

Preconditioning Strategies for Fully Implicit Radiation Diffusion with Material- Energy Transfer

Peter N. Brown and Carol S. Woodward

This article will appear in
SIAM Journal on Scientific Computing

May 2000
rev. November 2000

U.S. Department of Energy

Lawrence
Livermore
National
Laboratory

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

This report has been reproduced
directly from the best available copy.

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information
P.O. Box 62, Oak Ridge, TN 37831
Prices available from (423) 576-8401
<http://apollo.osti.gov/bridge/>

Available to the public from the
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd.,
Springfield, VA 22161
<http://www.ntis.gov/>

OR

Lawrence Livermore National Laboratory
Technical Information Department's Digital Library
<http://www.llnl.gov/tid/Library.html>

PRECONDITIONING STRATEGIES FOR FULLY IMPLICIT RADIATION DIFFUSION WITH MATERIAL-ENERGY TRANSFER*

PETER N. BROWN[†] AND CAROL S. WOODWARD[‡]

Abstract. In this paper, we present a comparison of four preconditioning strategies for Jacobian systems arising in the fully implicit solution of radiation diffusion coupled with material energy transfer. The four preconditioning methods are: block Jacobi, Schur complement, and operator splitting approaches that split the preconditioner solve into two steps. One splitting method includes the coupling of the radiation and material fields that appears in the matrix diagonal in the first solve, the other method puts this coupling into the second solve. All preconditioning approaches use multigrid methods to invert blocks of the matrix formed from the diffusion operator. The Schur complement approach is clearly seen to be the most effective for a large range of weightings between the diffusion and energy coupling terms. In addition, tabulated opacity studies were conducted where, again, the Schur preconditioner performed well. Lastly, a parallel scaling study was done showing algorithmic scalability of the Schur preconditioner.

Keywords: preconditioning, Newton-Krylov, operator splitting, nonlinear solvers, radiation diffusion

AMS(MOS) subject classification: 65F10, 65N40, 65N55, 65H10

1. Introduction. In this paper, we present a comparison of preconditioners for a new numerical approach to the solution of very large-scale radiation diffusion problems. In this model, energy can be transferred to a material through coupling terms in both the radiation and material energy equations. These problems are important in modeling photon energy progression through an optically thick regime, a situation common in some laser and stellar fusion applications. Traditionally, solutions for these problems have been developed using operator split and time-lag techniques to reduce the coupled system of nonlinear equations to the solution of a series of linear problems. These solution techniques, however, lead to requirements of unacceptably small time steps. Furthermore, as computers have become faster, researchers have attempted to simulate larger problems, despite existing solution methods that did not scale well for increased numbers of unknowns.

For these reasons, we have developed a solution method for solving radiation diffusion problems formulated in a fully implicit manner [5]. The fully implicit formulation allows larger time steps to be taken without sacrificing accuracy. Furthermore, recent work in iterative methods has provided computational scientists with new tools for solving these problems — tools that scale well to large numbers of unknowns. In order to solve this fully implicit formulation, we employ ODE time integration techniques which then require an implicit solve for the solution at each time step. We use an inexact Newton method for these solves, with a preconditioned Krylov method for solving the linear Jacobian systems that arise within the Newton iterations. The Newton method provides fast nonlinear convergence, and the Krylov method gives a robust linear solver.

We consider four methods for solving the Jacobian preconditioning step in the Krylov method. All four schemes neglect the nonlinearity in the diffusion coefficient. The four methods primarily differ in how they approximate the coupling between the

* This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under contract number W-7405-Eng-48.

[†]Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, Livermore, CA 94551; E-mail: pnbrown@llnl.gov.

[‡]Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, Livermore, CA 94551; E-Mail: cswoodward@llnl.gov.

radiation and material energies. The first scheme is just to use the block diagonal part of the Jacobian matrix, thereby neglecting the majority of the coupling between the two fields. The second scheme is to factor the Jacobian and use a Schur complement preconditioner. The remaining methods use operator splitting approaches to split the preconditioner solve into two steps. One method includes the coupling of the radiation and material fields that appears in the matrix diagonal in the first solve, the other method puts this coupling into the second solve. All preconditioning approaches use multigrid methods to invert blocks of the matrix formed from the diffusion operator.

Recent work by Mousseau, Knoll and Rider [15] has considered the fully implicit formulation of radiation diffusion using an operator splitting preconditioner similar to the second splitting method mentioned above. They saw this preconditioner to be quite effective in solving one and two dimensional problems. In addition, earlier work by Knoll, Rider and Olson [14] showed that the fully implicit form of the one dimensional radiation diffusion problem gave greater accuracy in shorter times than did traditional methods. In previous work, we found that the block Jacobi preconditioning method was effective for test problems in three dimensions on parallel computers [5]. Further work has shown that this method is not as effective as we would like for cases where the material coupling dominates the diffusion operator. In this paper, we will compare the four preconditioning methods mentioned above on three-dimensional test problems. We will show an effective, fully implicit, parallel solution strategy for these problems.

The rest of this paper is organized as follows. In the next section, we present the mathematical models we are considering for this work. In Sections 3 and 4 we discuss the spatial and temporal discretization techniques used, and in Section 5 we detail the four preconditioning methods compared and show analysis indicating what qualitative behavior we expect of each for different problem parameters. We briefly discuss multigrid methods in Section 6. In Section 7 we give some numerical results showing algorithm performance on problems with various degrees of difficulty and in parallel on problems in three dimensions. Section 8 provides some concluding remarks.

2. Problem Formulation. Our model for radiation diffusion is a simplification of the full radiation transport equation given in [16]. We assume isotropic radiation (no angular dependence), Fick's Law of diffusion, no scattering effects, and that the photon energy is Planckian, and then integrate the transport equation in frequency to get the diffusion model [2],

$$(2.1) \quad \frac{\partial E_R}{\partial t} = \nabla \cdot \left(\frac{c}{3\rho\kappa_R(T_R)} \nabla E_R \right) + c\rho\kappa_P(T_M) \cdot (aT_M^4 - E_R),$$

where $E_R(\mathbf{x}, t)$ is the radiation energy density ($\mathbf{x} = (x, y, z)$), $T_M(\mathbf{x}, t)$ is the material temperature, $\rho(\mathbf{x})$ is the material density, c is the speed of light, and $a = 4\sigma/c$ where σ is the Stephan-Boltzmann constant. The Rosseland opacity, κ_R , is a nonlinear function of the radiation temperature, T_R , which is defined by the relation $E_R = aT_R^4$. The Planck opacity, κ_P , is a nonlinear function of material temperature, T_M , which is related to the material energy through an equation of state, $E_M = EOS(T_M)$. In many instances, the two opacities will take on similar values.

We also consider a spatially dependent source term in this equation expressing sources or sinks in the radiation field given by,

$$(2.2) \quad \chi(\mathbf{x})caT_{\text{source}}^4,$$

where T_{source} is a given source temperature and $\chi(\mathbf{x})$ is a function of the spatial variable \mathbf{x} . Computed solutions to (2.1) may result in photon velocities which exceed the speed of light. To prevent this nonphysical phenomena, a flux-limiter is often added to the diffusion term [2]. We use a flux-limiter of the form $\frac{\|\nabla E_R\|}{E_R}$, where the norm $\|\cdot\|$ is just the l^2 norm of the gradient vector.

