Numerical Heat Transfer, Part B, 46: 19–38, 2005

ISSN: 1040-7790 print/1521-0626 online
DOI: 10.1080/10407790490487451

Taylor & Francis
Taylor & Francis Group

PARALLEL COMPUTATIONS OF RADIATIVE HEAT TRANSFER USING THE DISCRETE ORDINATES METHOD

Gautham Krishnamoorthy, Rajesh Rawat, and Philip J. Smith
Department of Chemical and Fuels Engineering, University of Utah, Salt Lake City, Utah, USA

The discrete ordinates method is spatially decomposed to solve the radiative transport equation on parallel computers. Mathematical libraries developed by third parties are used to solve the matrices that result during the solution procedure. The radiation component is verified by comparing computed values against a benchmark. Fixed and scaled problem size efficiencies are examined. Contrary to most previous studies, the parallel efficiencies did not depend strongly on the optical thickness of the medium for our model problem. Timing studies show that GMRES, BiCGSTAB iterative methods with block Jacobi preconditioning perform the best for solving these matrix systems.

INTRODUCTION

Radiation is often the dominant mode of heat transfer in hydrocarbon fires. With the advent of parallel computers, performing realistic computations of participating media radiative transfer is becoming increasingly tractable. Issues relating to the use of high-performance computing in participating media heat transfer were identified at a National Science Foundation workshop [1]. In order to spatially resolve the important flow characteristics in a fire, grids containing 10^6–10^8 computational cells need to be used at every time step associated with the calculation. Performing radiation calculations at these scales is extremely expensive due to its dependence on space, direction, as well as the energy of the photons. Consequently, parallelism in radiation calculations may be achieved by decomposing the radiation solution domain in spatial, angular, or the energy domains, respectively. We have incorporated a finite-volume-based discrete ordinates radiation model decomposed

Received 29 January 2004; accepted 17 April 2004.

This work was supported by the Center for the Simulation of Accidental Fires and Explosions (C-SAFE) at the University of Utah. C-SAFE is a part of U.S. Department of Energy, Academic Strategic Alliance Partners (ASAP)—Advanced Simulation and Computing (ASC, formerly known as ASCI) alliance. The authors would like to thank Dr. Stanislav Borodai, Divya Ramachandran, and our collaborators in the Scientific Computing and Imaging Institute at the University of Utah for assistance rendered during various stages of this effort. The authors would also like to express their thanks to the reviewers for their feedback.

Address correspondence to Philip J. Smith, Department of Chemical and Fuels Engineering, 50 S Central Campus Dr, Room 3290 MEB, University of Utah, Salt Lake City, UT 84112-9203, USA. E-mail: smith@crsim.utah.edu

NOMENCLATURE

A	surface area of the face of the volume element normal to the x axis	V	volume of a grid element
B	surface area of the face of the volume element normal to the y axis	w	angular weights
C	surface area of the face of the volume element normal to the z axis	x, y, z	coordinate directions
G	incident radiation	ε	total emissivity
I	radiative intensity	ζ	direction cosine in the x direction
k	absorption coefficient	η	direction cosine in the z direction
L_m	mean beam length	η_f	incremental fixed problem size efficiency
M	preconditioning matrix	η_s	scaled problem size efficiency
n_n	number of computational nodes along each dimension	μ	direction cosine in the y direction
N_p	number of processors	σ	Stefan-Boltzmann constant
P	square coefficient matrix	τ	optical thickness
$\mathbf{q}$	radiative heat flux vector	Ω	solid angle
Q	right-hand-side vector of the linear matrix system	**Subscripts**	
r	residual	b	blackbody
$\mathbf{r}$	location vector	i, j, k	surface element indices
R	condition number of matrix	m	discrete angular direction of intensity
$\hat{\mathbf{s}}$	unit vector in the direction of radiation intensity	n	order of the angular quadrature set
t_g	wall clock time	p	node at which the directional intensity is being solved
T	temperature	s	index that runs over all neighboring surface nodes
		t	iteration number

in the spatial domain. The inputs to this model are the concentrations of the important radiatively active species (CO_2, H_2O, soot) and temperature, which are calculated on the spatially decomposed flow grid as well as at the boundaries. The adoption of a spatial decomposition strategy for the radiation component allows for its easy integration with other combustion components. This radiation algorithm is then used with the combustion code to perform transient simulations of fires with nongray properties.

First we summarize previous applications of spatial decomposition strategies to finite-volume-based radiation models. Burns and Christon [2, 3] used massively parallel supercomputers to examine the scaled and fixed problem efficiencies that resulted from solving a 3-D benchmark radiation problem using the discrete ordinates method. They observed that the global nature of radiative transport resulted in a decrease in parallel efficiency but provided useful speedups for large computational grids. Angular as well as spatial parallelization techniques were applied to the discrete ordinates method by Gonçalves and Coelho [4] and to the finite-volume method by Coelho and Gonçalves [5]. The influences of the angular and spatial resolutions and absorption coefficient of the medium on the parallel efficiencies were investigated in both these studies. Liu et al. [6] parallelized an unstructured finite-volume method by decomposing the spatial domain along the longer geometric dimension. More recently, Tal et al. [7] parallelized a discrete ordinates code with pseudo-time-stepping to achieve high efficiencies employing the spatial decomposition technique. Table 1 summarizes these works along with other recent efforts in parallelizing radiation calculations. Though the table might seem to suggest ray or

angular decomposition to be more efficient than spatial decomposition, angular or ray decomposition techniques have only been achieved with coarse grids because they require the nondecomposed spatial domain to be stored in each processor in its entirety.

The main features that distinguish this work from the ones described above are that we solve the discretized radiative transport equation (RTE) by interfacing our radiation component to robust, scalable, nonlinear, and linear solvers developed by third parties. This provides us with easy access to a suite of direct and iterative solvers and preconditioners that can be suitably selected depending on the problem being solved. The comparative performances of different solvers and preconditioners are reported in the Results and Discussion section of this study. Second, we report the parallel performance of the radiation algorithm up to 960 processors during the course of a multiphysics fire simulation. The authors are not aware of any other work that has demonstrated the scaling of radiative heat transfer computations carried out within the context of a combustion simulation with this many processors (cf. Table 1). Finally, in contrast to most previous studies that employed domain decomposition to solve the RTE, we did not notice a degradation of the parallel efficiencies at low optical thickness for our model problem.

The fire simulations reported in this work were computed on a 68-node, 16-processors-per-node, IBM SP system at the Lawrence Livermore National Laboratory (“Frost”). All other calculations were carried out on a distributed memory Linux cluster at the University of Utah that has 128 dual processor nodes.

THE DISCRETE ORDINATES METHOD

In this work the discrete ordinates method is used to solve the differential form of the RTE. The discrete ordinates method is based on the numerical solution of the RTE along specified directions. The total solid angle about a location is divided into a number of ordinate directions, each assumed to have uniform intensities. Each transport equation that is solved corresponds to an ordinate direction (selected from an angular quadrature set that discretizes the unit sphere) and describes the variation of directional intensity throughout the domain. If ζ_m, μ_m, and η_m represent the direction cosines associated with each ordinate direction, k represents the absorption coefficient and I_b represents the blackbody emissive power, then the differential equation governing the discrete ordinates method in the absence of scattering can be written for each direction m as [12]

$$\zeta_m \frac{\partial I_m}{\partial x} + \mu_m \frac{\partial I_m}{\partial y} + \eta_m \frac{\partial I_m}{\partial z} = -kI_m + kI_b \tag{1}$$

The boundary condition associated with the above equation, considering the surrounding surfaces to be black, is

$$I_m = I_b \tag{2}$$

If the absorption coefficient and temperature within the domain as well as at the boundaries are specified, then Eq. (1) can be solved iteratively for the directional intensities (I_m) throughout the domain. Equation (1) needs to be solved for each direction associated with the discrete ordinates method.

