

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 front, or Marshak wave, under suitable initial and boundary conditions. The front can vary several orders of magnitude over a very short distance. Resolving the shape of the Marshak wave 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, combustion applications, atmospheric dynamics, and medical imaging. When photon mean free paths are much shorter than characteristic length scales, a diffusion approximation provides a reasonably accurate description of radiation penetrating from a hot source to a cold medium. This approximation features a nonlinear conduction coefficient that leads to formation of a sharp front, in which the solution can vary several orders of magnitude over a very short distance. The shape of the 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 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 that precludes effective use of higher order time integration methods, and requires small time steps to maintain time accuracy. Analytic and computational results that demonstrate

*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, an affirmative action/equal opportunity employer. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U. S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness. LA-UR-05-1776.

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

degradation in time accuracy associated with linearization in the presence of strong nonlinear coefficients can be found in [24] and [15]. 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 [26].

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 on AMR grids.

Prior efforts in implicit AMR have not adequately addressed the algorithmic scalability issue. Early work by Winkler, Norman, and Mihalas [31] 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 two- or three-dimensional problems, and even in one dimension difficulties with gridpoint movement are encountered. We use h -refinement, in particular structured AMR (SAMR), to avoid these problems. Both Trompert and Verwer [29, 30] and Li, Petzold, and Hyman [17] use incomplete factorization to precondition their linear solvers. Incomplete factorization can be expensive to set up, can have high storage overhead, scales poorly with increasing resolution, and is difficult to implement in parallel. 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 [19, 20] 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 [20]. Similar approaches for elasto-plasticity [5] and equilibrium radiation diffusion [28] use multilevel methods on unstructured grids to solve the linearized equations, but still require formation of the Jacobian. More recently, Howell and Greenough [14] solve a *linearized* radiation diffusion problem on a SAMR grid as part of a semi-implicit radiation-hydrodynamics algorithm. However, 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 and to produce solutions with accuracy commensurate with the finest resolution on the AMR grid.

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 conclusions are summarized in §7.

2. Equilibrium Radiation Diffusion. We provide a brief description of the equations that govern equilibrium radiation diffusion. See [21] 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

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

with initial conditions $E(t = 0) = E_0$. We use a flux-limited model of radiative conductivity

$$(2.2) \quad D(E) = \left(\frac{1}{Z^\alpha E^\beta} + \frac{\|\nabla E\|}{|E|} \right)^{-1},$$

where $\alpha < 0$, $\beta \in [0, 1]$, and Z is the atomic number of the medium. The term $\|\nabla E\|/|E|$ is Wilson’s form for flux limiting [21, 32], which is an *ad hoc* adjustment introduced to prevent nonphysical behavior in the solution, and which reproduces the correct asymptotic behavior. The parameters α and β can be determined experimentally or through simulation [22]. In this work $\alpha = -3$ and $\beta = 0.75$.

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

$$(2.3) \quad \begin{aligned} \frac{1}{2}\mathbf{n} \cdot D(E)\nabla E + \frac{1}{4}E &= R \text{ on } \partial\Omega_{\mathcal{R}}, \quad t > 0, \\ \mathbf{n} \cdot D(E)\nabla E &= 0 \text{ on } \partial\Omega_{\mathcal{N}}, \quad t > 0, \end{aligned}$$

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 [18].

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) is *h*-refinement with local fine grid patches placed over coarser grid patches providing




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

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 structured grids is possible.

A SAMR grid consists of a collection of grid patches at different grid resolutions, which together cover the computational domain. Patches with the same mesh resolution are grouped together and form a refinement level. Patches 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 patch, with the underlying coarse patches being the parents. A child patch can have several parent patches and vice versa. The boundaries of child and parent patches are not assumed 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 ℓ are defined in cells not covered by cells belonging to level $\ell + 1$. 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 ℓ , $\ell = 1, 2, \dots, J$, denoted by $\Omega_\ell^{h_\ell}$ or more simply as Ω_ℓ^h , consists of a collection $\{\mathcal{P}\}$ of grid patches at the same grid resolution h_ℓ . The refinement levels are ordered by increasing grid resolution with $h_\ell = r h_{\ell+1}$ $\ell = 1, 2, \dots, J - 1$, 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 $\Omega_{\ell+1}^h$ fully nest within the subdomain(s) covered by Ω_ℓ^h . Figure 3.1 illustrates this multilevel structure; Figures 6.4 and 6.6 in §6 provide more examples of multilevel grid configurations.

