement ratio. Each component of r is fixed at 2 in this application, typical values being 2, 3, and 4. The subdomain(s) covered by Ω_{k+1}^h fully nest within the subdomain(s) covered by Ω_k^h . Figure 3.1 illustrates this multilevel structure; Figures 6.2 and 6.5 in §6 provide more examples of multilevel grid configurations.

4. Discretization. A method of lines approach is used in which (2.1) is first discretized in time. Let $t_{n+1} = t_n + \Delta t_n$, $n = 0, 1, \dots$, with $t_0 = 0$. Δt_n is a variable time step determined adaptively during the course of the simulation. At $t = t_{n+1}$ an equation of the form

$$(4.1) \quad E^{n+1} - \gamma \Delta t_n \nabla \cdot D(E^{n+1}) \nabla E^{n+1} - g_n = 0$$

is solved, where E^{n+1} denotes an approximation to $E(t_{n+1})$, $\gamma \in \mathbb{R}$, and g_n is a real-valued function of past information that is determined by the specific choice of the

discretization scheme. We use backward Euler (BE), where $\gamma = 1$ and $g_n(E^n) = E^n$, and the second-order backward differentiation formula (BDF2), where $\gamma = 2/3$ and $g_n(E^n, E^{n-1}) = \frac{4}{3}E^n - \frac{1}{3}E^{n-1}$ for constant Δt_n .

Spatial discretization of (4.1) on uniform meshes is done by subdividing Ω into rectangular control volumes and employing standard finite volume discretization techniques. Cell centers are indexed with integer pairs (i, j) , and fluxes are computed at faces $(i - \frac{1}{2}, j)$ and $(i, j - \frac{1}{2})$ via

$$(4.2) \quad \begin{aligned} (D(E)E_x)_{i-\frac{1}{2},j} &\approx F_{i-\frac{1}{2},j} \equiv D(E)_{i-\frac{1}{2},j} \frac{E_{i,j} - E_{i-1,j}}{\Delta x} \Delta y, \\ (D(E)E_y)_{i,j-\frac{1}{2}} &\approx F_{i,j-\frac{1}{2}} \equiv D(E)_{i,j-\frac{1}{2}} \frac{E_{i,j} - E_{i,j-1}}{\Delta y} \Delta x. \end{aligned}$$

We employ a simple scheme for specifying different materials in which Z in (2.2) is specified at cell centers. With this convention, it is natural to also define $D(E)$ at cell centers, so we need to define face-centered values to compute these fluxes. For this, we use harmonic averages

$$(4.3) \quad \begin{aligned} D(E)_{i-\frac{1}{2},j} &= 2 \left(\frac{1}{D(E)_{i-1,j}} + \frac{1}{D(E)_{i,j}} \right)^{-1} \\ D(E)_{i,j-\frac{1}{2}} &= 2 \left(\frac{1}{D(E)_{i,j-1}} + \frac{1}{D(E)_{i,j}} \right)^{-1} \end{aligned}$$

to define face-centered conductivities. Finally, the discrete spatial operator is obtained by differencing the fluxes on opposite faces and summing the result:

$$(4.4) \quad \nabla \cdot D(E)\nabla E \approx \left(F_{i+\frac{1}{2},j} - F_{i-\frac{1}{2},j} \right) + \left(F_{i,j+\frac{1}{2}} - F_{i,j-\frac{1}{2}} \right).$$

On a SAMR grid, we use the same spatial discretization on patch interiors. There are numerous approaches to handling discretization near changes in resolution. A typical situation for finite volume discretization is depicted in Figure 4. In general, data at cell centers is not properly aligned across the change in resolution. The necessary alignment can be achieved by interpolating data on the coarse side of the coarse/fine interface. Fluxes at fine resolution can then be calculated directly from the aligned data. Alternatively, the aligned data on the coarse side of the interface can be centered in a ghost cell at fine resolution by interpolating in a direction normal to the coarse/fine interface. Finally, the fluxes that reside on fine faces that coincide with a coarse face are summed to obtain a flux on the coarse face. Piecewise linear interpolation leads to the symmetric scheme in [10]; piecewise quadratic interpolation is also commonly used [1].