Table 1. Summary of previous efforts in parallelizing radiative heat transfer

Reference	Radiation model (decomposition strategy)	Problem description (maximum number of processor used), (maximum grid size)	Comments
Burns and Christon [2, 3]	Discrete ordinates (spatial)	3-D benchmark radiation problem (512), ($305 \times 305 \times 305$)	Observed degradation in parallel efficiency with increase in number of processors but obtained useful speed-ups for large computational grids.
Gonçalves and Coelho [4]	Discrete ordinates (spatial and angular)	3-D rectangular enclosures and 2-D square enclosure (80), ($36 \times 12 \times 12$, 80×80)	Obtained higher efficiencies with the angular decomposition strategy. Efficiencies increased with increase in absorption coefficient when spatial decomposition strategy was used for prescribed medium temperature problem.
Coelho and Gonçalves [5]	Finite volume (spatial and angular)	3-D rectangular enclosures (80), ($36 \times 12 \times 12$, 80×80)	Obtained higher efficiencies with the angular decomposition strategy. Efficiencies increased with increase in absorption coefficient and grid resolution when spatial decomposition strategy was used for prescribed medium temperature problem.
Liu et al. [6]	Unstructured finite volume (spatial decomposition along longer geometric dimension)	2-D axisymmetric furnace and an idealized 3-D furnace (18), (8,140 triangular volume cells, 65,968 prism volume cells)	Found degradation in parallel performance with increase in processor number with increase in absorption coefficient for a prescribed temperature problem. Parallel performance improved with increase in grid number.

Tal et al. [7]	Discrete ordinates with pseudo-time stepping (spatial)	2-D and 3-D benchmark problems (16), (128×128, $64 \times 64 \times 64$)	Attribute high parallel efficiency to explicit time-stepping scheme and high-bandwidth shared memory
Novo et al. [8]	Discrete transfer method (spatial and ray)	2-D and 3-D enclosures (16), (128×128, $64 \times 64 \times 64$)	Obtained higher efficiencies with the ray decomposition technique. Obtained higher efficiencies with increase in optical thickness when spatial decomposition was used.
Cumber and Beeri [9]	Discrete transfer radiation model (ray)	3-D test problem, methane jet fire in cross-flow (14) (46656)	Describe a semi-empirical model for the parallel run time, calibrated using the timing data on IBM SP2 architecture.
Yan [10]	Discrete transfer method (angular)	2-D turbulent buoyant diffusion flame, 3-D complex compartment fire (8), (46×50, 400,000 + gas-phase and solid-phase nodes)	Applied spatial decomposition to the computation of nonradiative turbulent reacting flow and solid-fuel pyrolysis. Angular decomposition applied to radiation computation. Obtained good parallel performance.
Marakis et al. [11]	Monte Carlo	3-D idealized furnance (9), ($22 \times 22 \times 42$)	Obtained extremely high efficiencies for the cases examined.

The variables of interest in most radiative transfer analysis are the distributions of radiative heat flux vectors [**q(r)**] and the radiative source terms [$-\nabla \bullet \mathbf{q}(\mathbf{r})$]. The radiative source term describes the conservation of radiative energy within a control volume and goes into the total energy equation, thereby coupling radiation with the other physical processes that occur in a multiphysics application. Both of these are direction-integrated quantities and are readily determined once the distributions of directional intensities (I_m) within the domain are known. In the discrete ordinates method, any integral over the solid angle is replaced by a quadrature of order n and an appropriate angular weight (w_m) associated with each direction m. The number of equations to be solved depends on the order of approximation n used. The radiative heat fluxes within the domain or at a surface and the radiative flux divergence are determined as [13]

$$\mathbf{q}(\mathbf{r}) = \int_{4\pi} I(\mathbf{r}, \hat{\mathbf{s}})\hat{\mathbf{s}}\, d\Omega \approx \sum_m w_m I_m(\mathbf{r})\hat{\mathbf{s}}_m \tag{3}$$

$$\nabla \bullet \mathbf{q}(\mathbf{r}) = k(\mathbf{r})[4\pi I_b(\mathbf{r}) - G(\mathbf{r})] \tag{4}$$

G, the incident radiation, is calculated as

$$G(\mathbf{r}) = \int_{4\pi} I(\mathbf{r}, \hat{\mathbf{s}})\, d\Omega \approx \sum_m w_m I_m(\mathbf{r}) \tag{5}$$

Multiplying both sides of Eq. (1) by $dV = dx\, dy\, dz$ and integrating over the volume elements, the intensity I_m at the center of each control volume in a uniform, structured grid is related to the face-centered values of the intensities as

$$\zeta_m A(I_{m,i+1} - I_{m,i}) + \mu_m B(I_{m,j+1} - I_{m,j}) + \eta_m C(I_{m,k+1} - I_{m,k}) = -kI_m V + kI_b V \tag{6}$$

where A, B, and C represent the surface areas of the faces of the volume element that are normal to the x, y, and z axes, respectively, and i, j, and k represent the surface element indices along the x, y, and z axes, respectively.

We employ the upwind differencing scheme to relate the face-centered values to the cell-centered values. For directions with positive direction cosines, this results in

$$I_{m,i+1} = I_{m,j+1} = I_{m,k+1} = I_m \tag{7}$$

A similar analysis can be carried out for other combinations of direction cosines as well. The upwind scheme is less accurate than higher-order schemes but guarantees non-negative intensities. Substituting Eq. (7) in Eq. (6), we get

$$(\zeta_m A + \mu_m B + \eta_m C + kV)I_m = \zeta_m A I_{m,i} + \mu_m B I_{m,j} + \eta_m C I_{m,k} + kI_b V \tag{8}$$

Thus, Eq. (8) corresponds to a discretized set of linear algebraic equations of the form

$$P_p I_{m,p} + \sum_s P_s I_{m,s} = Q_p \tag{9}$$

where the subscript p corresponds to the node at which Eq. (1) is being approximated, the index s runs over all the neighboring surface nodes in Eq. (8), P_s corresponds to the product of the direction cosines with the geometric properties (surface

areas), and Q_p contains all the terms that are known. Equation (9) can be written is matrix form as

$$PI_m = Q \tag{10}$$

where P corresponds to a square coefficient matrix. If $I_{m,i}, I_{m,j},$ or $I_{m,k}$ corresponds to intensity from a boundary, they must be suitably replaced by I_b according to the boundary condition, Eq. (2). They must then be multiplied by P_s and added to Q_p on the right-hand side of Eq. (9). Equation (10) corresponds to a linear system of equations for each direction in which the RTE is being solved. For the problems considered in this article, the square coefficient matrix (P) is a sparse nonsymmetric matrix with three other bands running parallel to the diagonal. The number of bands on the upper or lower triangular part of the matrix depend on the direction in which the RTE is being solved. A detailed derivation of the discretized form of Eq. (1) that includes the effect of scattering can be found in Siegel and Howell [14].