4. Discretization. We begin by first discretizing (2.1) in time. Let $t_{n+1} = t_n + \Delta t_n$, $n = 0, 1, \dots$, where $t_0 = 0$ and 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. For backward Euler (BE), $\gamma = 1$ and $g_n(E^n) = E^n$, and for the second-order backward differentiation formula (BDF2), $\gamma = 2/3$ 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

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

A simple scheme is employed for specifying different materials with Z in (2.2) specified at cell centers. With this convention, it is natural to also define $D(E)$ at cell centers, so face-centered values need to be defined to compute these fluxes. For this, harmonic averages

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

are used 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, the same spatial discretization is used in patch interiors. There are numerous approaches to handling discretization near changes in resolution. A typical situation for finite volume discretization is depicted in Figure 4.1. 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 with adjustments leads to the symmetric scheme in [11]; piecewise quadratic interpolation is also commonly used [1]. The results presented in this paper use piecewise quadratic interpolation, though similar performance results have been obtained with piecewise linear interpolation.

For ghost points at coarse/fine interfaces, $D(E)$ is interpolated piecewise constant from the coarser level. Because $D(E)$ can exhibit large jumps across material interfaces, refinement regions must be placed carefully to ensure the accuracy of the 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 [17] but in contrast to [29, 30]. This means that, on convergence, our solution conserves energy at coarse/fine interfaces, eliminating the need for re-flux operations when using local time stepping as in [14]. The system of nonlinear




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.

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

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

where F' is the Jacobian of F evaluated at x_k and $\theta \in (0, 1]$ is a damping parameter. 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 [10]. 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 [7]. 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 [16].

The damping parameter θ is determined by a combination of traditional linesearch backtracking [9] and problem-specific constraints. Physically, the radiation energy density E must be non-negative. However there is nothing to prevent the calculation of an inexact Newton step s_k that leads to $E < 0$ in some locations. We preclude this possibility by first scaling s_k to ensure that the updated solution satisfies the physical constraints. Linesearch backtracking is then invoked on the scaled step in the usual way.

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 (2.2).
- Compute the harmonic averages (4.3).
- Compute radiative fluxes (4.2).
- Difference fluxes to obtain (4.4).
- Assemble the nonlinear residual (4.1).

On a SAMR grid, computation is performed in each patch in the grid hierarchy. Since each patch is logically rectangular, software that implements the steps in Algorithm 1 can be used on each of the patches. However, data on each patch must be supplemented to satisfy inter-patch data dependencies. The considerations are similar to those needed for a parallel implementation of Algorithm 1, except that in the SAMR case the required data can come from three sources: other patches at the same resolution; data from a coarser level; or physical boundaries. As with parallel computation, these data dependencies are satisfied by supplementing the storage in each patch with a layer of ghost cells. Data is placed in these locations via copies from neighboring patches at the same refinement level, interpolating data from coarser levels as discussed in §4, and extrapolating data at physical boundaries. In addition, fluxes on coarse faces must be properly synchronized with fluxes on fine faces (cf. the discussion in §4), and the final result of Algorithm 1 must be synchronized in slave cells via interpolation. These considerations are summarized in Algorithm 2.

Algorithm 2: Nonlinear function evaluation

```

for  $\ell = J$  to 1
  Fill ghost cells on  $\Omega_\ell^h$ .
  foreach  $\mathcal{P} \in \Omega_\ell^h$ 
    Evaluate the radiation conductivity (2.2).
    Compute the harmonic averages (4.3).
    Compute radiative fluxes (4.2).
  if  $\ell < J$ 
    Coarsen fluxes from  $\Omega_{\ell+1}^h$  to  $\Omega_\ell^h$ .
for  $\ell = 1$  to  $J$ 
  foreach  $\mathcal{P} \in \Omega_\ell^h$ 
    Difference fluxes to obtain (4.4).
    Assemble the nonlinear residual (4.1).
  if  $\ell > 1$ 
    Interpolate the nonlinear residual from  $\Omega_\ell^h$  to  $\Omega_{\ell-1}^h$ .

```

Thus, function evaluation on a SAMR grid is a straightforward generalization of function evaluation on a regular grid. Moreover, JFNK has an added advantage for problems on SAMR grids, since the burden of determining the structure and entries of the Jacobian at changes in resolution is eliminated. The additional bookkeeping needed for a SAMR grid can be provided by a SAMR grid management package [25]. 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 [26] 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 [19] is employed.