At coarse/fine interfaces, $D(E)$ is calculated from interpolated values of E . Because $D(E)$ can exhibit large jumps across material interfaces, we must be careful about the placement of refinement regions to ensure the accuracy of our spatial interpolation schemes.

5. Algorithmic Components. The discretization described in §4 leads to a large-scale system of nonlinear equations defined on a SAMR grid. In particular, note that *all* levels in the grid hierarchy are advanced simultaneously with the same time step, similar to [15] but in contrast to [26, 27]. This means that, on convergence,




FIG. 4.1. Schematic of finite volume discretization at interfaces between coarse and fine regions. The leftmost figure shows the first step, which is to align data on both sides of the change in resolution through interpolation of data on the coarse side of the interface. The second step, shown in the middle figure, can either center the aligned data in a ghost cell at the fine resolution, or proceed directly to calculation of a flux on the fine face. Finally, fluxes computed at fine resolution must be synchronized with the flux on the underlying coarse face; this is depicted in the rightmost figure.

our solution conserves energy at coarse/fine interfaces, eliminating the need for re-flux operations when using local time stepping as in [12]. The system of nonlinear equations at each time step is solved with an inexact Newton method. The linearized problems that must be solved at each iteration of the inexact Newton method are preconditioned by the Fast Adaptive Composite grid (FAC) method. These algorithmic components are described in more detail in the following sections.

5.1. Jacobian-free Newton-Krylov Methods. Let $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a nonlinear function and consider calculating the solution $x^* \in \mathbb{R}^n$ of the system of nonlinear equations

$$(5.1) \quad F(x) = 0.$$

Starting with an initial approximation x_0 to x^* , classical Newton's method for solving (5.1) generates a sequence $\{x_k\}$ of approximations to x^* according to

$$(5.2) \quad \begin{aligned} F'(x_k)s_k &= -F(x_k) \\ x_{k+1} &= x_k + s_k \end{aligned}$$

where F' is the Jacobian of F evaluated at x_k . Newton's method is attractive because of its fast local convergence properties. However, for large-scale problems, it is impractical to determine the Newton step s_k in (5.2) with a direct method. Furthermore, when x_k is far from x^* , the linearization that leads to (5.2) may be a poor approximation to $F(x)$. It may be more efficient to require only that s_k satisfy

$$(5.3) \quad \|F(x_k) + F'(x_k)s_k\| \leq \eta \|F(x_k)\|$$

for some $\eta \in (0, 1)$ [8]. Appropriate selection of the forcing term η can lead to superlinear and even quadratic convergence [9]. While any iterative method can be used to find an s_k that satisfies (5.3), methods such as transpose-free Krylov subspace methods require only matrix-vector products, which can be approximated by finite differences

$$F'(x_k)v \approx \frac{F(x_k + \varepsilon v) - F(x_k)}{\varepsilon}.$$

In particular, we use GMRES because of its robustness in the presence of inexact matrix-vector products [6]. Such a *Jacobian-free Newton-Krylov* (JFNK) method is especially advantageous when F' is difficult to compute or expensive to store, and has proven to be effective on a wide variety of problems [14].

JFNK methods facilitate use of an inexact Newton method, since an application only needs to provide methods to evaluate F , set up a preconditioner, and apply the preconditioner. On a SAMR grid, these operations should ideally exploit the structure of the grid. Considerations for evaluating F are described next; corresponding considerations for operations involving the preconditioner are described in § 5.2.

On a regular grid, evaluation of the discrete form of (4.1) is straightforward, and there are many ways to organize this calculation. One approach is

Algorithm 1: Nonlinear function evaluation

- Evaluate the radiation conductivity (2.3).
- Compute the harmonic averages (4.3).
- Compute radiative fluxes (4.2).
- Difference fluxes to obtain (4.4).
- Assemble the nonlinear residual (4.1).