ALGORITHM DESCRIPTION

In order to efficiently solve the set of linear equations that result from each direction, we make use of the Portable Extensible Toolkit for Scientific Computation (PETSc) [15]. PETSc provides a comprehensive suite of data structures for parallel matrix and vector storage as well as unified interfaces to linear solvers and preconditioners for achieving scalable parallel computation. The matrices are computed and distributed among all the processors involved in the simulation. PETSc enables us to work only with global indices, even though, internally, local indices are used for accessing the distributed data structures. Once the mapping between the global and local indices is done, global matrices associated with every direction can readily be constructed. This eliminates the need to redo the mapping of indices, as encountered with solver strategies where the order of visiting the control volumes is appropriately selected according to the ordinate direction under consideration. PETSc provides various sparse matrix storage formats, all of which have a uniform interface to the matrix operations. PETSc thus offers the benefits of being a general-purpose solver, with many of its solvers appropriate for problems that are discretized using unstructured grids as well.

A preconditioned, iterative method results in the splitting of the matrix P in Eq. (10) of the form

$$P = M - N \tag{11}$$

where M is the preconditioning matrix (usually containing the diagonal or a triangular part of P) and N contains the negative of the remaining elements of P. The matrix equations to be solved then reduce to the form

$$MI_{m,t+1} = (M - P)I_{m,t} + Q \tag{12}$$

where t is the iteration number. It is best that the preconditioner M be convenient, sparse, invertible, and as close an approximation to P as possible.

Convergence is determined by the condition number or the ratio of the largest to smallest eigenvalues of the matrix given by [16]

$$R = M^{-1}(M - P) \tag{13}$$

The preconditioner approximates the solution on each subdomain with little or no communication between the partitions of the other domains. The outer solver then performs iterations and tries to resolve for the residual error over the global mesh and not just for the individual subdomains as might be encountered during the parallelization of direct methods. The criterion for convergence is the I_2 norm of the preconditioned residual [15]. The preconditioned residual is defined as

$$r_m \equiv M^{-1}Q - M^{-1}PI_m \tag{14}$$

Convergence is detected at iteration t if

$$\|r_t\|_2 < \max(\text{rtol} \times \|r_0\|_2, \text{atol}) \tag{15}$$

where rtol gives the relative decrease in the residual norm, and the absolute size of the residual norm is given by atol. The values of rtol for all the cases run in this article was set at 10^{-6}. A schematic of the sequential radiation algorithm is shown in Figure 1. As can be seen, the algorithm is designed in such a way that it can easily be extended to perform nongray analysis when absorption coefficients as a function of wavelength become available.

DESCRIPTION OF CASES

The first objective of this work was to carry out verification of the algorithm by comparing against benchmarks. The nonhomogeneous medium benchmark (case 1) that was studied for verification was first introduced by Hsu and Farmer [17]. The problem consists of an isothermal unit cube with cold, black walls. The interior of the cube consists of a gray nonscattering, absorbing/emitting material with an optical thickness ($\tau =$ absorption coefficient times the side length) distribution given by

$$\tau(x, y, z) = 0.9\left(1 - \frac{|x|}{0.5}\right)\left(1 - \frac{|y|}{0.5}\right)\left(1 - \frac{|z|}{0.5}\right) + 0.1 \tag{16}$$

A uniform blackbody emissive power of unity within the domain defines the distribution of temperature. Since the radiative property, temperature, and the boundary conditions for this problem are now known, the RTE can be solved to determine the distributions of the radiative fluxes and the radiative flux divergence. If we assume that the cube is centered on the origin and oriented so that the sides of the cube are orthogonal to the principal Cartesian axes, the benchmark values for the radiative heat flux along the line [$\mathbf{q}$ (x, 0.5, 0)] and the flux divergence along the line [$\nabla \bullet \mathbf{q}(\mathrm{x}, 0, 0)$] for this problem have been tabulated by Burns and Christon [2, 3]. A root-mean-square (RMS) error norm, also known as L2 error norm, was evaluated as defined by Burns and Christon [2, 3] in order to express the overall solution accuracy of the numerical method in a compact form. First, the relative error at each node is calculated as

$$e_{\varphi,p} = 100\frac{(\varphi_{\mathrm{Sn}} - \varphi_{\mathrm{b.m}})_p}{\|\varphi_{\mathrm{b.m}}\|_{\max}} \tag{17}$$

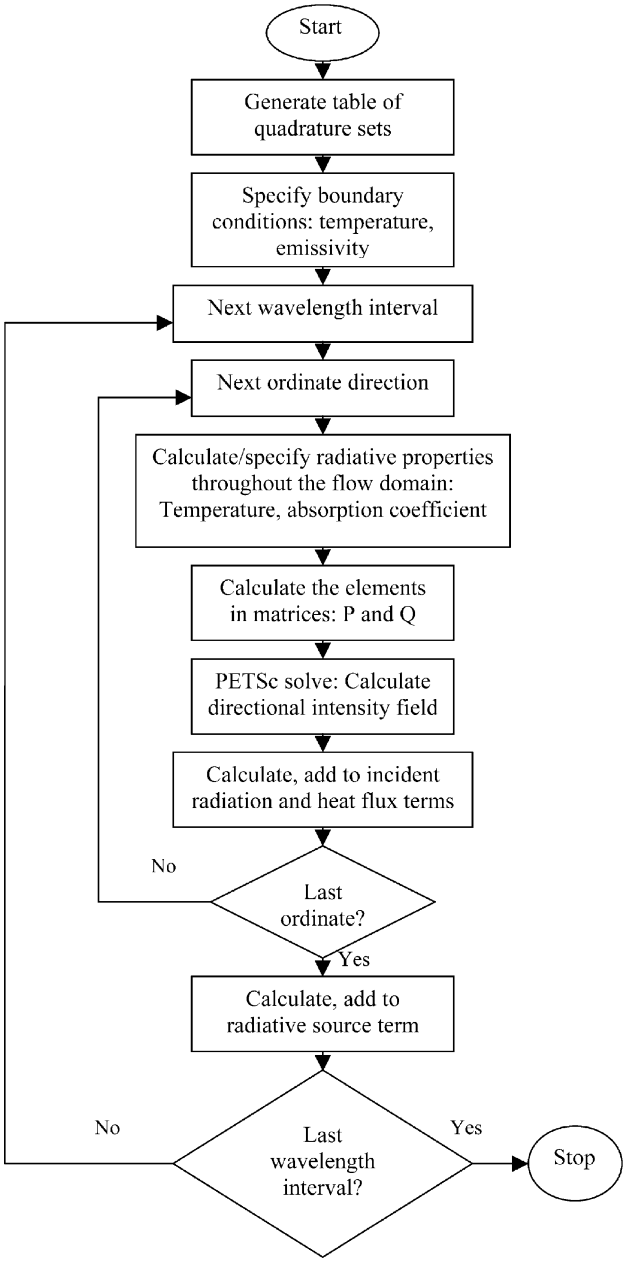

Figure 1. Logic flow diagram of the radiation algorithm.

where e is the relative error, φ refers to the radiative flux or flux divergence calculated at each nodal location (p), and Sn and b.m refer to discrete-ordinates solution values and benchmark values, respectively. $\|\varphi_{b.m}\|_{max}$ refers to the maximum benchmark value along the line.

The RMS error norm can then be calculated as

$$\|e\|_{\mathrm{RMS}} = \left\{ \frac{\sum_{i=1}^{n_n} e_{\varphi,i}{}^2}{n_n} \right\}^{1/2} \tag{18}$$

where n_n is the number of computational nodes along each dimension.