FAC extends techniques from multigrid on uniform grids to AMR grids. FAC is a multiplicative Schwarz method. A V-cycle implementation of the method is optimal requiring $\mathcal{O}(n)$ 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^ℓ and I_ℓ^c denote interlevel data transfer operators (restriction and interpolation, respectively) between the composite fine grid Ω^c and refinement level Ω_ℓ^h . For example, in this paper I_ℓ^c is based on bilinear interpolation and I_c^ℓ is based on a simple averaging of fine cell data to coarse cells.
- $I_{\ell+1}^\ell$ and $I_\ell^{\ell+1}$ denote interlevel data transfer operators (restriction and interpolation, respectively) between adjacent refinement levels Ω_ℓ^h and $\Omega_{\ell+1}^h$. Operators I_ℓ^c and I_c^ℓ may be considered as compositions of these operators.
- M^c is the composite fine grid discrete operator obtained by discretizing (5.4) on Ω^c , and M^ℓ approximates M^c on level ℓ .

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

Algorithm 3: Fast Adaptive Composite grid (FAC) Method

Initialize: $r^c = f^c - M^c x^c$; $f^\ell = I_c^\ell r^c$

foreach Ω_ℓ^h , $l = J, \dots, 2$

Smooth: $M^\ell e^\ell = f^\ell$

Correct: $x^c = x^c + I_\ell^c e^\ell$

Update: $r^c = f^c - M^c x^c$

Set : $f^{\ell-1} = I_c^{\ell-1} r^c$

Solve : $M^1 e^1 = f^1$

Correct: $x^c = x^c + I_1^c e^1$

foreach Ω_ℓ^h , $l = 2, \dots, J$

Update: $r^c = f^c - M^c x^c$

Set : $f^\ell = I_c^\ell r^c$

Smooth: $M^\ell e^\ell = f^\ell$

Correct: $x^c = x^c + I_\ell^c e^\ell$

6. Numerical Results. This section illustrates the performance of Newton-Krylov-FAC for equilibrium radiation diffusion problems on locally refined grids with a comparison to uniform grid calculations. Before describing the test cases, we first specify a few remaining details of the implementation.

Solvers and Software. As described earlier, a JFNK solver preconditioned by FAC is used to solve (4.1). This is accomplished using the infrastructure detailed in [25]. This infrastructure provides access to the capabilities of SAMRAI [13] to implement SAMR-specific aspects of Algorithms 2 and 3, and access to an implementation of JFNK from PETSc’s Scalable Nonlinear Equation Solver (SNES) package [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 $\epsilon_{step} = 10^{-10}$. For the linear systems that must be solved at each inexact Newton step, GMRES is used with a maximum Krylov subspace dimension of 20, and $\eta = 0.05$ is used in (5.3) to terminate the linear iterations. Note that it was never necessary to restart the GMRES iterations with these choices of maximum Krylov subspace dimension and stopping criteria. The preconditioner is one V-cycle of FAC with two pre- and post-smoothing sweeps of red-black Gauß-Seidel on refinement levels. One V-cycle of semi-coarsening multigrid (SMG) [27] implemented in *hypr* [12] 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. Equilibrium radiation diffusion problems (2.1) are rarely solved in isolation, and are more often encountered as part of a more sophisticated approximation to radiation transport or as one component in a larger multi-physics application. Thus, time step selection for solving (2.1) must account for other time scale restrictions that may be present. In such circumstances, a commonly-encountered strategy for choosing Δt is based on restricting the maximum fractional change of E to some prescribed percentage $\sigma \in (0, 1)$ [3, 26]:

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

For the time discretization schemes used (BE and BDF2), this leads to an effective time step control algorithm. For example, for BE, this leads to

$$(6.1) \quad \Delta t \leq \sigma \min_{(i,j)} \left(\frac{|E_{i,j}^{n+1}|}{|(\nabla \cdot D(E^{n+1}) \nabla E^{n+1})_{i,j}|} \right).$$

In practice, for a stand-alone solution of (2.1), this has several shortcomings. Initially, Δt based on (6.1) is too large, so instead the calculations begin with a smaller time step $\Delta t = 10^{-6}$, which is increased by 10% until (6.1) yields the smaller value for Δt . As the solution approaches steady state, (6.1) produces time steps that quickly become too large; consequently, the rate of growth of Δt is limited by 5% for $\Delta t \in (0.5, 10.0)$ and 3% for $\Delta t \geq 10.0$.