A parallel implementation of Algorithm 1 requires only minor changes that are necessary to satisfy off-processor dependencies for the calculation of (4.2) and (4.3). An additional minor modification of the data structures is also needed to provide storage for the off-processor data; a common approach is to supplement local storage with a layer of ghost cells. On a SAMR grid, data on each patch must likewise be supplemented to satisfy dependencies on data from another patch (which may or may not reside on the same processor). In this case, the required data can come from three sources: other patches at the same resolution; data from a coarser level; or physical boundaries. Once again, providing a layer of ghost cells with each patch is a convenient and effective way to provide storage for such data. Moreover, data from a coarser level must be properly interpolated, and fluxes on coarse faces must be synchronized with fluxes on fine faces (cf. the discussion in § 4). Because of this, JFNK has an added advantage for problems on SAMR grids, since the added burden of determining the structure of the Jacobian at changes in resolution is eliminated. In a parallel environment, some data needed to satisfy these dependencies for a given patch resides on the same processor as that patch. Neither the application nor the solver package should be responsible for this additional bookkeeping, and it is desirable to employ a SAMR grid management package to handle this additional complexity [23]. With such capabilities, we can re-use the computational kernels that implement Algorithm 1 (a principal advantage of SAMR), and Algorithm 1 can be recast as follows:

Algorithm 2: Nonlinear function evaluation

```

for  $k = J$  to 1
  Fill ghost cells on  $\Omega_k^{h_k}$ .
  foreach  $\mathcal{P} \in \Omega_k^{h_k}$ 
    Evaluate the radiation conductivity (2.3).
    Compute the harmonic averages (4.3).
    Compute radiative fluxes (4.2).
  if  $k < J$ 
    Coarsen fluxes from  $\Omega_{k+1}^{h_{k+1}}$ .
for  $k = 1$  to  $J$ 
  foreach  $\mathcal{P} \in \Omega_k^{h_k}$ 
    Difference fluxes to obtain (4.4).
    Assemble the nonlinear residual (4.1).

```

Thus, function evaluation on a SAMR grid is a straightforward generalization of function evaluation on a regular grid, and JFNK can readily be implemented on a SAMR grid, provided suitable representation of SAMR grid data as a vector is available [23]. However, the potential performance gains made possible by the use of AMR will be lost without effective preconditioning.

5.2. Preconditioning. Preconditioning is essential to make the JFNK method competitive. Yet, in the absence of an explicit Jacobian, selection of a preconditioning matrix can be problematic. A simple linearization of (4.1) leads to the choice

$$(5.4) \quad M = \mathbb{I} - \gamma \nabla \cdot D(E_k^{n+1}) \nabla$$

as a preconditioner, where $D(E_k^{n+1})$ is the k^{th} inexact Newton approximation to the time-advanced solution E^{n+1} . This choice was shown to be effective in [24] for problems on uniform grids. The task here is to solve systems of the form $Mz = r$ robustly, accurately, and efficiently when M is discretized on a SAMR grid. For this the Fast Adaptive Composite grid (FAC) method [17] is employed.

FAC extends techniques from multigrid on uniform grids to AMR grids. FAC is a multiplicative Schwarz method. Additive Schwarz variants of FAC exist and are a topic of future research for preconditioning NK methods on AMR grids. A V-cycle implementation of the method is optimal requiring $O(n)$ operations where n is the number of degrees of freedom. As a preconditioner for NK methods, FAC employs smoothing on refinement levels with a coarse grid solve using an approximate solver like multigrid. Below we describe the FAC algorithm for completeness after introducing necessary notation.

- I_c^k and I_k^c denote interlevel data transfer operators (restriction and interpolation, respectively) between composite grid Ω^c and refinement level Ω_k^h . For example, I_k^c could be based on linear interpolation and I_c^k could be defined as the adjoint of I_k^c .
- I_{k+1}^k and I_k^{k+1} denote interlevel data transfer operators (restriction and interpolation, respectively) between adjacent refinement levels Ω_k^h and Ω_{k+1}^h . Operators I_k^c and I_c^k may be considered as compositions of these operators.
- M^c is the composite grid discrete operator obtained by discretizing the PDE on Ω^c and M^k approximates M^c on level k .