Burns and Christon [2, 3] reported the overall solution accuracies from their algorithm and code for this problem employing a quadrature set proposed by Lathrop and Carlson (LC) [18]. They also rotated this quadrature set by 45° to evaluate the effect of the structure of the quadrature set on the solution accuracy. They found that rotating the quadrature set decreased the RMS error of the flux divergence significantly compared to the unrotated set. However, they did not notice any significant improvement in the error norms of the heat flux for the rotated over the unrotated quardrature set. We employed the spherical surface symmetrical equal dividing angular quadrature scheme (SSD) [19] to calculate the numerical solution accuracies for this problem. The SSD quadrature scheme has been previously shown to compete well with other traditional quadrature sets and provide good accuracies in black-walled rectangular enclosures, which are the type of configurations adopted in this study [19]. Moreover, the number of equations that need to be solved with the SSD_{1a}, SSD_{2a}, and SSD_{3b} schemes are exactly the same as those of the rotated LC_4, LC_6, and LC_8 quadrature sets, respectively. Consequently, the solve times for the two schemes for the corresponding ordered pairs are also nearly the same. This enables us to directly compare the solution accuracies of the two schemes when the same number of equations is being solved.

Case 2 was designed to study the variation of the parallel performance with the absorption coefficient of the medium. The problem consists of a 1-m-thick isothermal layer of a hypothetical gas at a temperature of 1,000 K and at a pressure of 1 atm, confined between two cold, black, parallel plates. The absorption coefficient (k) of the medium was then varied spanning four orders of magnitude from 0.01 to $10\,\mathrm{m}^{-1}$. The heat loss from the gas for this particular problem can be easily compared with the results obtained using a mean-beam-length approach. For this case after employing a mean beam length (L_m) of 1.76 m, the heat loss ($\mathbf{q}$) to the plates is seen to be

$$\mathbf{q} = 4\sigma\varepsilon T^4 \; W/m^2 \tag{19}$$

where T is the gas temperature and the emissivity ε is calculated as

$$\varepsilon = 1 - e^{-kL_m} \tag{20}$$

The 1-D problem of case 2 was modeled using a 3-D domain taking the geometric length and breadth to be 50 times the height (1 m). The four end walls were modeled as black surfaces at 0 K. This approach of employing a 3-D grid to approximate and solve a 1-D problem is similar to the one employed by Fiveland and Jamaluddin [20], who employed a grid whose geometric length and breadth was 24 times the height, to study radiative heat transfer between two infinitely parallel plates with a nongray medium between them. We therefore felt that our 3-D approximation to the 1-D problem was accurate enough that the center of the parallel plates was not only receiving its due share of the radiation but was also free

from the edge/boundary effects that result from modeling a 1-D problem on a 3-D domain. We employed a grid resolution of 121^3 and the SSD_{3b} angular quadrature scheme to solve this problem. The computed values of the heat fluxes at the center of the parallel plates were than compared to the exact results obtained from Eq. (19) as a verification step.

Computational experiments were performed to compare the relative speeds of preconditioned iterative techniques for solving the sparse, unsymmetric matrices in Eq. (10) that result during the solution of case 1. This effort is similar to the one by Brussino and Sonnad [21] where direct and preconditioned iterative techniques for the solution of nonsymmetric sparse systems of linear equations that arise from finite-element (FE) simulation of semiconductor devices were compared. The iterative techniques considered for comparison in their study were the bi-conjugate gradient method (BCG), the conjugate gradient squared method (CGS), the generalized minimal residual method (GMRES), the generalized conjugate residual method (GCR), and the method of orthogonal minimization (ORTHOMIN). They tested each of these methods using similar point-preconditioning (incomplete LU factorization) methods on a set of large sparse matrices. They also employed a direct method that used LU factorization to solve these problems. The results from their numerical experiments showed that preconditioned iterative methods offered significant savings in storage and CPU time. For the matrices considered in their study, they concluded that the CGS method performed best. However, the GMRES performed significantly better than the CGS on some ill-conditioned test problems. In this article we compare the performance of GMRES [15] and bi-conjugate gradient stabilized (BiCGSTAB) [15] iterative methods by employing the same preconditioner (block Jacobi) on both. In an analogous effort, we also investigate the effect of employing two different preconditioning techniques (point Jacobi and block Jacobi) on the GMRES iterative method. We also investigate the distribution of radiation solve times in solving the two cases described above and from this suggest possible ways of speeding up the algorithm.

RESULTS AND DISCUSSION

The RMS error norms of the radiative flux divergence and the heat flux term for case 1 are shown in Figure 2. The results obtained with the SSD quadrature scheme are shown as a function of angular and spatial resolution. The results obtained by Burns and Christon [2, 3] using the rotated LC quadrature scheme are also shown with open symbols. In general, we notice that the two schemes perform equally well and there is a decrease in error norms with increased spatial and angular resolutions. The reason for the significant increase in the numerical error of the radiative flux divergence with spatial resolution with the SSD_{1a} quadrature scheme is as follows. A high-degree polynomial fit was made to the available benchmark results of heat flux and the flux divergence to get distribution functions of the benchmark values along the lines under consideration. Then the relative error of the computed radiative flux and the flux divergence were computed at each grid point of the spatial domain according to Eq. (17). Now, the maximum value of the radiative flux divergence occurs at the center of the domain [2, 3]. The maximum errors associated with the flux divergence also sharply peak toward the center of the