Selection of refinement regions. At regular intervals (specifically, every tenth time step¹), regions are identified where enhanced resolution is desired. In this work, simple criteria are employed to identify features² of the solution that would benefit from enhanced resolution. A gradient detector is used to identify regions where the solution changes rapidly. We also use a curvature-based criterion [6]. It is found that, at least for this problem, the curvature-based criterion leads to slightly larger refinement regions, making it easier to ensure that the base of the thermal front is entirely contained within the finest refinement level without resorting to *ad hoc* procedures. However, the curvature-based criterion alone can miss relatively flat regions where the solution nevertheless changes rapidly. Finally, if a cell is marked for refinement, its eight nearest neighbors are also marked.

Test cases. Two test cases are used 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, $\Omega = [0, 1]^2$ and (2.3) is imposed at both the left $x = 0$ boundary and the right $x = 1$ boundary. Both cases are integrated to a final time of $t = 1500$. Case 1 is taken from [23], where it is used to illustrate the performance of a preconditioner for a non-equilibrium diffusion model. The material properties are defined by

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

¹This choice is based on trial and error, and represents a trade-off between the size of the refinement regions and ensuring that the base of the Marshak wave is always contained within the finest refinement level until the next regridding event occurs.

²Efforts to develop an error-based refinement criterion for this problem is deferred to later work.

With $\alpha = -3$, this leads to a jump of three orders of magnitude in D at the material interface.

Case 2 is adapted from the problem solved in [23]. In this, the material properties are given by

$$Z(x, y) = \begin{cases} 100, & (x, y) \in [\frac{3}{4}, 1] \times [0, \frac{1}{4}] \\ 50, & (x, y) \in [\frac{5}{8}, \frac{5}{8}] \times [\frac{7}{8}, \frac{7}{8}] \\ 20, & (x, y) \in [0, \frac{1}{2}]^2 \\ 10, & \text{otherwise.} \end{cases}$$

6.1. Accuracy. Before presenting performance results for Newton-Krylov-FAC methods on unsteady equilibrium radiation diffusion problems, accuracy measurements are presented. For these measurements, the pattern of local refinement in the dynamic grid calculations is not controlled, but allowed to evolve according to the strategies outlined above. Thus, this serves as a verification exercise for both the discretization scheme outlined in §4 as well as the regridding operations, which collectively identify a new computational grid and transfer the solution from the old grid to the new grid.

The accuracy studies are conducted for Case 2 with BDF2 as the time integrator and $\sigma = 0.05$ in the time step control algorithm. For boundary conditions (2.3), $R = 25.0$ at the $x = 0$ boundary and $R = 0.25$ at the $x = 1$ boundary. Since an analytic solution is not available for comparison, a reference solution on a 512×512 uniform grid is first computed, and results obtained on locally refined grids are compared against this reference solution. In regions where the locally refined grid is not at the same resolution as the uniform grid for the reference solution, the comparison is based on piecewise linear interpolation of the four nearest uniform grid neighbors.




FIG. 6.2. L_2 -norm of relative difference in locally refined solution against reference solution for a fixed base grid.

The first set of accuracy results uses a fixed base grid and fixed number of refinement levels, and examines errors over the course of the entire time integration. This appears in Figure 6.2. This shows the relative error over time in the solution computed on locally refined grids against the reference solution. The top curve is for a 32×32 grid with no refinement; each successive curve uses the same base grid but adds a level of refinement. As refinement levels are added, accuracy improvements that appear to be second order are observed. Once four refinement levels are added to

base grid, the finest level of the locally refined calculation is at the same resolution as the reference solution. There, at least three digits of agreement between the reference solution and solution obtained on the locally refined grid is observed.

The next set of accuracy results varies the base grid resolution as well as the number of refinement levels at fixed times. These results are given in Figure 6.3. Here we see excellent agreement among the solutions computed on different locally refined grids, at least until the finest grid level matches the resolution of the reference solution. Note that the finer base grids lead to better agreement with the reference solution.

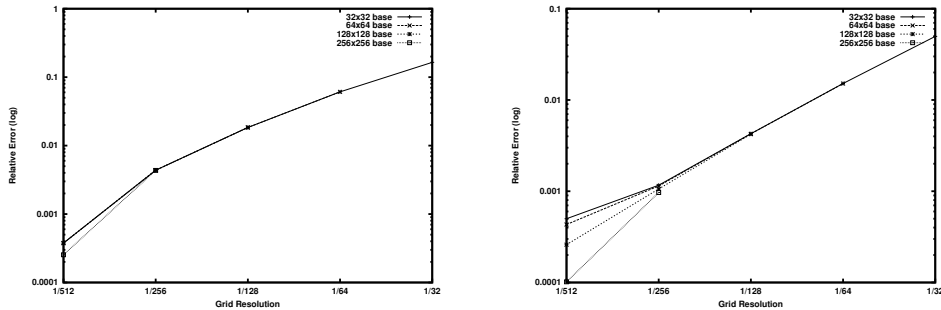


FIG. 6.3. L_2 -norm of relative difference in locally refined solution against reference solution for varying base grids. Horizontal axis shows the mesh size of the finest grid level in any of the configurations. Time $t = 0.5$ is on the left, and time $t = 2.5$ is on the right.