Expressed in the notation above, one iteration of a V-cycle FAC algorithm with smoothing on refinement levels and an approximate solve on the coarsest grid is:

Algorithm 3: Fast Adaptive Composite grid (FAC) Method

```

Initialize:  $r^c = f^c - M^c x^c$ ;  $f^k = I_c^k r^c$ 
foreach  $\Omega_k^{h_k}$ ,  $k = J, \dots, 2$ 
  Smooth:  $M^k e^k = f^k$ 
  Correct:  $x^c = x^c + I_k^c e^k$ 
  Update:  $r^c = f^c - M^c x^c$ 
  Set    :  $f^{k-1} = I_c^{k-1} r^c$ 
Solve   :  $M^1 e^1 = f^1$ 
Correct:  $x^c = x^c + I_1^c e^1$ 
foreach  $\Omega_k^{h_k}$ ,  $k = 2, \dots, J$ 
  Update:  $r^c = f^c - M^c x^c$ 
  Set    :  $f^k = I_c^k r^c$ 
  Smooth:  $M^k e^k = f^k$ 
  Correct:  $x^c = x^c + I_k^c e^k$ 

```

6. Numerical Results. In this section we describe the results of some numerical experiments that illustrate the performance of the methods described above. Before describing the test cases, we first specify a few remaining details of our implementation.

Solvers. As described earlier we use a JFNK solver preconditioned by FAC. For this we use the infrastructure detailed in [23] to access the SNES solver from PETSc [2]. SNES iterations are terminated when either an absolute tolerance $\|F(x_k)\| \leq \epsilon_{abs}$, a relative tolerance $\|F(x_k)\| \leq \epsilon_{rel} \|F(x_0)\|$, or a step tolerance $\|s_k\| \leq \epsilon_{step}$ is satisfied; we use $\epsilon_{abs} = \epsilon_{rel} = 10^{-8}$ and $\epsilon_{step} = 10^{-10}$. We found that we had to supplement linesearch backtracking with scaling of the inexact Newton step to ensure positive values in the solution. We use a maximum Krylov subspace dimension of 20, but our linear iteration count in each time step never exceeds this value. The preconditioner is one V-cycle of symmetric FAC with two pre- and post-smoothing sweeps of symmetric red-black Gauss-Seidel on refinement levels. One V-cycle of the SMG solver from the *hypra* [7, 11] library is used on the coarsest level. All calculations were performed on Linux workstations with Intel Xeon 2.4 GHz processors in double precision arithmetic.

Time step control. We base our selection of Δt on the following criterion:

$$\max_{(i,j)} \left(\frac{|E_{i,j}^{n+1} - E_{i,j}^n|}{|E_{i,j}^{n+1}|} \right) \leq CFL_t ,$$

where $CFL_t \in (0, 1)$. This criterion controls the amount of relative change in E from one time step to the next, and is frequently used in applications. For each of the time discretization schemes we use, this leads to an effective time step control algorithm. For example, for BE, this leads to

$$\Delta t \leq CFL_t \min_{(i,j)} \left(\frac{|E_{i,j}^{n+1}|}{|(\nabla \cdot D(E^{n+1}) \nabla E^{n+1})_{i,j}|} \right) .$$

We use $CFL_t = 0.6$. In addition, we control the rate of growth of Δt so that it never increases by more than 10% from one time step to the next.

Selection of refinement regions. At regular intervals (specifically, every fourth time step), we identify regions where enhanced resolution is desired. In this work, we employ simple criteria to identify features of the solution that would benefit from enhanced resolution. A gradient detector is used to identify regions where the solution

changes rapidly. However, a gradient-based refinement criterion can be ineffective for problems that include second-order derivatives (consider the solution of $u_{xx} = f$ on $\Omega = [-1, 1]$ where $f(x) = \cosh^{-2}(x/\epsilon)$ for $0 < \epsilon \ll 1$). Thus, we also use a curvature-based criterion [5]. We have found that, at least for this problem, neither criterion alone is adequate, but that the combination works quite well. Finally, if a cell is marked for refinement, we also mark its eight nearest neighbors. We do not maintain that this approach leads to the desired accuracy result, but our illustrations will show that they do generate AMR configurations that provide a reasonable test of our solution strategy.