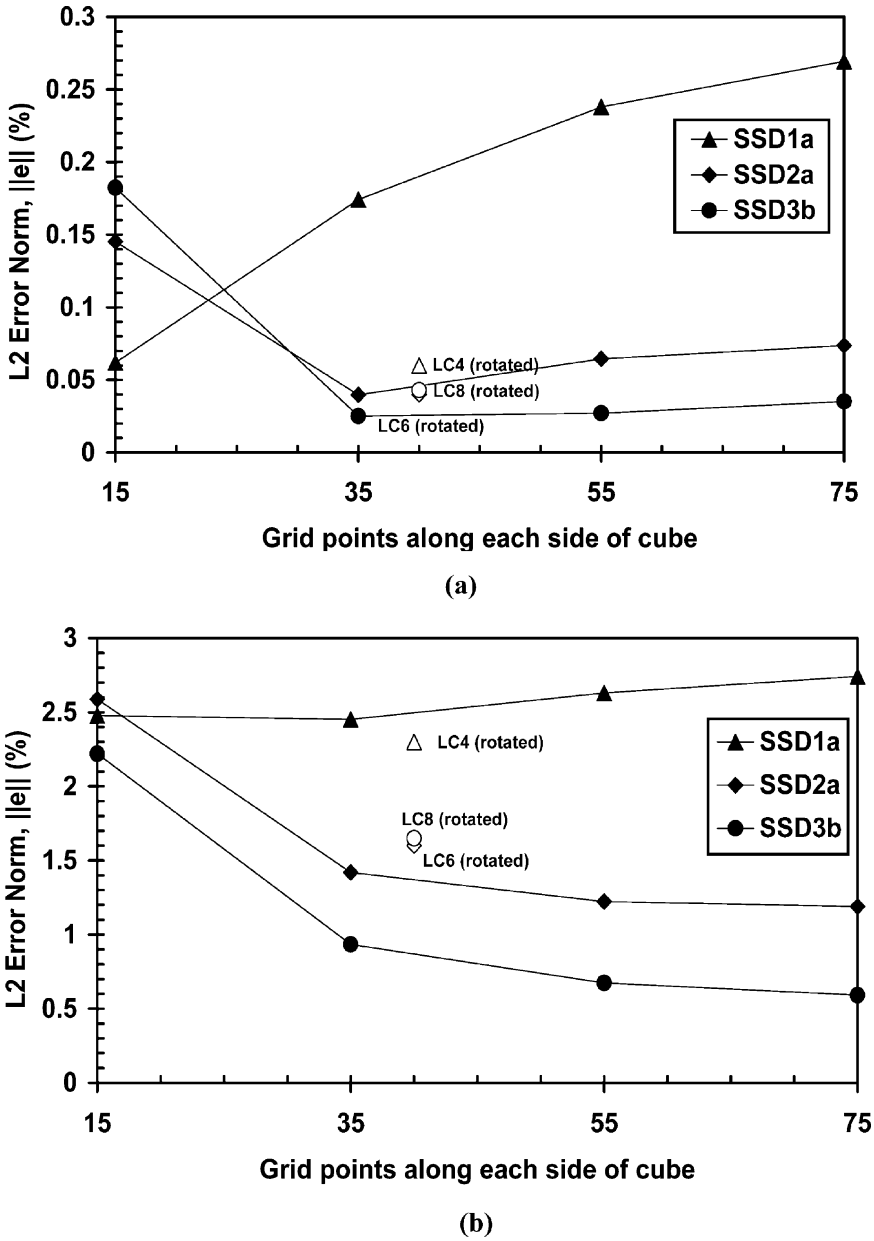

Figure 2. Spatial and angular dependence of the numerical accuracies in case 1: (*a*) predicted radiative flux divergence along (x, 0, 0); (*b*) radiative heat flux along (x, 0.5, 0).

domain [3]. With increased spatial resolution there are more grid points that fall in regions near the center, hence the contribution to the resulting relative error norm grows. This effect is more pronounced only at low angular resolutions because case 1 has been formulated in such a manner that the distribution of optical thickness in the

medium [cf. Eq. (16)] can under certain conditions increase the "ray effects" [22], which can reduce the accuracy of the discrete ordinates solution procedure, particularly at low angular resolutions.

As mentioned previously, this algorithm has been designed to perform nongray calculations also. Nongray calculations with this gray benchmark were also carried out for further verification of the algorithm. The spectral region of interest was divided into intervals and equal fractions of the blackbody function were allocated to each interval while ensuring that the sum of the fractions was unity. The error norms from the nongray calculations (not shown in the figure) fell right on top of the gray-calculation error norms for this problem, as they should. However, for actual engineering calculations, a series summation such as the one proposed by Chang and Rhee [23] can be used to determine the fractions of the blackbody functions corresponding to each spectral interval.

In accordance with Burns and Christon [2, 3], we define an incremental fixed problem size efficiency as

$$\eta_f = \frac{2t_g(2)}{N_p t_g(N_p)} \tag{21}$$

where $t_g(N_p)$ is the wall clock time taken to perform the entire calculation except the problem specification (generating table of quadrature sets and specifying boundary conditions) using N_p processors. The problem specifications take only a small fraction of the total solve time and hence has a negligible effect on the reported timing results. The wall clock time, however, includes calculating and assembling the elements of the matrix as well as solving the matrix systems. Equation (21) enables us to obtain incremental fixed problem size efficiencies for all the spatial resolutions considered in this article. The run time results for case 1 for three different problem sizes (indicated in brackets) are plotted in Figure 3*a* against the number of processors. Also shown are straight lines that correspond to 100% fixed problem size efficiencies (linear scaling). For the problem with 121^3 grid points we notice near-linear scaling performance. This is due to the robust, scalable, linear solvers provided by PETSc. However, for problem sizes with 37^3 and 61^3 grid points, the efficiencies start to deviate from the linear behavior at high processor number as the local on-processor matrix size becomes too small to be solved efficiently in parallel. A similar plot of run time results versus the number of processors is shown in Figure 3*b* for case 2. Case 2 was run with specified absorption coefficient of 0.01 and $10\,\mathrm{m}^{-1}$, respectively, employing a grid resolution of 121^3 and the SSD_{3b} angular quadrature scheme. The compute time does decrease with increase in absorption coefficient, but comparing the results with the straight lines that correspond to linear scaling reveals that the parallel efficiencies for this model problem do not depend strongly on the absorption coefficient. This independence of the parallel performance with the optical thickness is in contrast to previous studies [4, 5, 8]. The use of a global solver with a robust preconditioner such as block Jocobi does not severely affect solver performance when the radiative properties of the medium change. This is a critical requirement in fire simulations, where the medium properties vary significantly across the domain. The incremental fixed size efficiencies for cases 1 and 2 have been tabulated in Table 2. The numbers of iterations are directly proportional to the solve times and have therefore not been reported. From Figures 3*a* and 3*b* we also

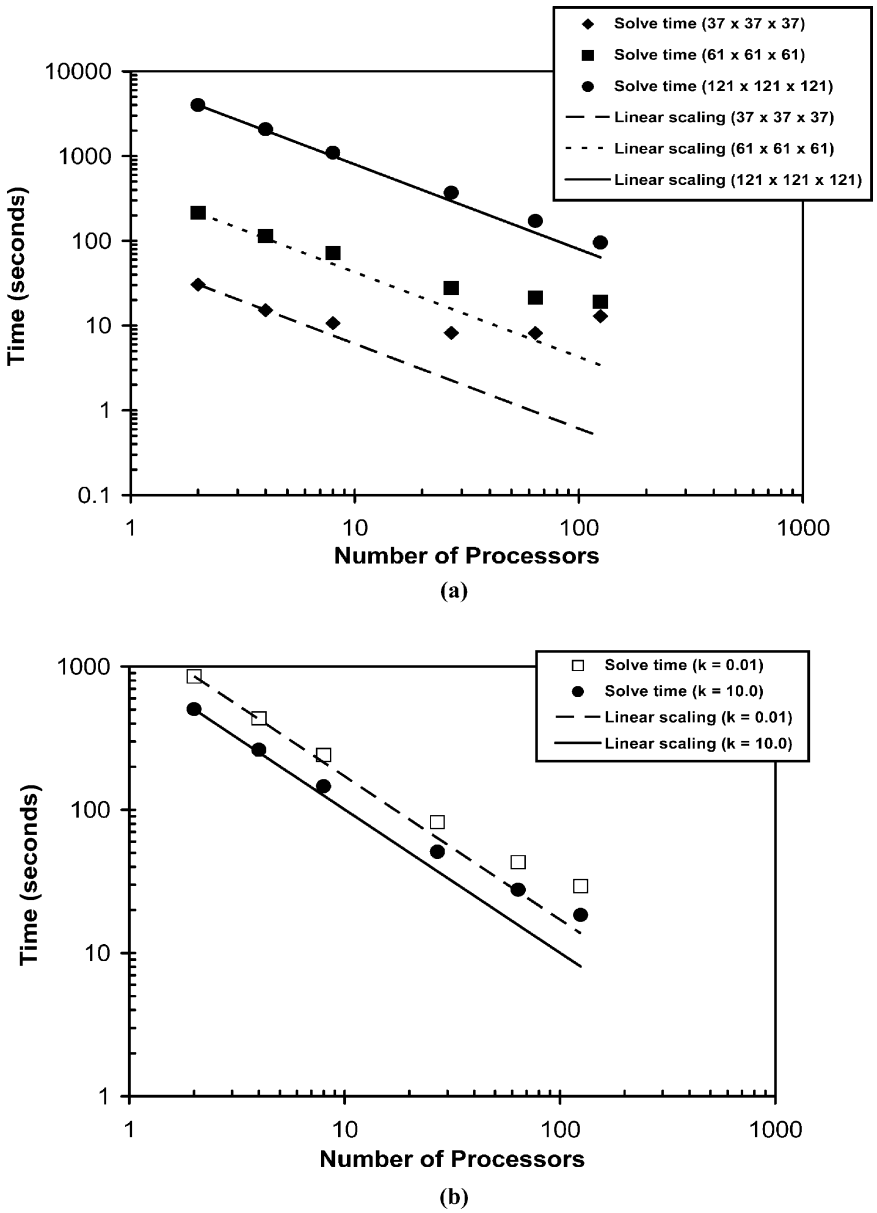

Figure 3. Variations in the run time with the number of processors for (*a*) case 1 run with the SSD_{3b} scheme; (*b*) case 2 run with the SSD_{3b} scheme and 121^3 spatial resolution.