6.2. Performance. Performance results are now presented for the two test cases described previously. In order to stress the performance of the solvers, we set $R = 2500$ at the $x = 0$ boundary and $R = 0.25$ at the $x = 1$ boundary in (2.3), and set $\sigma = 0.1$ within the time step control algorithm. Performance results are presented for both BE and BDF2 computations.

6.2.1. Case 1. Time evolution of the solution and the grid hierarchy is shown in Figure 6.4. By $t = 0.25$, the incident energy has heated up the left side of the domain and a Marshak wave has propagated into the domain and begun to interact with the second material. Note that the finest level tracks the steepest part of the 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 $t = 1$.

A comparison of the problem size needed for adaptive and uniform mesh calculations appears in Figure 6.5. In this example, the locally refined grid starts with a 32×32 coarse grid plus 4 refinement levels. Initially, refinement is automatically introduced only along the $x = 0$ boundary to resolve the increase in energy density due to the imposed radiation field. During this initial period the adaptive calculation requires less than 6% of the storage of the uniform grid calculation. As the front propagates into the domain, the refinement levels grow to capture the evolving front, and the storage required for the adaptive method increases until a maximum of 22% 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 2% of the uniform grid calculation. The size of the adaptive calculation averages

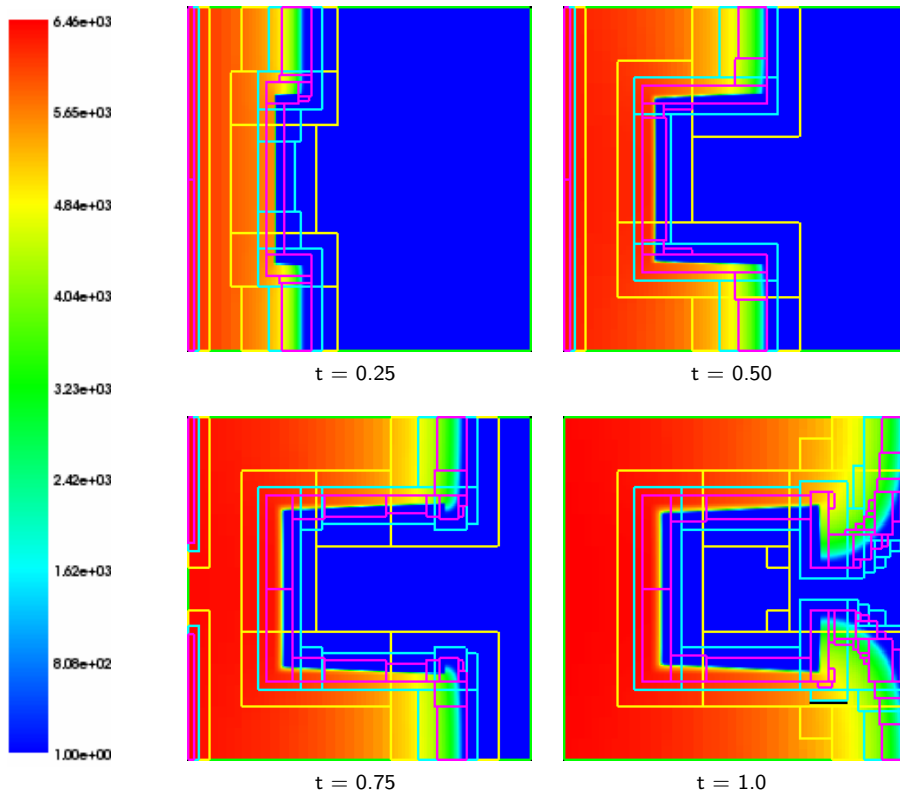


FIG. 6.4. Evolution of solution and grid for Case 1, using a 32×32 base grid plus 3 refinement levels. Boundaries of refinement patches are superimposed on a pseudocolor plot of the solution using a logarithmic scale. The coarsest level is outlined in green; level 1: yellow; level 2: light blue; and level 3: magenta.