Test cases. We use two test cases to illustrate the performance of Newton-

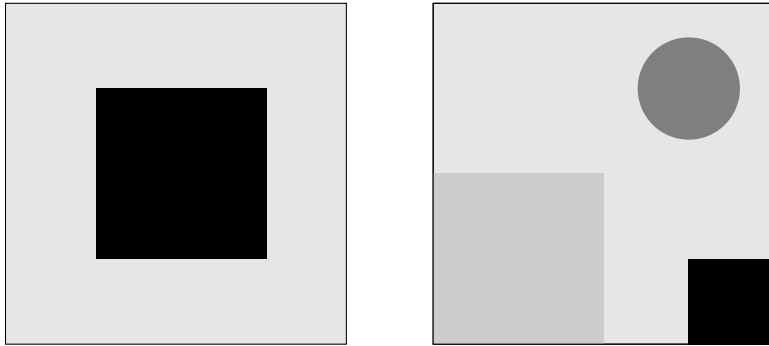


FIG. 6.1. *Material configuration for test cases. Case 1 is on the left and Case 2 is on the right.*

Krylov-FAC for solving (2.1), which are depicted in Figure 6.1. In both test cases, $\Omega = [0, 1]^2$, $\Omega_{\mathcal{R}} = \{x = 0\} \cup \{x = 1\}$, we take $R = 2500$ at $x = 0$ and $R = 0.25$ at $x = 1$, and we integrate until $t = 1500$. The first test case is taken from [21], where it is used to illustrate the performance of a preconditioner for a non-equilibrium diffusion model. The material properties are defined by

$$Z(x, y) = \begin{cases} 100 & (x, y) \in [\frac{1}{4}, \frac{3}{4}]^2 \\ 10 & \text{otherwise} \end{cases}.$$

With our choice of α this leads to a jump of three orders of magnitude in D at the material interface.

Case 2 is taken from [24], where the material properties are given by

$$Z(x, y) = \begin{cases} 100 & (x, y) \in [\frac{3}{4}, 1] \times [0, \frac{1}{4}] \\ 50 & (x - \frac{3}{4})^2 + (y - \frac{3}{4})^2 \leq 0.0225 \\ 20 & (x, y) \in [0.0, \frac{1}{4}]^2 \\ 10 & \text{otherwise} \end{cases}.$$

6.1. Case 1. Time evolution of the solution and the grid hierarchy is shown in Figure 6.2. By $t = 0.25$, the incident energy has heated up the left side of the domain and a thermal front has propagated into the domain and begun to interact with the second material. Note that the finest level tracks the steepest part of the thermal front, where the solution rapidly changes by three orders of magnitude. The finest level continues to track the front as it propagates further into the domain, and follows it as it wraps around the second material at time $t = 1$.




FIG. 6.2. Evolution of solution and grid for Case 1, using a 16×16 base grid plus 4 refinement levels. Boundaries of refinement patches are superimposed on a color contour plot of the solution. The coarsest level is outlined in green; level 1: yellow; level 2: light blue; level 3: magenta; and level 4: peach.

A comparison of the problem size needed for adaptive and uniform mesh calculations appears in Figure 6.3. For the uniform mesh calculation, we use a global 256×256 fine mesh, at the same resolution as the finest level of the adaptive mesh calculation. Initially, as the incident energy heats up the left boundary, a global coarse mesh is sufficient, and the adaptive calculation requires less than 1% of the storage of the uniform grid calculation. As the thermal front propagates into the domain, more refinement levels are needed to resolve it, and the storage required for the adaptive method increases until a maximum of 45% of the uniform grid calculation is needed. As the radiation energy in the interior of the material increases, de-refinement takes place, until the simulation reaches close to steady state at the end of the calculation, and the adaptive calculation requires only 4% of the uniform grid calculation. The adaptive calculation averages about 22.5% over the course of the simulation. Similar behavior is observed when BDF2 is used for the time integration.