The resulting radiation diffusion equation we use as our model is,

$$(2.3) \quad \frac{\partial E_R}{\partial t} = \nabla \cdot \left(\frac{c}{3\rho\kappa_R(T_R) + \frac{\|\nabla E_R\|}{E_R}} \nabla E_R \right) + c\rho\kappa_P(T_M) \cdot (aT_M^4 - E_R) + \chi(\mathbf{x})caT_{\text{source}}^4.$$

This equation is coupled to an equation expressing conservation of material energy given by,

$$(2.4) \quad \frac{\partial E_M}{\partial t} = -c\rho\kappa_P(T_M) \cdot (aT_M^4 - E_R).$$

We will focus on the development of solution methods for the system (2.3)–(2.4) in what follows.

3. Solution Method. We apply a method of lines approach to the solution of (2.3)–(2.4). The spatial discretization used is as follows. We use a tensor product grid with N_x , N_y and N_z cells in the x , y and z directions, respectively. Define $E_{R,i,j,k}(t) \approx E_R(\mathbf{x}_{i,j,k}, t)$ and $E_{M,i,j,k}(t) \approx E_M(\mathbf{x}_{i,j,k}, t)$, with $\mathbf{x}_{i,j,k} = (x_i, y_j, z_k)$ the cell centers. Next, define

$$\mathbf{E}_R \equiv \begin{pmatrix} E_{R,1,1,1} \\ \vdots \\ E_{R,N_x,N_y,N_z} \end{pmatrix} \text{ and } \mathbf{E}_M \equiv \begin{pmatrix} E_{M,1,1,1} \\ \vdots \\ E_{M,N_x,N_y,N_z} \end{pmatrix}.$$

We employ a cell-centered finite difference scheme over the computational mesh and write our discrete equations in terms of a discrete diffusion operator given by $\mathbf{L}(\mathbf{E}_R) \equiv (L_{1,1,1}(\mathbf{E}_R), \dots, L_{N_x,N_y,N_z}(\mathbf{E}_R))^T$, where

$$(3.1) \quad L_{i,j,k}(\mathbf{E}_R) = \left(\frac{c}{3\rho_{i+1/2,j,k}\kappa_{R,i+1/2,j,k} + \frac{\|\nabla E_R\|_{i+1/2,j,k}}{E_{R,i+1/2,j,k}}} \frac{E_{R,i+1,j,k} - E_{R,i,j,k}}{x_{i+1} - x_i} - \frac{c}{3\rho_{i-1/2,j,k}\kappa_{R,i-1/2,j,k} + \frac{\|\nabla E_R\|_{i-1/2,j,k}}{E_{R,i-1/2,j,k}}} \frac{E_{R,i,j,k} - E_{R,i-1,j,k}}{x_i - x_{i-1}} \right) / (x_{i+1/2} - x_{i-1/2}) + y \text{ and } z \text{ terms},$$

and a local operator $\mathbf{S}(\mathbf{E}_R, \mathbf{E}_M) \equiv (S_{1,1,1}(\mathbf{E}_R, \mathbf{E}_M), \dots, S_{N_x,N_y,N_z}(\mathbf{E}_R, \mathbf{E}_M))^T$, where

$$(3.2) \quad S_{i,j,k}(E_{R,i,j,k}, E_{M,i,j,k}) = c\rho_{i,j,k}\kappa_{P,i,j,k} (aT_{M,i,j,k}^4 - E_{R,i,j,k}).$$

Thus, our discrete scheme is to find $\mathbf{E}_R(t)$ and $\mathbf{E}_M(t)$ such that,

$$(3.3) \quad \frac{d\mathbf{E}_R}{dt} = \mathbf{L}(\mathbf{E}_R) + \mathbf{S}(\mathbf{E}_R, \mathbf{E}_M) + \mathbf{Q},$$

$$(3.4) \quad \frac{d\mathbf{E}_M}{dt} = -\mathbf{S}(\mathbf{E}_R, \mathbf{E}_M),$$

where $\mathbf{Q} \equiv caT_{\text{source}}^4(\chi(\mathbf{x}_{1,1,1}), \dots, \chi(\mathbf{x}_{N_x, N_y, N_z}))^T$. The system (3.3)–(3.4) is an ODE system and our time integration technique will be based on ODE time integration methods.

4. Time Integration. We have developed a 3D, parallel simulator that employs a fully implicit formulation and solution process for the radiation diffusion model in (3.3)–(3.4) and with which algorithms for radiation diffusion can be studied. In order to allow for accurate time-stepping as well as larger steps than what traditional methods allow, we use an ODE time integrator to handle the temporal discretization. This simulator uses the parallel ODE solver, PVODE [8], developed at Lawrence Livermore National Laboratory and based on the VODPK package [7]. PVODE employs the fixed leading coefficient variant of the Backward Differentiation Formula (BDF) method [4, 12] and allows for variation in the order of the time discretization as well as in the time step size. Time step sizes are chosen to minimize the local truncation error, and thus give a solution that obeys a user-specified accuracy bound.

This time integration technique leads to a coupled, nonlinear system of equations that must be solved at each time step. For the solution of this system, we use an inexact Newton-Krylov method with Jacobian-vector products approximated by finite differences. As the methods in PVODE are Predictor-Corrector in nature, an explicit predictor is used for an initial guess in the nonlinear solve.

In the methods discussed above, we use the scaling technique incorporated into PVODE. Thus, we include an absolute tolerance (ATOL) for each unknown and a relative tolerance (RTOL) which is applied to all unknowns. These tolerances are then used to form a weight which is applied to each solution component during the time step from t_{n-1} to t_n . This weight is given as,

$$(4.1) \quad w_i = RTOL|y_{n-1}^i| + ATOL_i,$$

and is also used to weight a root mean square norm which is applied to all error-like vectors within the solution process. This scaling gives each vector component equal weight when calculating norms. For our application, we supply two absolute tolerances, one to be used with the radiation energy unknowns and one to be used with the material energy unknowns.

5. Preconditioners. The use of Newton-Krylov methods necessitates the use of preconditioning, and we consider several strategies. Before detailing the four preconditioning strategies we compare in this work, we consider the content and structure of the Jacobian matrix we are trying to precondition. We formulate our system of ODEs as $\dot{y} = f(t, y)$, set $y = (\mathbf{E}_R^T, \mathbf{E}_M^T)^T$, and then form f using the right-hand-sides of (3.3)–(3.4). The Jacobian matrices used in the Newton method are of the general form $F'(y) = (I - \gamma J)$, where $J = \partial f / \partial y$ is the Jacobian of the nonlinear function f , and the parameter $\gamma \equiv \Delta t \beta$ with Δt the current time step value and β a coefficient depending on the order of the BDF method. Recalling the definitions of the discrete divergence and source operators, defined in (3.1) and (3.2), the block form of the Jacobian of f is

$$J = \begin{pmatrix} \partial \mathbf{L} / \partial \mathbf{E}_R + \partial \mathbf{S} / \partial \mathbf{E}_R & \partial \mathbf{S} / \partial \mathbf{E}_M \\ -\partial \mathbf{S} / \partial \mathbf{E}_R & -\partial \mathbf{S} / \partial \mathbf{E}_M \end{pmatrix} = \begin{pmatrix} A + G & B \\ C & D \end{pmatrix},$$

where $A = \partial \mathbf{L} / \partial \mathbf{E}_R$, $G = \partial \mathbf{S} / \partial \mathbf{E}_R$, $B = \partial \mathbf{S} / \partial \mathbf{E}_M$, $C = -\partial \mathbf{S} / \partial \mathbf{E}_R$ and $D = -\partial \mathbf{S} / \partial \mathbf{E}_M$. We note that G, B, C and D are all diagonal matrices.

Since Jacobian approximations can be expensive to compute, the preconditioner is not updated with every Newton iteration. Preconditioner updates occur only when the Newton iteration fails to converge, 20 time steps pass without an update, or when there is a significant change in the time step size and order of the ODE method.