less than 10% of the size of the uniform calculation 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 BE and BDF2 time integration. This is summarized as averages in Table 6.1. There is very little variation in the number of nonlinear iterations for a fixed base grid, and a slight decrease as the base grid is refined for a given number of refinement levels. Fixing resolution, we see a slight increase in the number of nonlinear iterations as the resolution of the base grid is increased. The number of linear iterations, which is a better reflection of the amount of work performed, shows similar trends, though with greater variation. This is not very surprising, as SMG's very robust semi-coarsening/line relaxation algorithm on the global fine grid is being replaced with simple point relaxation, geometric interlevel transfers, and rediscritized operators at each refinement level. It must also be kept in mind that the larger iteration counts occur on much smaller problems. For example, we are comparing a combination of expensive SMG iterations on a 32×32 grid plus point Gauß-Seidel iterations on intermediate levels against SMG iterations on a 512×512 uniform grid.

6.2.2. Case 2. Time evolution of the adaptive computation appears in Figure 6.6. By $t = 0.5$, the lower conductivity in Region 1 ($Z = 20$) has slowed the

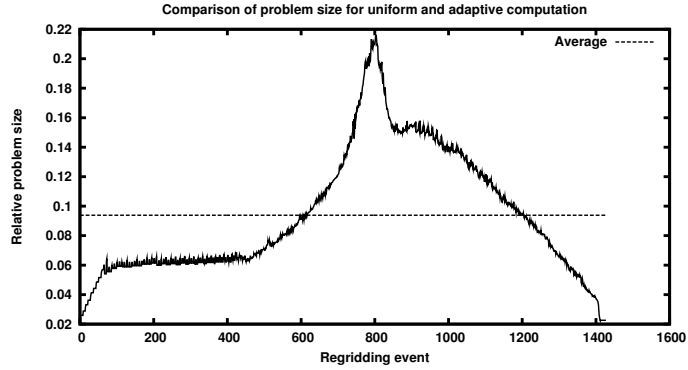


FIG. 6.5. Comparison of problem size for adaptive vs. uniformly fine calculation.

TABLE 6.1

Summary of performance for Case 1. The first column gives the size of the base grid; J is the number of refinement levels. Performance at a fixed finest resolution is obtained by reading diagonally from lower left to upper right. Grid configurations that were not run are denoted by $-$.

Average number of nonlinear iterations per time step										
J	BE					BDF2				
	1	2	3	4	5	1	2	3	4	5
32×32	4.4	4.2	4.2	4.3	4.5	4.4	4.3	4.3	4.2	4.4
64×64	4.3	4.0	4.1	4.4	-	4.4	4.1	4.1	4.3	-
128×128	3.9	3.9	4.2	-	-	4.1	3.9	4.1	-	-
256×256	3.7	3.9	-	-	-	3.8	3.9	-	-	-
512×512	3.8	-	-	-	-	3.6	-	-	-	-

Average number of linear iterations per time step										
J	BE					BDF2				
	1	2	3	4	5	1	2	3	4	5
32×32	5.9	6.6	7.2	8.4	9.6	5.5	6.1	6.4	7.4	8.4
64×64	5.8	6.7	7.6	9.1	-	5.4	6.1	6.8	7.8	-
128×128	5.8	7.0	8.2	-	-	5.3	6.3	7.3	-	-
256×256	5.9	7.1	-	-	-	5.3	6.5	-	-	-
512×512	6.0	-	-	-	-	5.5	-	-	-	-

Marshak wave in comparison to the background material ($Z = 10$), which has been recognized by our refinement criteria. By $t = 0.75$, the 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 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 Marshak wave 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 6.7. The situation is similar to what was seen for Case 1. This time, the number of cells in the adapted grid