Next, we compare the number of iterations required by the adaptive and uniform mesh calculations for backward Euler time integration. This is depicted graphically over the course of the simulation in Figure 6.4 and summarized as averages in Table 6.1. The adaptive mesh calculation requires about the same number of nonlinear iterations as the uniform mesh calculation, and slightly more linear iterations. It is interesting to note the greater variation in the iteration counts for the adaptive calculation, which is likely caused by regriding and the fact that we make no special




FIG. 6.3. Comparison of problem size for adaptive vs. uniformly fine calculation.

provision for discontinuous coefficients on the finer levels. Similar results are obtained when BDF2 is used for time integration.




FIG. 6.4. Comparison of performance on a uniform grid and a locally refined grid, using backward Euler time integration. *NNI* is the number of JFNK iterations and *NLI* is the number of linear iterations, plotted as a function of time.

6.2. Case 2. Time evolution of the adaptive computation appears in Figure 6.5. By $t = 0.5$, the lower conductivity in Region 1 ($Z = 20$) has slowed the thermal front in comparison to the background material ($Z = 10$), which has been recognized by our refinement criteria. By $t = 0.75$, the thermal front has begun to interact with Region 2 ($Z = 50$) and has begun to propagate towards Region 3 ($Z = 100$). Again, the finest region tracks this behavior. At $t = 1.0$, the thermal front continues towards Region 3 and the thermal front has begun to wrap around Region 2 due to its lower conductivity. By $t = 1.25$, the thermal front has begun to interact with Region 3, while the radiation energy in the interior of Region 2 begins increasing.

Comparison of problem sizes is presented in Figure 6.6. Once again, a global coarse grid is sufficient until the thermal wave propagates into the domain, at which

TABLE 6.1

Summary of performance for Case 1. *NNI*: average number of nonlinear iterations; *NLI*: average number of linear iterations.

	BE		BDF2	
	uniform	AMR	uniform	AMR
NNI	4.9	5.1	4.8	5.2
NLI	12.2	14.6	10.4	12.3




FIG. 6.5. Evolution of solution and grid for Case 2. Conditions for this test are the same as in Figure 6.2.

point refinement levels are created and extended into the domain. The adaptive calculation peaks at about 50% of the uniform calculation as the thermal wave breaks around Region 1 and interacts with region Region 2. This is followed by a steady reduction in problem size as the solution increases inside Region 2 and the thermal wave begins to interact with Region 3. At the end of the calculation, only Region 3 requires enhanced resolution, and most of the domain is de-refined. On average, this calculation required about 21% of the size of the uniform calculation. In this case, there is a lot more local variation in the size of adaptive calculation. This is due to our regridding strategy, which leads to behavior in which regions are refined and de-refined in successive regridding events. This behavior can be alleviated by fine tuning the thresholds used in our refinement criteria or by using additional criteria based on truncation error estimates.




FIG. 6.6. Comparison of problem size for adaptive vs. uniformly fine calculation.

Finally, Figure 6.7 compares iteration counts over the course of the simulation when BDF2 time integration is used. Once again, the locally refined calculation requires roughly the same number of nonlinear iterations, and slightly more linear iterations, than the uniform grid calculation. Backward Euler time integration produces similar results. These data are summarized as averages in Table 6.2.




FIG. 6.7. Comparison of performance on a uniform grid and a locally refined grid. Notation is the same as for Figure 6.4.

7. Conclusions. We have demonstrated the feasibility of combining implicit time integration with adaptive mesh refinement for equilibrium radiation diffusion problems. The numerical examples presented show that, for both first and second

TABLE 6.2
Summary of performance for Case 2. Notation is the same as in Table 6.1.

	BE		BDF2	
	uniform	AMR	uniform	AMR
NNI	5.2	5.1	4.8	5.4
NLI	11.2	13.5	9.7	12.0

order time integration schemes, the multilevel FAC preconditioner effectively controls the number of linear iterations per time step.