On close inspection of the nonlinear diffusion operator $\mathbf{L}(\mathbf{E}_R)$, we can write

$$\mathbf{L}(\mathbf{E}_R) = \hat{\mathbf{L}}(\mathbf{E}_R)\mathbf{E}_R,$$

where $\hat{\mathbf{L}}$ is a nonlinear matrix-valued function of \mathbf{E}_R . In all of our preconditioning strategies, we neglect the nonlinearity in the diffusion term and use the approximation

$$A = \partial\mathbf{L}(\hat{\mathbf{E}}_R)/\partial\mathbf{E}_R \approx \hat{\mathbf{L}}(\hat{\mathbf{E}}_R) \equiv \tilde{A},$$

where $\partial\mathbf{L}(\hat{\mathbf{E}}_R)/\partial\mathbf{E}_R$ is the Jacobian of \mathbf{L} evaluated at a radiation energy, $\hat{\mathbf{E}}_R$. The size of the neglected term is related to the derivatives of the Rosseland opacity and the flux-limiter. Our motivation for neglecting this term arises from the fact that $-\tilde{A}$ is symmetric and positive definite, whereas A has a first order term that leads to nonsymmetries in its discretized form. This first order term includes derivatives of the Rosseland opacity and flux-limiter. By neglecting the nonlinearity, this term is removed from the preconditioner, and a symmetric approximation results. In addition, calculation of \tilde{A} is much cheaper than for A , as no derivatives of the flux-limiter need be computed. We also note that computation of the derivative of the flux-limiter may lead to numerical errors if $\nabla\mathbf{E}_R$ approaches 0.

Our preconditioning strategies differ in how they approximately solve systems with the matrix,

$$M = I - \gamma \begin{pmatrix} \tilde{A} + G & B \\ C & D \end{pmatrix},$$

In all cases, multigrid methods are used to invert the \tilde{A} blocks. We will discuss the specifics of the multigrid scheme after detailing the four strategies.

In the following discussion, we will be examining the preconditioner performance responses to changes in the κ_P, κ_R and Δt parameters in the problem. For these discussions, it will be useful to consider the same radiation diffusion model as given in (3.3)-(3.4) but neglecting flux-limiting and assuming κ_P and κ_R constant. In this case, we see that $\tilde{A} = O(1/\kappa_R)$ and that B, C , and D are all $O(\kappa_P)$.

5.1. Block Jacobi. Our first strategy is to approximate the Jacobian system with,

$$M_{\text{Jacobi}} = I - \gamma \begin{pmatrix} \tilde{A} + G & 0 \\ 0 & D \end{pmatrix}.$$

This method effectively neglects the coupling between the radiation and material energy fields.

We now examine the error, $Err_{\text{Jacobi}} = (I - \gamma J) - M_{\text{Jacobi}}$. We see that

$$Err_{\text{Jacobi}} = \gamma \begin{pmatrix} \tilde{A} - A & -B \\ -C & 0 \end{pmatrix}.$$

Thus, the error for this block Jacobi preconditioning strategy is $O(\Delta t)$. In addition, for the non-flux-limited, constant opacity case, we see that $\tilde{A} = A$ and the error is $O(\Delta t \times \kappa_P)$ since both B and C are $O(\kappa_P)$. When $\Delta t \times \kappa_P$ gets large, we would expect this preconditioner to perform poorly.

5.2. Schur Complement. Our second preconditioning strategy is to factor the matrix

$$\begin{pmatrix} P & Q \\ R & T \end{pmatrix} \equiv \begin{pmatrix} I - \gamma(\tilde{A} + G) & -\gamma B \\ -\gamma C & I - \gamma D \end{pmatrix} = M$$

into the following,

$$M_{\text{Schur}} = \begin{pmatrix} I & QT^{-1} \\ 0 & I \end{pmatrix} \begin{pmatrix} P - QT^{-1}R & 0 \\ 0 & T \end{pmatrix} \begin{pmatrix} I & 0 \\ T^{-1}R & I \end{pmatrix}.$$

Letting $S = P - QT^{-1}R$, we write the solution to $M_{\text{Schur}}x = b$ as

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} S^{-1}(b_1 - QT^{-1}b_2) \\ T^{-1}(-Rx_1 + b_2) \end{pmatrix}.$$

The error in this preconditioner, $Err_{\text{Schur}} = (I - \gamma J) - M_{\text{Schur}}$, is given by,

$$Err_{\text{Schur}} = \gamma \begin{pmatrix} \tilde{A} - A & 0 \\ 0 & 0 \end{pmatrix}.$$

If the Schur complement, S , is exactly inverted, there will be no error associated with this preconditioner for the non-flux-limited, constant opacity case. In addition, because D and hence T is diagonal, there is no penalty associated with inverting T for every iteration of a method that inverts S , as there would be if a material energy diffusion term were added to the equations. Also note that S is formed by modifying the diagonal of P , so we can still employ multigrid methods to invert this Schur complement, as we would to invert the \tilde{A} matrix.

5.3. Matrix Split. Our third strategy is motivated by a preconditioner developed in [11] where a splitting of the Jacobian matrix is used. Our preconditioner is written as,

$$(5.1) \quad M_{\text{matrix_split}} = (I - \gamma J_{\text{diag}})(I - \gamma J_{\text{border}}),$$

where

$$J_{\text{diag}} = \begin{pmatrix} \tilde{A} + G & 0 \\ 0 & 0 \end{pmatrix}, \quad \text{and} \quad J_{\text{border}} = \begin{pmatrix} 0 & B \\ C & D \end{pmatrix}.$$

Solving systems of the form $M_{\text{matrix_split}}x = b$ requires two steps. The first step consists of a solve with the system $(I - \gamma J_{\text{diag}})y = b$, and the second step consists of a solve with the system $(I - \gamma J_{\text{border}})x = y$. Multigrid methods can be used to solve the first system.

To see how the second system can be easily inverted, we consider a re-ordering of the unknowns and equations of the system, so that unknowns are first ordered by space and then by energy type for each spatial point. Equations are re-ordered similarly. With this new ordering, $I - \gamma J_{\text{border}}$ will be a block diagonal matrix with 2×2 blocks. Each of these blocks can be written as,

$$(I - \gamma J_{\text{border}})_i = \begin{pmatrix} 1 & b_i \\ c_i & d_i \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ c_i & 1 \end{pmatrix} \begin{pmatrix} 1 & b_i \\ 0 & d' \end{pmatrix},$$

where $d' = d_i - b_i c_i$. Solutions of the second step in the application of $M_{\text{matrix_split}}$ are easily obtained with this factorization.

The error in this preconditioner, $Err_{\text{matrix_split}} = (I - \gamma J) - M_{\text{matrix_split}}$, is given by,

$$(5.2) \quad Err_{\text{matrix_split}} = \gamma \begin{pmatrix} -A + \tilde{A} & 0 \\ 0 & 0 \end{pmatrix} - \gamma^2 \begin{pmatrix} 0 & (\tilde{A} + G)B \\ 0 & 0 \end{pmatrix}.$$

For the non-flux-limited, constant coefficient case, $\tilde{A} = A$. The term $(\tilde{A} + G)B$ is of order $O((\frac{1}{\kappa_R} + \kappa_P)\kappa_P) \approx O(\kappa_P^2)$. Thus, this preconditioner has error $O((\Delta t)^2 \times \kappa_P^2)$. So, as $\Delta t \times \kappa_P$ gets large, we would expect this preconditioner's effectiveness to deteriorate rapidly.