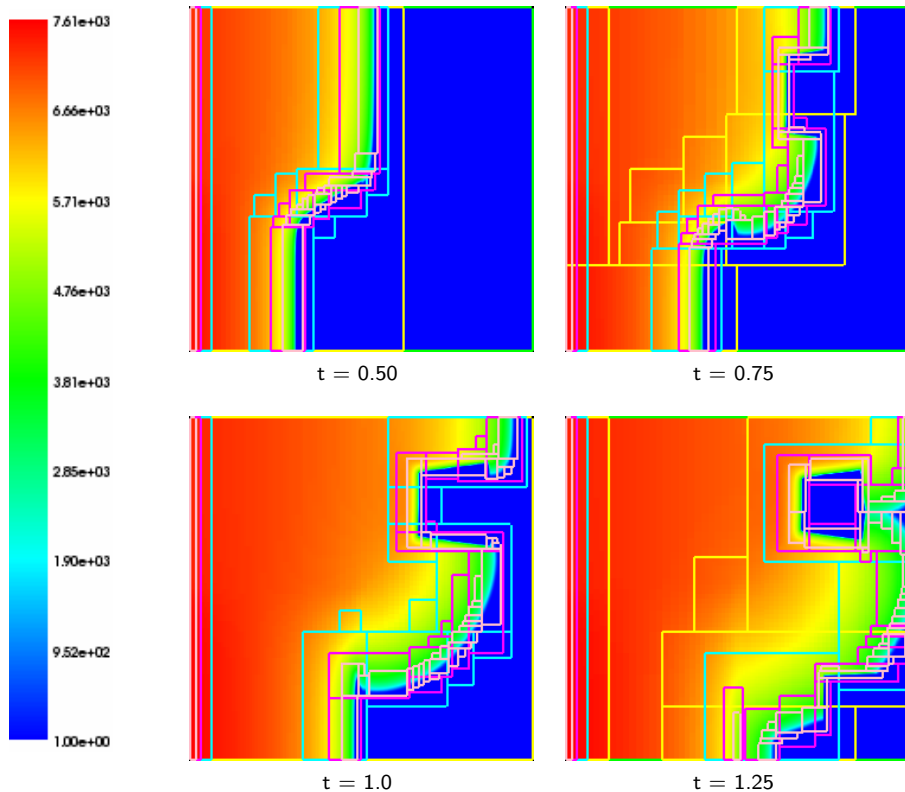


FIG. 6.6. Evolution of solution and grid for Case 2, using a 32×32 base grid plus 4 refinement levels. Boundaries of refinement patches are superimposed on a pseudocolor plot of the solution using a logarithmic color scale. The coarsest level is outlined in green; level 1: yellow; level 2: light blue; level 3: magenta; level 4: peach.

increases more quickly due to the presence Region 1, adjacent to the $x = 0$ boundary. Eventually there is a decrease in the size of the dynamic calculation as Region 1 is de-refined and before resolution is increased in Region 2. Two inflection points are seen in the size of the locally refined calculation, initially as Region 2 is fully resolved and resolution is increased around Region 3, and subsequently as Regions 2 and 3 are de-refined. The number of cells in the dynamic calculation peaks at less than 20% of the uniform grid calculation, then decreases steadily. On average the dynamic calculation is around 8% of the size of the uniform grid calculation.

Table 6.2 compares nonlinear and linear iteration counts per time step. Once again little variation is seen in the number of nonlinear iterations per time step for a fixed base grid size or for fixed finest resolution, and a small decrease in this iteration count for a fixed number of refinement levels. In contrast, the number of linear iterations per time step increases slowly as more refinement levels are added, and increases by nearly half as we fix resolution and move from a global fine grid to a locally refined calculation. Again, this is likely due to the fact that operators on refinement levels are simply obtained by rediscretization, and interlevel transfer operators are purely geometric.

7. Conclusions and Future Work. The results presented demonstrate the feasibility of combining implicit time integration with adaptive mesh refinement for




FIG. 6.7. Comparison of problem size for adaptive vs. uniformly fine calculation.

TABLE 6.2

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

Average number of nonlinear iterations per time step										
J	BE					BDF2				
	1	2	3	4	5	1	2	3	4	5
32×32	4.4	4.3	4.0	4.1	4.2	4.4	4.3	4.0	4.0	4.1
64×64	4.3	3.9	3.9	4.1	–	4.3	3.9	3.9	4.0	–
128×128	4.0	3.7	3.9	–	–	4.1	3.7	3.9	–	–
256×256	3.8	3.7	–	–	–	3.8	3.7	–	–	–
512×512	3.7	–	–	–	–	3.6	–	–	–	–

Average number of linear iterations per time step										
J	BE					BDF2				
	1	2	3	4	5	1	2	3	4	5
32×32	5.8	6.3	7.1	7.8	8.7	5.4	5.8	6.4	7.1	7.8
64×64	5.6	6.6	7.1	8.4	–	5.2	6.0	6.7	7.6	–
128×128	5.7	6.6	7.8	–	–	5.1	6.1	7.0	–	–
256×256	5.6	7.0	–	–	–	5.0	6.3	–	–	–
512×512	5.7	–	–	–	–	5.2	–	–	–	–

equilibrium radiation diffusion problems. The numerical examples show that, for both first and second order time integration schemes, the multilevel FAC preconditioner effectively controls the number of linear iterations per time step. While good scalability with the number of refinement levels is also demonstrated, these results should be considered preliminary, since the refinement regions change as the number of levels is increased and the resolution of the coarsest level is decreased.