There are several issues in our implicit AMR strategy that merit further exploration. In particular, our interlevel transfers use geometric interpolation, an approach that is known to lead to robustness problems in the presence of discontinuous coefficients. This may contribute to some of the convergence failures of JFNK that we have observed. Our criteria for selecting refinement regions are based on feature detection, and not estimation of spatial errors. A systematic evaluation of temporal and spatial convergence orders is needed to determine the efficacy of this approach. We also need to evaluate the parallel performance of our code and explore the use of asynchronous multilevel preconditioners.

A major issue that we have not addressed to our satisfaction is the fact that previous information is no longer a solution after regridding. This was observed in [3], where it was argued that higher-order spatial interpolation could solve the problem, and in [15], where it is handled by allowing an underlying variable order, variable step integrator to reduce the order of integration and the size of the time step. While it was noted that this behavior leads to some loss of efficiency, the issue of the effect on the order of temporal accuracy was not explored. We handle this issue directly (albeit in a brute force manner) by re-solving (2.1) on the new grid to update the most recent solution information. At early times, when there is not much change in the grid hierarchy, no additional work is necessary, since in these circumstances the interpolated past information still satisfies (2.1) on the new grid to the specified tolerances. However, at later times, past information that has been interpolated to the new grid does not satisfy (2.1) on the new grid, and a full solve is needed. We plan to explore this issue in more detail in future work.

REFERENCES

- [1] A. S. ALMGREN, J. B. BELL, P. COLELLA, L. H. HOWELL, AND M. L. WELCOME, *A conservative adaptive projection method for variable density incompressible Navier-Stokes equations*, J. Comput. Phys., 142 (1998), pp. 1–46.
- [2] S. BALAY, W. D. GROPP, L. C. MCINNES, AND B. F. SMITH, *PETSc Users Manual*, Tech. Rep. ANL-95/11 - Revision 2.1.0, Argonne National Laboratory, 2001.
- [3] M. BERZINS, P. CAPON, AND P. JIMACK, *On spatial adaptivity and interpolation when using the method of lines*, Appl. Num. Math., 26 (1998), pp. 117–133.
- [4] R. BLAHEETA, *Adaptive composite grid methods for problems of plasticity*, Math. Comp. Sim., 50 (1999), pp. 123–134.
- [5] J. G. BLOM, R. A. TROMPERT, AND J. G. VERWER, *Algorithm 758: VLUGR2: a vectorizable adaptive-grid solver for PDEs in 2D*, ACM-TOMS, 22 (1996), pp. 302–329.
- [6] P. N. BROWN, *A local convergence theory for combined inexact-Newton/finite-difference projection methods*, SIAM J. Numer. Anal., 24 (1987), pp. 407–434.
- [7] E. CHOW, A. C. CLEARY, AND R. D. FALGOUT, *Design of the hypre preconditioner library*, in Object Oriented Methods for Inter-operable Scientific and Engineering Computing., M. E.