5.4. Operator Split. Our last strategy is motivated by a preconditioner developed in [6] where an operator splitting of the Jacobian operator is used to split the preconditioning into two steps. This preconditioner is very similar to the previous strategy except that the G term is part of the second step, rather than the first. This preconditioner is written as,

$$(5.3) \quad M_{\text{operator_split}} = (I - \gamma J_{\text{diff}})(I - \gamma J_{\text{coupling}}),$$

where

$$J_{\text{diff}} = \begin{pmatrix} \tilde{A} & 0 \\ 0 & 0 \end{pmatrix}, \quad \text{and} \quad J_{\text{coupling}} = \begin{pmatrix} G & B \\ C & D \end{pmatrix}.$$

Again, the first step consists of a solve with the system $(I - \gamma J_{\text{diff}})y = b$, and the second step consists of a solve with the system $(I - \gamma J_{\text{coupling}})x = y$. Multigrid methods can be used to solve the first system.

The second system here is also easily inverted by a simple LU decomposition of a re-ordered problem. However, the 2×2 blocks have a non-identity upper left entry, so that the decomposition is:

$$(I - \gamma J_{\text{coupling}})_i = \begin{pmatrix} g_i & b_i \\ c_i & d_i \end{pmatrix} = \begin{pmatrix} g_i & 0 \\ c_i & 1 \end{pmatrix} \begin{pmatrix} 1 & b_i/g_i \\ 0 & d' \end{pmatrix},$$

where $d' = d_i - (b_i c_i)/g_i$.

The error associated with this preconditioner, $Err_{\text{operator_split}} = (I - \gamma J) - M_{\text{operator_split}}$, is given by,

$$(5.4) \quad Err_{\text{operator_split}} = \gamma \begin{pmatrix} -A + \tilde{A} & 0 \\ 0 & 0 \end{pmatrix} - \gamma^2 \begin{pmatrix} \tilde{A}G & \tilde{A}B \\ 0 & 0 \end{pmatrix}.$$

Again, for the non-flux-limited, constant coefficient case, $\tilde{A} = A$ and the terms $\tilde{A}G$ and $\tilde{A}B$ are of order $O(\frac{\kappa_P}{\kappa_R})$. Thus, this preconditioner has $O((\Delta t)^2 \times \frac{\kappa_P}{\kappa_R})$ error. Thus, if $\kappa_R \approx \kappa_P$, we would expect this preconditioner to show minimal deterioration in effectiveness as the opacities get large.

Note that $G = -C$ and $B = -D$ so that as $\Delta t \rightarrow \infty$, the second stage of this preconditioner $I - \gamma J_{\text{coupling}}$ becomes singular. As a result, we may expect this preconditioner to deteriorate for extremely large time steps.

6. Multigrid Methods. The Rosseland opacity will exhibit large changes where material interfaces exist in the domain. The temperature dependence gives rise to large value changes as well. These changes imply that the problem can be very heterogeneous. As a result, to invert matrix blocks formed from the diffusion operator,

we use a multigrid method designed to handle large changes in problem coefficients. In particular, we use 1 V-cycle of the PFMG algorithm developed by Ashby and Falgout [1] as our multigrid solver. Other multigrid methods have been developed for highly heterogeneous problems. A comparison of PFMG and another of these methods can be found in [13]. We use PFMG here because it is fast and scales extremely well. More information about multigrid methods can be found in [3].

7. Numerical Results. To understand how these preconditioning strategies perform for problems with varying degrees of difficulty, we performed a number of studies. Our first study looked at the effects of changing the relative weighting of the diffusion and coupling terms in the radiation equation by setting the two opacities constant and equal and then investigating preconditioner responses to increasing the value. Our second study looked at the effects of tabulated opacities, as are currently used in applications of interest, on these preconditioners. Lastly, we performed a parallel scaling study with the most effective preconditioner to verify algorithmic scalability of the solution method.

For all runs, we used the PVODE package default settings with the following exception. Some of our tests led to nonstable solutions with higher order methods, so we have limited the ODE method order to 2 for all cases, except where explicitly noted. We are presently looking into why these situations occur. Some work has been done in the area of avoiding these sorts of instabilities [10], and we will investigate its applicability here. Note that the default setting for the maximum number of GMRES iterations for PVODE is 5. No restart is performed.

7.1. Constant Opacity Results. In our first study, we set the Rosseland and Planck opacities equal to a single parameter, κ . We then changed the value of this parameter from $\kappa = 1$ to $\kappa = 100,000$.

For this problem, we set $E_M = T_M$. The system (2.3)–(2.4) is solved on the box $\mathcal{D} \equiv \{\mathbf{x} = (x, y, z) : 0 \leq x, y, z \leq 1\text{cm}\}$. The function $\chi(\mathbf{x})$ in (2.3) is defined by

$$(7.1) \quad \chi(\mathbf{x}) = \begin{cases} 1, & \text{if } 0.4 \leq x, y, z \leq 0.6, \text{ and} \\ 0, & \text{otherwise.} \end{cases}$$

The parameter T_{source} was $3,481,440 \text{ }^\circ\text{K}$ (approximately 300eV), and the initial conditions were taken as $E_R = aT_{R,0}^4$ and $E_M = T_{M,0}$, where $T_{R,0} = T_{M,0} = 300 \text{ }^\circ\text{K}$. Dirichlet boundary conditions were consistent with the initial conditions. The density was taken to be 1.0g/cc . The spatial grid was uniform with $N_x = N_y = N_z = 20$.

We examined solver statistics at 50 intervals of about 0.1002s with a final simulation time of about 5.0104s . Flux-limiting was applied to the problem as discussed above. We asked for a relative tolerance on each solution component of 10^{-4} and an absolute tolerance on each of the energies of 200.

Figures 7.1 and 7.2 show the solutions over the x -line for $y = z = 0.475\text{cm}$ at times 0.1002s and 1.5031s for the value $\kappa = 100\text{cm}^2/\text{g}$. The time 0.1002s is the first output time we recorded. The solutions have felt the effects of the source at this point. By 1.5031s , the energies have increased due to the source at the domain center. We also see the effects of the diffusion operator spreading out the radiation energy.

Figures 7.3 and 7.4 show the solutions at these times but for $\kappa = 10^5\text{cm}^2/\text{g}$. Here the “bump” in the domain center is much more pronounced. This difference from the lower κ case is due to the increased coupling between the two energy fields. For higher values of κ , the radiation energy diffuses much less, and more of its energy is transferred to the material. For these cases, the local physics is clearly dominating the calculation.

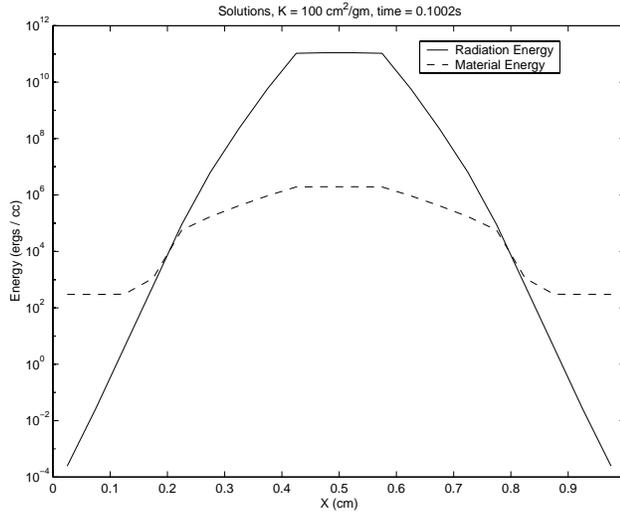


FIG. 7.1. Radiation and material energies for $\kappa = 100$, time = 0.1002s, $y = z = 0.475\text{cm}$.