There are several issues in implicit AMR that merit improvement and further exploration. In particular, the interlevel transfers use geometric interpolation, an approach that is known to lead to robustness problems in the presence of discontinuous coefficients. Coarse grid operators are defined via rediscrretization, rather than Galerkin coarsening. Evaluation of the parallel performance of the code is necessary,

as is exploring the use of asynchronous multilevel preconditioners. The 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. Before this can occur, the fact that previous information is no longer a solution after regridding must be addressed. We handle this by re-solving (2.1) on the new grid to update the most recent solution information. In [4] it was argued that this can be addressed with higher-order spatial interpolation, but this approach fails for this problem due to the presence of large spatial gradients. In [17] it is handled by reducing the size of the time step and the order of the integrator, but the effect on the order of temporal accuracy was not explored. We plan to investigate these issues 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 2.1.6, Argonne National Laboratory, 2004.
- [3] C. BALDWIN, P. N. BROWN, R. FALGOUT, F. GRAZIANI, AND J. JONES, *Iterative linear solvers in a 2D radiation-hydrodynamics code: methods and performance*, J. Comput. Phys., 154 (1999), pp. 1–40.
- [4] 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.
- [5] R. BLAHETA, *Adaptive composite grid methods for problems of plasticity*, Math. Comp. Sim., 50 (1999), pp. 123–134.
- [6] 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.
- [7] P. N. BROWN, *A local convergence theory for combined inexact-Newton/finite-difference projection methods*, SIAM J. Numer. Anal., 24 (1987), pp. 407–434.
- [8] R. S. DEMBO, S. C. EISENSTAT, AND T. STEIHAUG, *Inexact Newton methods*, SIAM J. Numer. Anal., 19 (1982), pp. 400–408.
- [9] J. E. DENNIS JR. AND R. B. SCHNABEL, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Prentice-Hall, Inc., 1983.
- [10] S. C. EISENSTAT AND H. F. WALKER, *Globally convergent inexact Newton methods*, SIAM J. Optimization, 4 (1994), pp. 393–422.
- [11] 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.
- [12] R. D. FALGOUT AND U. M. YANG, *hypre: 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.
- [13] R. D. HORNUNG AND S. KOHN, *Managing application complexity in the SAMRAI object-oriented framework*, Concurrency Comput.: Pract. Exp., 14 (2002), pp. 347–368.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] 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 Computational Science and Engineering, Heidelberg, 2003, Springer-Verlag.
- [18] R. E. MARSHAK, *Effect of radiation on shock wave behavior*, Physics of Fluids, 1 (1958), pp. 24–29.
- [19] S. F. MCCORMICK, *Multilevel Adaptive Methods for Partial Differential Equations*, SIAM, Philadelphia, PA, 1989.

- [20] S. F. McCORMICK AND J. W. THOMAS, *The Fast Adaptive Composite grid (FAC) method for elliptic equations*, Math. Comp., 46 (1986), pp. 439–456.
- [21] D. MIHALAS AND B. WEIBEL-MIHALAS, *Foundations of Radiation Hydrodynamics*, Dover Publications, Inc., Mineola, NY, 1999.
- [22] 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.
- [23] 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.
- [24] 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.
- [25] M. PERNICE AND R. D. HORNUNG, *Newton-Krylov-FAC methods for problems discretized on locally refined grids*, Comput. Visual. Sci., (2005). (to appear).
- [26] 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.
- [27] S. SCHAFFER, *A semicoarsening multigrid method for elliptic partial differential equations with highly discontinuous and anisotropic coefficients*, SIAM J. Sci. Comput., 20 (1999), pp. 228–242.
- [28] L. STALS, *Comparison of non-linear solvers for the solution of radiation transport equations*, Elec. Trans. Num. Anal., 15 (2003), pp. 78–93.
- [29] 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.
- [30] ———, *Runge-Kutta methods and local uniform grid refinement*, Math. Comp., 60 (1993), pp. 591–616.
- [31] K.-H. WINKLER, M. NORMAN, AND D. MIHALAS, *Implicit adaptive-grid radiation-hydrodynamics*, in Multiple Time Scales, Academic Press, 1985.
- [32] A. M. WINSLOW, *Extensions of asymptotic neutron diffusion theory*, Nucl. Sci. and Eng., 32 (1968), pp. 101–110.