notice that for the same spatial and angular resolution, case 2 takes less time to solve than case 1. This aspect of the radiation solve time will be discussed later in this section.

Another important measure of the parallel efficiency is the scaled problem size, where the size of the local on-processor grid remains constant and therefore the size

Table 2. Incremental fixed problem size efficiencies for cases 1 and 2

	Case 1			Case 2	
N_p	37^3	61^3	121^3	$K=0.01$	$K=10.0$
2	100	100	100	100	100
4	100	94	96	98	96
8	71	74	91	89	86
27	28	57	80	77	73
64	12	31	73	62	57
125	4	18	67	47	44

of the distributed global grid increases with increase in processor number. Theoretically, in scaled problem size problems, as each processor has the same amount of work to do, the solve time should remain constant as the number of processors increases. Thus, the scaled problem size efficiency is defined as

$$\eta_s = \frac{t_g(1)}{t_g(N_p)} \tag{22}$$

Figure 4 shows a plot of the solve time and the number of processors for case 1, represented by open symbols run with scaled problem sizes of 25^3 and 50^3 and the SSD_{3b} angular quadrature scheme. Also shown in Figure 4 are scaled problem size (25^3), radiation solve time results during a transient 10-m heptane pool fire simulation with all the solve times computed at exactly the same stage of fire simulation. First, we notice from Figure 4 that the solve time for all three sets of data points show similar trends (increase in a similar manner). Second, for the fire simulation the

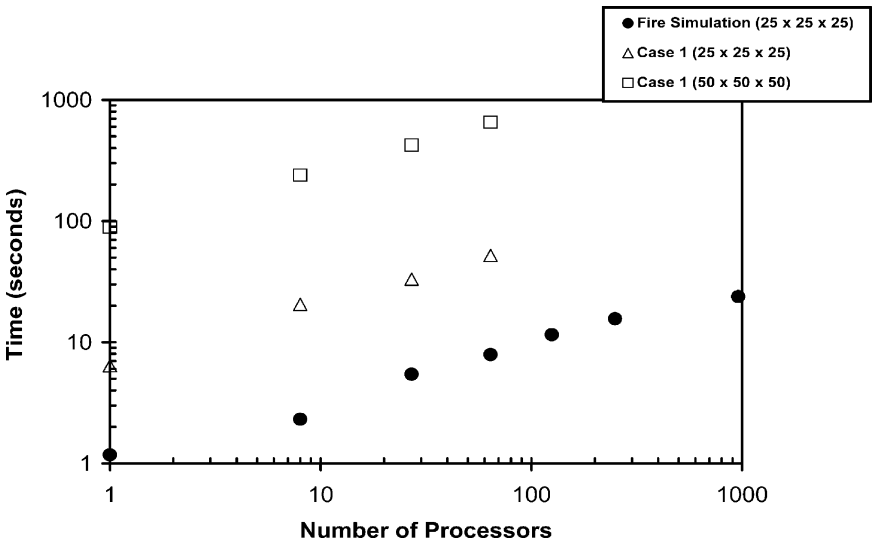

Figure 4. Variation in the run time with the number of processors for scaled problem sizes indicated within brackets of the legend.

solve time increases by about 20 times as the number of processors (and consequently the global problem size) increases 960 times. This corresponds to very useful speed-ups in massively parallel simulations that employ large computational grids. However, we continue to explore ways to improve the scaled problem size efficiencies. The scaled problem size efficiencies for case 1 and the fire simulation have been tabulated in Table 3. We believe that we can extend this algorithm to handle problems with reflecting boundaries as well as scattering media without serious degradation in the parallel efficiencies.

The performance of the GMRES and BiCGSTAB iterative methods [15] during the solution of case 1 employing a block Jacobi preconditioner [where the inverse of the matrix M in Eq. (11) contains the diagonal elements of matrix P] on both is shown in Figure 5*a*. A grid resolution of 61^3 and the SSD_{3b} angular quadrature scheme was used to generate this figure. The BiCGSTAB method is faster for this problem at low processor number, but there appears to be little difference in the relative speeds of the two iterative methods at higher processor numbers. In a similar vein, block Jacobi and point Jacobi preconditioning [where the matrix M in Eq. (11) contains the diagonal elements of matrix P] to the GMRES method are compared in Figure 5*b*. The spatial and angular resolutions are the same as those in Figure 5*a*. Clearly, the point Jacobi preconditioning results in nearly an order-of-magnitude increases in the compute time when compared to block Jacobi preconditioning. The condition number of the matrix $M^{-1}P$ [cf. Eq. (13)] is known to be smaller for block preconditioning than for point preconditioning when the matrix P results from the discretization of the Poisson equation [24]. Smaller condition numbers correspond to faster convergence. We apply the same reasoning to the matrices that arise during the solution of case 1, to explain the differences in compute times between these two preconditioners. Figure 6*a* shows the amount of time spent by the linear solver in each octant of the discrete ordinates method for case 1 solved on a 121^3 domain and the SSD_{3b} angular quadrature scheme with 1 processor as well as 125 processors (indicated in brackets). The direction cosines are all positive and all negative in the first and eighth octants, respectively. Consequently, the sparse nonsymmetric matrices that result during the solution procedure involving these direction cosines have all the three bands that run parallel to the diagonal either completely on the upper triangular portion (for the directions representative of the first octant) or the lower triangular portion (for the directions that are in the eighth octant) [cf. Eqs. (8)

Table 3. Scaled problem size efficiencies for case 1 and the fire simulation

	Case 1		Fire simulation
N_p	25^3	50^3	25^3
1	100	100	100
8	31	37	51
27	19	21	22
64	12	14	15
125			10
250			8
960			5