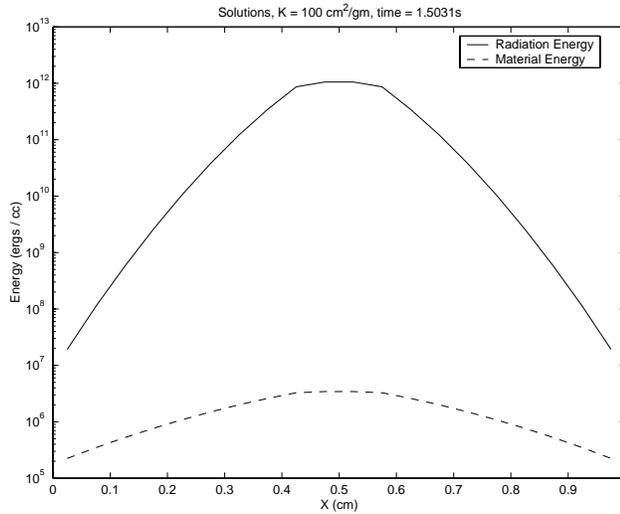


FIG. 7.2. Radiation and material energies for $\kappa = 100$, time = 1.5031s, $y = z = 0.475\text{cm}$.

Table 7.1 shows the cumulative solver statistics for these runs. In this and subsequent tables,

- S = Schur Preconditioner,
- BJ = Block Jacobi Preconditioner,
- MS = Matrix Split Preconditioner, and
- OS = Operator Split Preconditioner,

and the statistical counters are

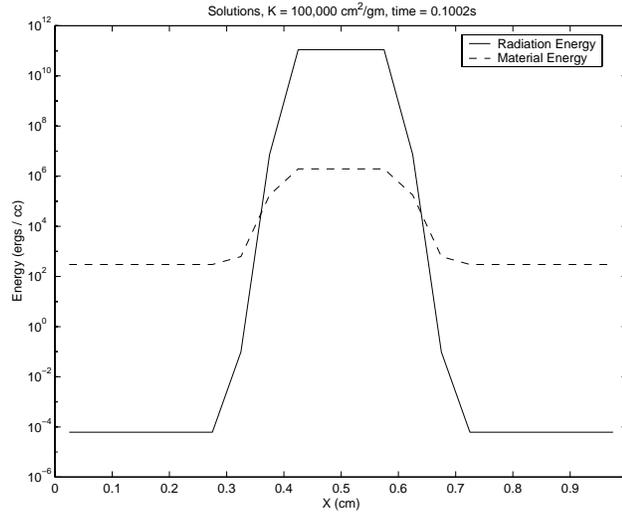


FIG. 7.3. Radiation and material energies for $\kappa = 10^5$, time = 0.1002s, $y = z = 0.475$ cm.

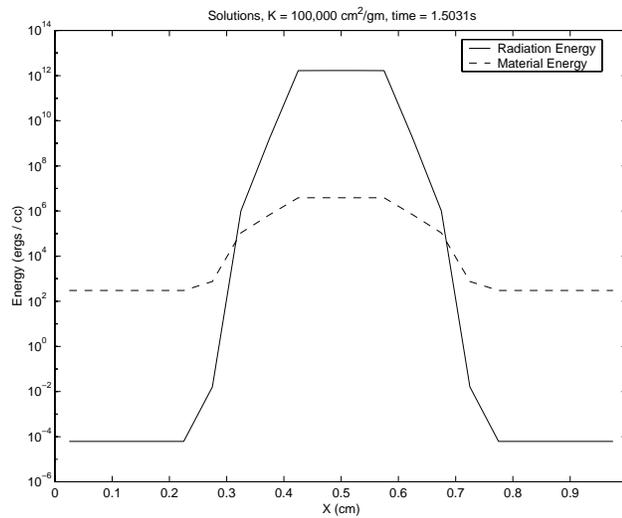


FIG. 7.4. Radiation and material energies for $\kappa = 10^5$, time = 1.5031s, $y = z = 0.475$ cm.

- NST = total number of time steps,
- NNI = total number of nonlinear iterations,
- NLI = total number of linear iterations,
- NFE = total number of $f(t, y)$ evaluations,
- NPE = total number of preconditioner evaluations,
- NPS = total number of preconditioner solves,
- HU = step size that was used on the last step (scaled by c),
- RT = run time in seconds,
- NCFN = total number of nonlinear convergence failures, and
- NCFL = total number of linear convergence failures.

We see that the Schur preconditioner is consistently performing better and faster than

TABLE 7.1
Solver statistics for constant opacity problem

PC	κ	NST	NNI	NLI	RT	NCFN	NCFL
S	1	1,831	1,884	2,168	2,648	0	0
BJ	1	1,830	1,886	2,167	2,711	0	0
MS	1	1,830	1,884	2,171	2,617	0	0
OS	1	1,830	1,884	2,170	2,626	0	0
S	100	1,782	1,875	3,065	3,145	0	0
BJ	100	1,778	1,874	3,273	3,366	0	0
MS	100	1,786	1,894	3,141	3,213	3	17
OS	100	1,782	1,877	2,995	3,129	0	0
S	10,000	660	723	1,271	1,315	0	0
BJ	10,000	694	831	2,299	2,116	8	88
MS	10,000	2,828	4,939	19,343	17,371	930	2,893
OS	10,000	670	733	2,285	1,950	0	134
S	100,000	424	474	786	846	0	0
BJ	100,000	1,288	2,084	6,010	5,822	88	439
MS	100,000	2,359	3,787	14,780	12,995	518	2,090
OS	100,000	650	949	2,912	2,609	8	214

the others. As the coupling term between the radiation and material grows in weight relative to the diffusion term, we see that the matrix split preconditioner is the first to show significant signs of struggle. This degradation in the matrix split preconditioner performance is expected since its error was on the order of κ_P^2 . However, all but the Schur preconditioner are struggling for $\kappa = 10^4$ and $\kappa = 10^5$. The operator split preconditioner shows the second best performance, which is also expected since its error is on the order of κ_P/κ_R . For lower opacity values, the preconditioners all perform fairly well.

Figure 7.5 shows the cumulative numbers of nonlinear iterations taken by each of the preconditioners for the 50 output times for the two cases of $\kappa = 100\text{cm}^2/g$ and $\kappa = 10^5\text{cm}^2/g$. For the lower κ value, all preconditioners result in about the same number of nonlinear iterations at each time step. For the higher κ value, however, the preconditioners show distinctly different performances. During the transition to steady state, the matrix split preconditioner has the most degradation with the block Jacobi preconditioner also showing degradation, but less. These results bear out the analysis given above in that these two preconditioner have the strongest dependence on the κ value with the matrix split the strongest. After the solution gets close to steady state, however, the four preconditioners all require very few nonlinear iterations to resolve the physics. Similar results bear out for the linear iteration counts.