- Henderson, C. R. Anderson, and S. L. Lyons, eds., SIAM, 1999, pp. 106–116.
- [8] R. S. DEMBO, S. C. EISENSTAT, AND T. STEIHAUG, *Inexact Newton methods*, SIAM J. Numer. Anal., 19 (1982), pp. 400–408.
- [9] S. C. EISENSTAT AND H. F. WALKER, *Globally convergent inexact Newton methods*, SIAM J. Optimization, 4 (1994), pp. 393–422.
- [10] R. E. EWING, R. D. LAZAROV, AND P. S. VASSILEVSKI, *Local refinement techniques for elliptic problems on cell-centered grids. I: Error analysis*, Math. Comp., 56 (1991), pp. 437–461.
- [11] R. D. FALGOUT AND U. M. YANG, *hypr: a library of high performance preconditioners*, in Computational Science - CARS 2002 Part III, P. M. A. Sloot, C. J. K. Tan, J. J. Dongarra, and A. G. Hoekstra, eds., vol. 2331 of Lecture Notes in Computer Science, New York, 2002, Springer-Verlag, pp. 632–641.
- [12] L. H. HOWELL AND J. A. GREENOUGH, *Radiation diffusion for multi-fluid Eulerian hydrodynamics with adaptive mesh refinement*, J. Comput. Phys., 184 (2003), pp. 53–78.
- [13] D. A. KNOLL, L. CHACÓN, L. G. MARGOLIN, AND V. A. MOUSSEAU, *On balanced approximations for time integration of multiple time scale systems*, J. Comput. Phys., 185 (2003), pp. 583–611.
- [14] D. A. KNOLL AND D. E. KEYES, *Jacobian-free Newton-Krylov methods: a survey of approaches and applications*, J. Comput. Phys., 193 (2004), pp. 357–397.
- [15] S. LI, L. PETZOLD, AND J. M. HYMAN, *Solution adapted mesh refinement and sensitivity analysis for parabolic partial differential equation systems*, in Large-Scale PDE-Constrained Optimization, vol. 30 of Lecture Notes in Comp. Sc. and Engrg., Heidelberg, 2003, Springer-Verlag.
- [16] R. E. MARSHAK, *Effect of radiation on shock wave behavior*, Physics of Fluids, 1 (1958), pp. 24–29.
- [17] S. F. MCCORMICK, *Multilevel Adaptive Methods for Partial Differential Equations*, SIAM, Philadelphia, PA, 1989.
- [18] S. F. MCCORMICK AND J. W. THOMAS, *The Fast Adaptive Composite grid (FAC) method for elliptic equations*, Math. Comp., 46 (1986), pp. 439–456.
- [19] D. MIHALAS AND B. WEIBEL-MIHALAS, *Foundations of Radiation Hydrodynamics*, Dover Publications, Inc., Mineola, NY, 1999.
- [20] E. MINGUEZ, P. MARTEL, M. GIL, J. G. RUBIANO, AND R. RODRIGUEZ, *Analytic opacity formulas for ICF elements*, Fusion Eng. Des., 60 (2002), pp. 17–25.
- [21] V. A. MOUSSEAU, D. A. KNOLL, AND W. J. RIDER, *Physics-based preconditioning and the Newton-Krylov method for non-equilibrium radiation diffusion*, J. Comput. Phys., 160 (2000), pp. 743–765.
- [22] C. C. OBER AND J. N. SHADID, *Studies on the accuracy of time integration methods for the radiation diffusion equations*, J. Comput. Phys., 195 (2004), pp. 743–772.
- [23] M. PERNICE AND R. D. HORNUNG, *Newton-Krylov-FAC methods for problems discretized on locally refined grids*, Tech. Rep. LAUR 03-0619, UCRL-JC-151614, Los Alamos National Laboratory and Lawrence Livermore National Laboratory, 2003. (submitted to Computing and Visualization in Science).
- [24] W. J. RIDER, D. A. KNOLL, AND G. L. OLSON, *A multigrid Newton-Krylov method for multi-material equilibrium radiation diffusion*, J. Comput. Phys., 152 (1999), pp. 164–191.
- [25] L. STALS, *Comparison of non-linear solvers for the solution of radiation transport equations*, Elec. Trans. Num. Anal., 15 (2003), pp. 78–93.
- [26] R. A. TROMPERT AND J. G. VERWER, *Analysis of the implicit Euler local uniform grid refinement method*, SIAM J. Sci. Comput., 14 (1993), pp. 259–278.
- [27] ———, *Runge-Kutta methods and local uniform grid refinement*, Math. Comp., 60 (1993), pp. 591–616.
- [28] K.-H. WINKLER, M. NORMAN, AND D. MIHALAS, *Implicit adaptive-grid radiation-hydrodynamics*, in Multiple Time Scales, Academic Press, 1985.
- [29] A. M. WINSLOW, *Extensions of asymptotic neutron diffusion theory*, Nucl. Sci. and Eng., 32 (1968), pp. 101–110.