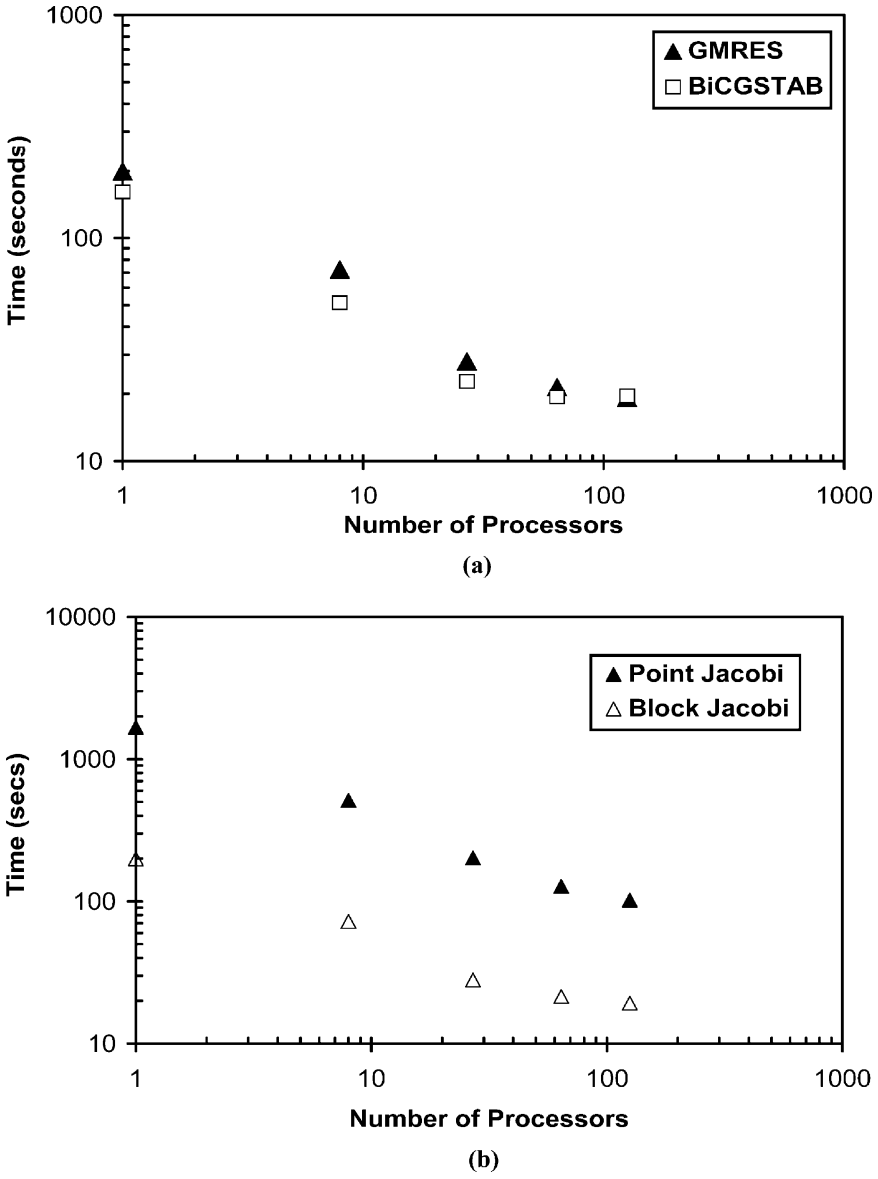

Figure 5. Comparison of solver run times for case 1 run with a spatial resolution 61^3 and the SSD_{3b} scheme: (*a*) GMRES and Bi-CGSTAB linear solvers with block Jacobi preconditioning on both; (*b*) GMRES linear solver with block Jacobi and point Jacobi preconditioning.

and (9)]. The distributions of the radiation solve times for case 1 clearly show that the iterative method takes much less time in solving upper and lower triangular systems than those with bands on both sides of the diagonal. Figure 6*b* shows similar run-time distributions between different octants for case 2 run with specified absorption coefficients of 0.01 and 10, solved on a 121^3 domain, SSD_{3b} angular resolution, with 1 processor and 125 processors, respectively. In comparison to

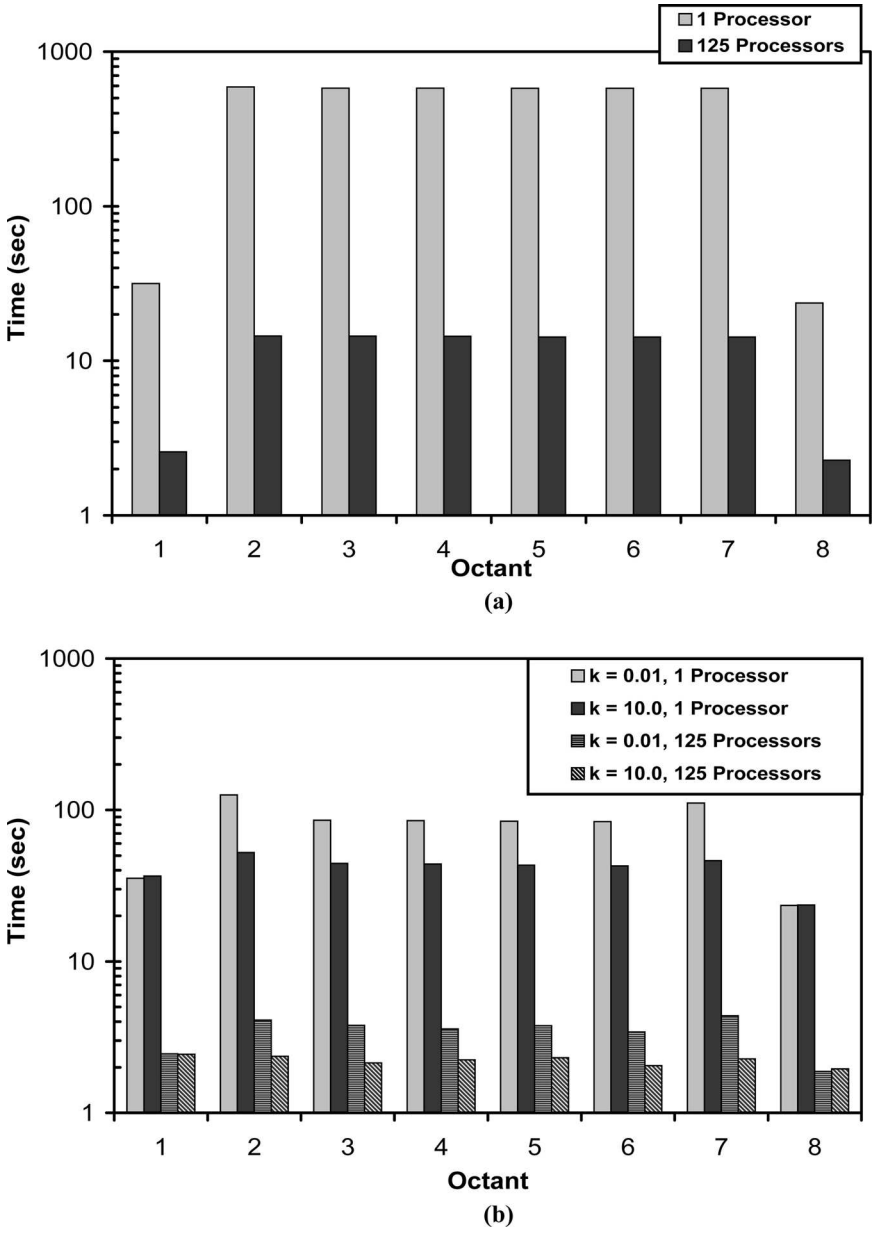

Figure 6. Distribution of time taken to solve the RTE in all the directions corresponding to each octant of the discretized sphere for problem size 121^3 and the SSD_{3b} scheme: (*a*) case 1; (*b*) case 2 with the corresponding absorption coefficient indicated in the figure.