7.2. Tabular Opacities Results. In this section we give results of using the above preconditioners on several problems involving the use of tabular opacities. We use the LEOS [9] package to give the Rosseland and Planck opacities as nonlinear functions of the radiation temperature T_R and material temperature T_M , respectively. The system (2.3)–(2.4) is solved on the box $\mathcal{D} \equiv \{\mathbf{x} = (x, y, z) : 0 \leq x, y, z \leq 1\text{cm}\}$ with Dirichlet boundary conditions. The function $\chi(\mathbf{x})$ in (2.3) is defined by

$$(7.2) \quad \chi(\mathbf{x}) = \begin{cases} 1, & \text{if } 0.3 \leq x, y, z \leq 0.7, \text{ and} \\ 0, & \text{otherwise.} \end{cases}$$

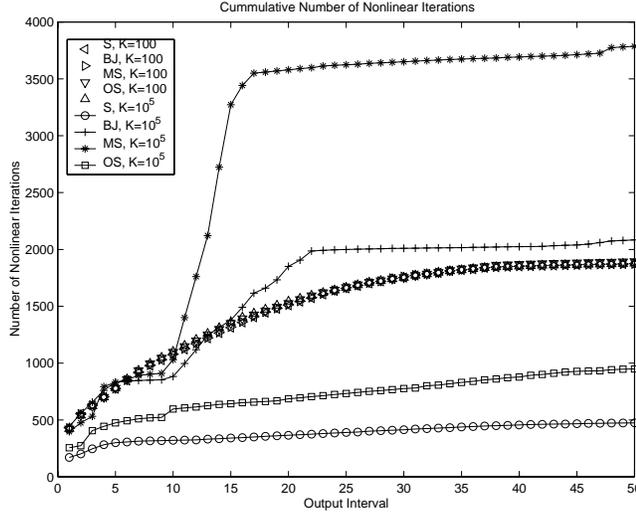


FIG. 7.5. Cumulative nonlinear iteration counts for all four preconditioners for $\kappa = 100\text{cm}^2/\text{g}$ and $\kappa = 10^5\text{cm}^2/\text{g}$.

The parameter T_{source} was 3,481,440 °K (approximately 300eV), and the initial conditions were taken as $E_R = aT_{R,0}^4$ and $E_M = EOS(\rho, T_{M,0})$, where $T_{R,0} = T_{M,0} = 116,100$ °K (approximately 10eV), and $EOS(\rho, T_M)$ is the equation of state function in the LEOS package giving E_M as a function of ρ and T_M . The Dirichlet boundary values for E_R and E_M are taken to be consistent with the initial conditions. The material used was carbon at a reference density of $\rho = 1.05\text{g}/\text{cc}$. The spatial grid was uniform with $N_x = N_y = N_z = 20$.

For this problem, the time behavior consists of an initial transient in which the material heats up in the region of the source (from 0 to .01 microseconds), followed by a radiation front traveling to the boundary (continuing to .41 microseconds), and then a final phase in which the solution approaches a steady state (integrated to about 1.33 microseconds). Figures 7.6 and 7.7 show the solutions plotted on the line $y = z = 0.475\text{cm}$ at .01 and .03 microseconds. The only preconditioner that was effective for the entire course of the simulation was the Schur preconditioner. We note that for the initial conditions, the starting values of the Rosseland and Planck opacities are on the order of 10^4 and 10^5 , respectively. Table 7.2 compares the statistics of the PVODE solver at .01 microseconds.

At this early output time, the matrix split preconditioner is actually performing the best. However, in the next phase of the solution all the preconditioners start having large numbers of linear convergence failures except for the Schur preconditioner. It has zero linear and nonlinear convergence failures for this problem. When a linear convergence failure occurs and the preconditioner is current, the PVODE solver reduces the step size and tries the step over. This has the effect of increasing the total number of steps for the simulation. Also note that the step sizes used by PVODE are much larger than one would expect for the split and Jacobi preconditioners to be effective. With step sizes of order 1, the errors in the split and Jacobi preconditioners are extremely large. Hence, it is not hard to understand the failure of these preconditioners (or their high cost since the step sizes must be kept small) for the latter part of the simulation. Figures 7.8 and 7.9 show the step size behavior as a function of output

times for the Schur and Block Jacobi preconditioners. While the step sizes change fairly smoothly for the Schur preconditioner, the behavior has a sawtooth flavor for the Block Jacobi preconditioner. The final statistics for the Schur preconditioner are given in Table 7.3, as well as the final computed statistics for the other preconditioners. For this problem, we requested 400 output snapshots. The Schur preconditioner finished the computation in under a 2 hour limit, while the Block Jacobi reached 285 output points, the Operator Split reached 14, and the Matrix Split reached only 7. For the Schur preconditioner, there were on average 1.1 nonlinear iterations per time step, and 2.0 linear iterations per nonlinear iteration.

TABLE 7.2
Statistics for LEOS Problem 1 at .01 Microseconds

PC	NST	NNI	NLI	NFE	NPE	NPS	HU	RT	NCFN	NCFL
S	323	420	562	985	59	975	1.52	298.43	0	0
BJ	305	518	1023	1544	68	1534	.95	488.59	0	14
MS	226	308	280	591	54	581	6.54	188.91	0	0
OS	493	663	1511	2179	103	2164	.78	663.87	0	33

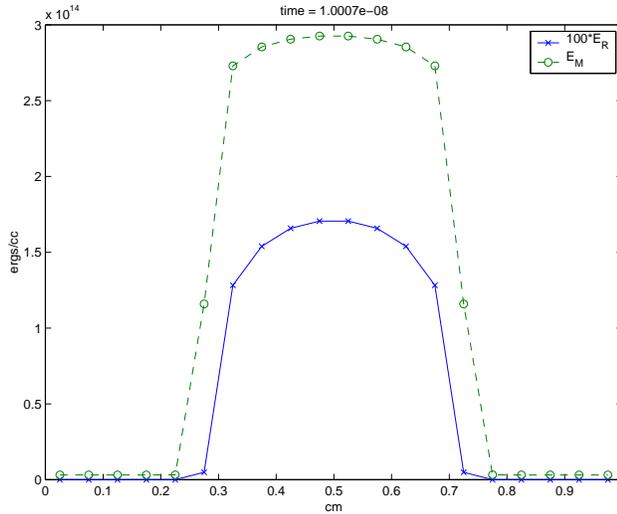


FIG. 7.6. Snapshot of E_R and E_M for the first LEOS problem on the line $y = z = 0.475\text{cm}$ at .01 microseconds

TABLE 7.3
Statistics for LEOS Problem 1 at 1.33 Microseconds

PC	NST	NNI	NLI	NPE	HU	RT	NCFN	NCFL
S	2654	2942	5930	189	4421	2592	0	0
BJ	>4197	>4969	>23387	>451	15.11	>7208	>55	>1720
MS	>3226	>5261	>20691	>1653	.00166	>6559	>473	>1680
OS	>4743	>6720	>22528	>962	.321	>6867	>10	>185

For this first problem, we also investigated the effect of restricting the ODE solver to first order (i.e., backward Euler). While this had a significant effect on

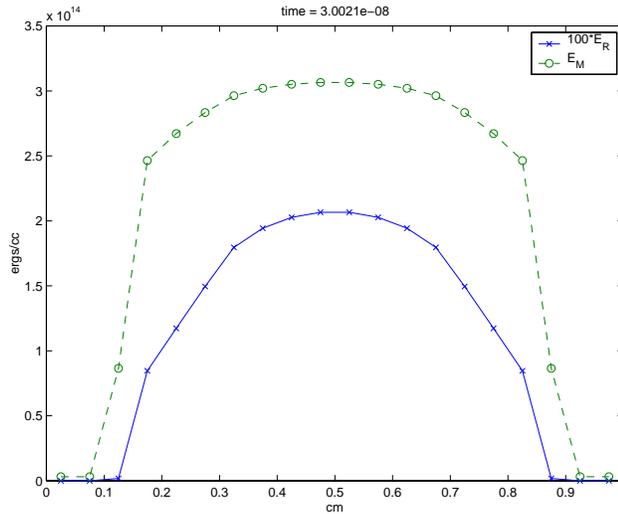


FIG. 7.7. Snapshot of E_R and E_M for the first LEOS problem on the line $y = z = 0.475\text{cm}$ at .03 microseconds

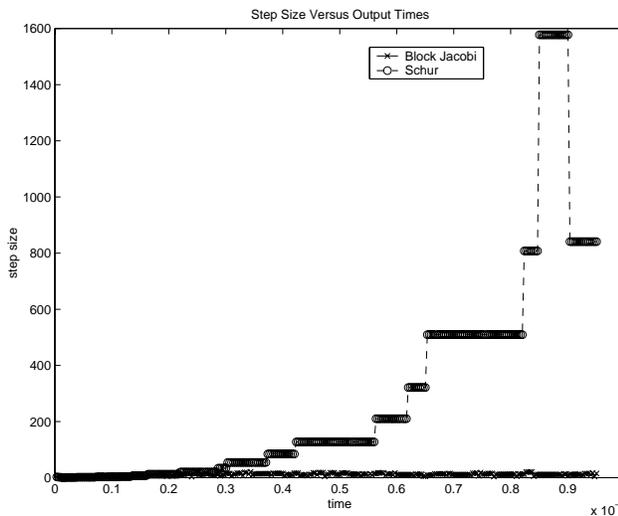


FIG. 7.8. Plot of step size comparisons for first LEOS problem. Note that the time steps are actually $c\Delta t$.

reducing the number of linear and nonlinear convergence failures for all but the matrix split preconditioner, the number of time steps increased dramatically. As a result, all of the preconditioners failed to produce the requested 400 output points within the two hour run time limit. Table 7.4 gives statistics and final times reached for each preconditioner. From these results, it is apparent that the higher order time integration methods can be extremely effective in reducing overall run time costs.

A second problem was run with hydrogen as the material at a reference density of $\rho = .874\text{g/cc}$. This problem has the same general behavior as the first, except that the time to reach steady state is an order of magnitude lower, i.e., the simulation was

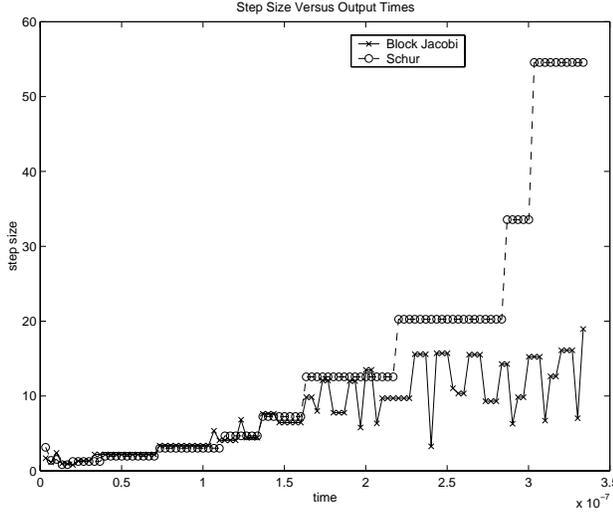


FIG. 7.9. Plot of step size comparisons for first LEOS problem. Note that the time steps are actually $c\Delta t$.

TABLE 7.4
Statistics for LEOS Problem 1 Restricted to First Order

PC	NST	NNI	NLI	Final Time Reached (μs)	RT	NCFN	NCFL
S	9236	9726	14725	.0097	7142	0	0
BJ	5970	6280	16256	.0047	6862	0	18
MS	6564	8177	13717	.0031	6894	191	701
OS	5629	5968	18706	.0053	7161	0	64

run to .133 microseconds. Figures 7.10, 7.11 and 7.12 show the solutions plotted on the line $y = z = 0.475cm$ at .01, .02 and .133 microseconds. Table 7.5 contains the statistics for this problem. As before, the Schur preconditioner was the only effective preconditioner for the entire run. There were on average 1.1 nonlinear iterations per time step, and 2.15 linear iterations per nonlinear iteration. Note that there were 13 linear convergence failures and 1 nonlinear failure.

TABLE 7.5
Statistics for LEOS Problem 2 at .133 Microseconds

PC	NST	NNI	NLI	NFE	NPE	NPS	HU	RT	NCFN	NCFL
S	834	943	2029	2976	81	2967	26.69	879	1	13

7.3. Scalability Study with Tabulated Opacities. A scalability study was performed on a third problem. The system (2.3)–(2.4) was solved on the box $\mathcal{D} \equiv \{ \mathbf{x} = (x, y, z) : 0 \leq x, y, z \leq 1cm \}$ with Dirichlet boundary conditions. The function $\chi(\mathbf{x})$ in (2.3) is defined by (7.1). The parameter T_{source} was 3,481,440 $^{\circ}K$, and the initial conditions were obtained using $T_{R,0} = T_{M,0} = 300^{\circ}K$. The Dirichlet boundary values were taken to be consistent with the initial conditions. The material used was carbon at a reference density of $\rho = 1.05g/cc$, and the spatial grid per processor was

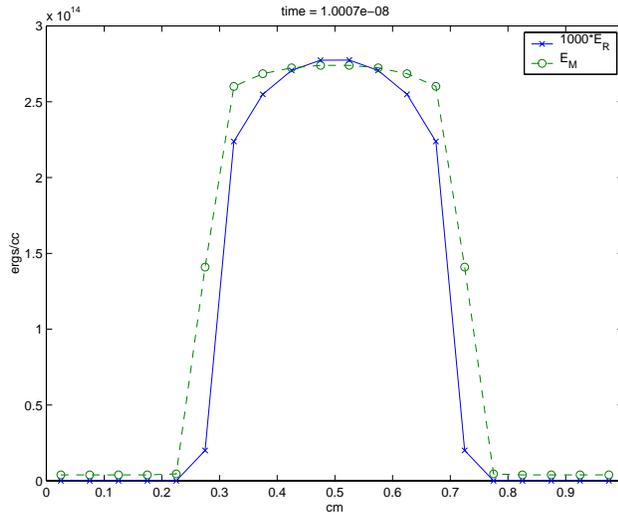


FIG. 7.10. Snapshot of E_R and E_M for the second LEOS problem on the line $y = z = 0.475$ cm at $.01$ microseconds

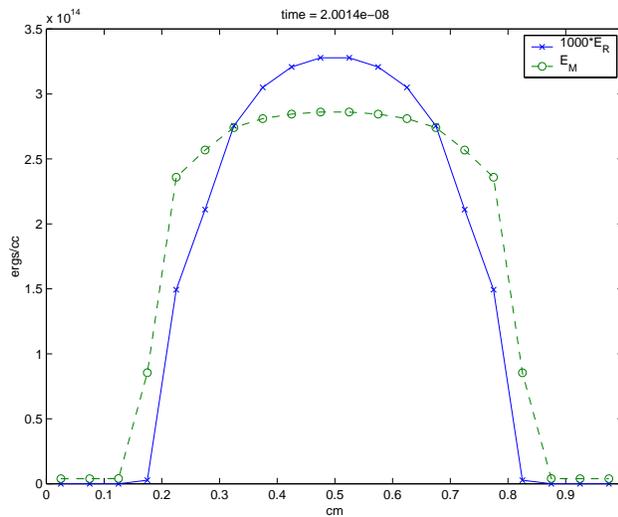


FIG. 7.11. Snapshot of E_R and E_M for the second LEOS problem on the line $y = z = 0.475$ cm at $.02$ microseconds

uniform with $N_x = N_y = N_z = 40$. Thus, problem size and computational resources were simultaneously increased for this study. Only the Schur preconditioner was used. Table 7.6 contains the results of the scalability study. The reported scaled efficiency for a run on N processors was calculated by dividing the run time for the single processor case by the run time for the N processor case. As can be seen, except for the run times, all the statistics scaled extremely well. (We note that when this study was performed, the run time environment on the IBM ASCI Blue Pacific machine at LLNL was under a state of flux. Earlier scalability studies performed showed a much better scalability of run times.) The simulation was run until approximately