Figure 6*a*, we still notice that the first and eighth octant solve times are less than those of the other octants, but the difference is not as marked as it is in Figure 6*a*. It is important to realize that case 2 is solved on a much larger domain ($50\,\mathrm{m} \times 50\,\mathrm{m} \times 1\,\mathrm{m}$) than case 1 ($1\,\mathrm{m} \times 1\,\mathrm{m} \times 1\,\mathrm{m}$). And since all the geometric

information goes into the diagonal [cf. Eqs. (8) and (9)], the diagonal elements in case 2 are larger and hence more "dominant" than they are in case 1. When the matrix is diagonally dominant the matrix becomes easier to solve and hence there is little difference in the solve time among the various octants. With increasing absorption coefficient the matrix becomes even more diagonally dominant [cf. Eqs. (8) and (9)], and consequently this explains why the solve time decreases with increase in absorption coefficient as observed in Figure 3*b*, as well as the lesser difference observed in Figure 6*b* for the solve time among various octants. This is also the reason behind the difference in the solve times of case 1 and case 2 observed in Figures 3*a* and 3*b* for the same angular and spatial resolution. Effort is currently underway to implement a parallel algebraic multigrid preconditioner. We anticipate the successful implementation of the preconditioner to further reduce radiation solve times and improve parallel performance.

CONCLUSIONS

A domain decomposition paradigm was applied to the radiation component that solves the differential form of the RTE using the discrete ordinates method. Robust, parallel linear solver components developed by third parties were used to solve the sparse, nonsymmetric matrices that result during the solution procedure. The global nature of radiative transfer does cause degradations in the parallel efficiencies with increase in the number of processors, as observed in previous studies that employed spatial decomposition. However, it also provides useful speed-ups in massively parallel combustion simulations that employ large computational grids. For our model problem the parallel efficiencies of our algorithm did not depend strongly on the optical thickness, which contradicts the observations of many previous studies. Numerical experiments have been performed to compare the performances of different iterative techniques and preconditioners for solving the matrix systems that result from these problems. The experiments also helped identify regions of the solution procedure that consumed the most computational time.

REFERENCES

1. L. A. Gritzo, R. D. Skocypec, and T. W. Tong, The Use of High-Performance Computing to Solve Participating Media Radiative Heat Transfer Problems—Results of an NSF Workshop, Sandia Rep. SAND95-0225, Sandia, NM, 1995.
2. S. P. Burns and M. A. Christon, Spatial Domain-Based Parallelism in Large-Scale, Participating-Media, Radiative Transport Applications, *Numer. Heat Transfer B*, vol. 31, pp. 401–421, 1997.
3. S. P. Burns and M. A. Christon, Spatial Domain-Based Parallelism in Large-Scale, Participating-Media, Radiative Transport Applications, Sandia Rep. SAND96-2485, Sandia, NM, 1996.
4. J. Gonçalves and P. J. Coelho, Parallelization of the Discrete Ordinates Method, *Numer. Heat Transfer B*, vol. 32, pp. 151–173, 1997.
5. P. J. Coelho and J. Gonçalves, Parallelization of the Finite Volume Method for Radiation Heat Transfer, *Int. J. Numer. Meth. Heat Fluid Flow*, vol. 9, pp. 388–404, 1999.
6. J. Liu, H. M. Shang, and Y. S. Chen, Parallel Simulation of Radiative Heat Transfer Using an Unstructured Finite-Volume Method, *Numer. Heat Transfer B*, vol. 36, pp. 115–137, 1999.

7. J. Tal, R. Ben-Zvi, and A. Kribus, A High-Efficiency Parallel Solution of the Radiative Transfer Equation, *Numer. Heat Transfer B*, vol. 44, pp. 295–308, 2003.
8. P. J. Novo, P. J. Coelho, and M. G. Carvalho, Parallelization of the Discrete Transfer Method, *Numer. Heat Transfer B*, vol. 35, pp. 137–161, 1999.
9. P. S. Cumber and Z. Beeri, A Parallelization Strategy for the Discrete Transfer Radiation Model, *Numer. Heat Transfer B*, vol. 34, pp. 287–302, 1998.
10. Z. H. Yan, Parallel Computation of Turbulent Combustion and Flame Spread in Fires, *Numer. Heat Transfer B*, vol. 39, pp. 585–602, 2001.
11. J. G. Marakis, J. Chamico, G. Brenner, and F. Durst, Parallel Ray Tracing for Radiative Heat Transfer—Application in a Distributed Computing Environment, *Int. J. Numer. Meth. Heat Fluid Flow*, vol. 11, pp. 663–681, 2001.
12. A. S. Jamaluddin and P. J. Smith, Predicting Radiative Transfer in Rectangular Enclosures Using the Discrete Ordinates Method, *Combustion Sci. Technol.*, vol. 59, pp. 321–340, 1988.
13. M. F. Modest, *Radiative Heat Transfer*, pp. 543–544, McGraw-Hill, New York, 1993.
14. R. Siegel and J. R. Howell, *Thermal Radiation Heat Transfer*, 4th ed., pp. 687–688, Taylor & Francis, New York, 2002.
15. S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith, The Portable Extensible Toolkit for Scientific Computation (PETSc), Version 28, www.mcs.anl.gov/petsc/petsc.html, 2000.
16. G. Strang, *Introduction to Applied Mathematics*, p. 405, Wellesley-Cambridge Press, Wellesley, MA, 1986.
17. P. Hsu and J. T. Farmer, Benchmark Solutions of Radiative Heat Transfer within Nonhomogeneous Participating Media Using the Monte Carlo and YIX Methods, *ASME J. Heat Transfer*, vol. 119, pp. 185–188, 1997.
18. K. D. Lathrop and B. G. Carlson, Discrete Ordinates Angular Quadrature of the Neutron Transport Equation, Los Alamos Sci. Lab. Rep. LA-3186, Los Alamos, NM, 1965.
19. Ben-Wen Li, Hai-Geng Chen, Jun-Hu Zhou, Xin-Yu Cao, and Ke-Fa Cen, The Spherical Surface Symmetrical Equal Dividing Angular Quadrature Scheme for Discrete Ordinates Method, *ASME J. Heat Transfer*, vol. 124, pp. 482–490, 2002.
20. W. A. Fiveland and A. S. Jamaluddin, Three-Dimensional Spectral Radiative Heat Transfer Solutions by the Discrete-Ordinates Method, in *Heat Transfer Phenomena in Radiation, Combustion and Fires: National Heat Transfer Conference, Philadelphia, PA, HTD*, vol. 106, pp. 43–48, ASME, Heat Transfer Division, New York, 1989.
21. G. Brussino and V. Sonnad, A Comparison of Direct and Preconditioned Iterative Techniques for Sparse, Unsymmetric, Systems of Linear Equations, *Int. J. Numer. Meth. Eng.*, vol. 28, no. 4, pp. 801–815, 1989.
22. J. C. Chai, H. S. Lee, and S. V. Patankar, Ray Effect and False Scattering in the Discrete Ordinates Method, *Numer. Heat Transfer B*, vol. 24, pp. 373–389, 1993.
23. S. L. Chang and K. T. Rhee, Blackbody Radiation Functions, *Int. Commun. Heat Mass Transfer*, vol. 11, pp. 451–455, 1984.
24. M. Hegland and P. E. Saylor, Block Jacobi Preconditioning of the Conjugate Gradient Method on a Vector Processor, in D. J. Evans (ed.), *Preconditioned Iterative Methods*, pp. 121–140, Topics in Computer Mathematics, vol. 4, Gordon & Breach, Yverdon, Switzerland, Langhome, PA, 1994.