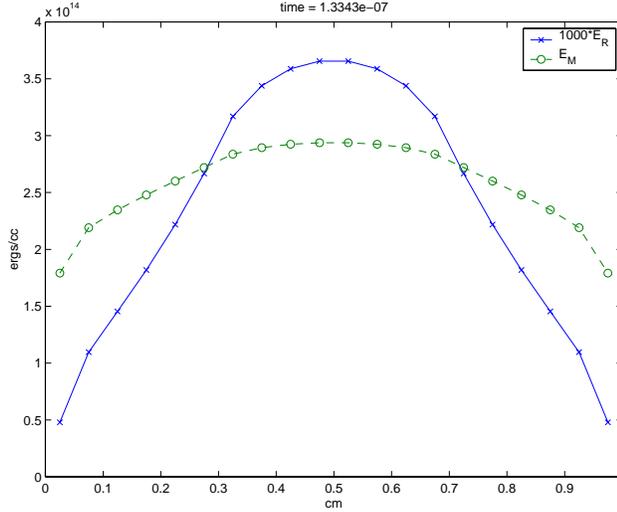


FIG. 7.12. Snapshot of E_R and E_M for the second LEOS problem on the line $y = z = 0.475\text{cm}$ at .133 microseconds

.001 microseconds, which is very early in the time history.

TABLE 7.6
Statistics for Scalability Study

Processor Topology	NST	NNI	NLI	NPE	RT	RT Scaled Efficiency	Avg. Cost per Step
$1 \times 1 \times 1$	217	329	423	70	2015	—	9.3
$2 \times 2 \times 2$	214	324	411	75	2287	88.1%	10.7
$4 \times 4 \times 4$	196	295	378	66	2220	90.7%	11.3
$8 \times 8 \times 8$	197	273	374	57	2575	78.2%	13.1
$16 \times 8 \times 8$	190	273	376	60	3106	64.8%	16.4

8. Conclusions. We have presented a comparison of four preconditioning strategies for Jacobian systems arising in the fully implicit solution of radiation diffusion coupled with material energy transfer. The four preconditioning methods are: block Jacobi, Schur complement, and operator splitting approaches that split the preconditioner solve into two steps. From our results, it is apparent that the Schur complement approach is clearly seen to be the most effective for a large range of weightings between the diffusion and energy coupling terms. For problems using tabulated opacities, the Schur preconditioner outperformed the other preconditioners by a wide margin. One conclusion we can draw from our studies is that it appears to be more effective to use full matrix approaches to developing preconditioners for radiation transport problems rather than approaches based on splittings or on only parts of the matrix.

While limiting the time integration methods to first order helps lower the number of step failures, there is a marked increase in the number of steps. At least for the problems we have considered, the better preconditioner allows for much larger step sizes within the allowed error bounds used by the PVODE solver, and this significantly reduces the overall work. Our parallel scaling study demonstrated good algorithmic

scalability of our solution approach. Finally, when material diffusion is added to equation (2.4), the Schur approach may fail to be competitive. In addition, extending the physics in the problem to include multigroup diffusion (where the radiation energy spectrum is resolved) leads to a Schur complement that is a full matrix. Our future work will include exploring the use of system-based multigrid solvers as preconditioners, as well as other multilevel methods, to address material energy diffusion and multigroup energy resolution.

REFERENCES

- [1] S. F. ASHBY AND R. D. FALGOUT, *A parallel multigrid preconditioned conjugate gradient algorithm for groundwater flow simulations*, Nuclear Science and Engineering, 124 (1996), pp. 145–159.
- [2] R. L. BOWERS AND J. R. WILSON, *Numerical Modeling in Applied Physics and Astrophysics*, Jones and Bartlett, Boston, 1991.
- [3] W. F. BRIGGS, V. E. HENSON, AND S. F. McCORMICK, *A Multigrid Tutorial, 2nd. Edition*, SIAM, Philadelphia, PA, 2000.
- [4] P. N. BROWN, G. D. BYRNE, AND A. C. HINDMARSH, *VODE: A variable-coefficient ODE solver*, SIAM J. Sci. Stat. Comput., 10 (1989), pp. 1038–1051.
- [5] P. N. BROWN, B. CHANG, F. GRAZIANI, AND C. S. WOODWARD, *Implicit solution of large-scale radiation-material energy transfer problems*, in Iterative Methods in Scientific Computation IV, D. R. Kincaid and A. C. Elster, eds., vol. 5 of Series in Computational and Applied Mathematics, New Brunswick, NJ, 1999, International Association for Mathematics and Computers in Simulations, pp. 343–356.
- [6] P. N. BROWN AND A. C. HINDMARSH, *Reduced storage matrix methods in stiff ODE systems*, J. of Appl. Math. and Comp., 31 (1989), pp. 40–91.
- [7] G. D. BYRNE, *Pragmatic experiments with Krylov methods in the stiff ODE setting*, in Computational Ordinary Differential Equations, J. R. Cash and I. Gladwell, eds., Oxford, 1992, Oxford University Press, pp. 323–356.
- [8] G. D. BYRNE AND A. C. HINDMARSH, *PVODE, an ODE solver for parallel computers*, Int. J. High Perf. Comput. Appl., 13 (1999), pp. 354–365. Also available as LLNL Tech. Report UCRL-JC-132361.
- [9] E. M. COREY AND D. A. YOUNG, *A new prototype equation of state data library*, Tech. Rep. UCRL-JC-127698, Lawrence Livermore National Laboratory, 1997. Submitted to American Physical Society Meeting, 1997.
- [10] A. C. HINDMARSH, *Avoiding BDF stability barriers in the MOL solution of advection-dominated problems*, Applied Numerical Mathematics, 17 (1995), pp. 311–318.
- [11] A. C. HINDMARSH AND M. D. ROTTER, *Using an ODE solver for a class of integro-differential systems*, Tech. Rep. UCRL-JC-138647, Lawrence Livermore National Laboratory, 2000. Submitted to J. Comp. Phys.
- [12] K. R. JACKSON AND R. SACKS-DAVIS, *An alternative implementation of variable step-size multistep formulas for stiff ODEs*, ACM Trans. Math. Software, 6 (1980), pp. 295–318.
- [13] J. E. JONES AND C. S. WOODWARD, *Newton-Krylov-multigrid solvers for large-scale, highly heterogeneous, variably saturated flow problems*, Center for Applied Scientific Computing UCRL-JC-138618, Lawrence Livermore National Laboratory, Livermore, CA, Apr. 2000. Submitted to Advances in Water Resources.
- [14] D. A. KNOLL, W. J. RIDER, AND G. L. OLSON, *An efficient nonlinear solution method for nonequilibrium radiation diffusion*, J. Quant. Spec. and Rad. Trans., 63 (1999), pp. 15–29.
- [15] V. A. MOUSSEAU, D. A. KNOLL, AND W. J. RIDER, *Physics-based preconditioning and the newton-krylov method for non-equilibrium radiation diffusion*, J. of Comput. Phys., 160 (2000), pp. 743–765.
- [16] G. C. POMRANING, *The Equations of Radiation Hydrodynamics*, Pergamon, New York, 1973.

University of California
Lawrence Livermore National Laboratory
Technical Information Department
Livermore, CA 94551

