

PARALLELIZATION OF THE P-1 RADIATION MODEL

Gautham Krishnamoorthy, Rajesh Rawat, and Philip J. Smith

Department of Chemical Engineering, University of Utah, Salt Lake City,
Utah, USA

The P-1 radiation model is spatially decomposed to solve the radiative transport equation on parallel computers. Mathematical libraries developed by third parties are employed to solve the linear systems that result during the solution procedure. Multigrid preconditioning accelerated the convergence of iterative methods. The parallel performance did not depend strongly on the radiative properties of the medium or the boundary conditions. Predictions from coupling the weighted-sum-of-gray-gases model with the P-1 approximation are compared against benchmarks for model problems. The P-1 approximation resulted in only a moderate loss in accuracy while being significantly faster than the discrete ordinates method.

INTRODUCTION

Radiation is often the dominant mode of heat transfer in hydrocarbon pool fires. Performing realistic simulations of fires require the representation of the relevant physical processes such as turbulent reacting flows, heat transfer, and chemistry that occur over a wide range of continuum length scales and their corresponding time scales. Large-eddy simulations (LES) attempt to capture the important features of the fire physics by resolving the large length and time scales that are responsible for controlling the dynamics of the fire. In order to spatially resolve the important flow characteristics in a fire, grids containing 10^6 to 10^8 computational cells need to be used at every time step associated with the calculation. Massively parallel computations are required to perform this task. Performing radiation calculations at these spatial and temporal scales is computationally prohibitive even with parallel computers when the same grid is used for both fluid flow as well as radiation calculations. This is due to the dependence of the radiation calculations at each time step on space, direction, as well as the energy of photons. This article describes the parallelization of the P-1 approximation to the radiative transport equation (RTE) and the

Received 24 May 2005; accepted 13 August 2005.

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 their colleagues at C-SAFE and Dr. Wing K. Yee (currently with Bechtel Corporation) for additional collaboration and assistance related to this effort. The authors would also like to thank Prof. Brent W. Webb, who kindly made his prediction data available.

Address correspondence to Gautham Krishnamoorthy, Fluent Incorporated, 10 Cavendish Court, Centerra Resource Park, Lebanon, NH 03766, USA. E-mail: gk@fluent.com

NOMENCLATURE

a_ϵ, a_κ	emissivity and absorptivity weighting factors	x, y, z	coordinate directions
A	square coefficient matrix	Y	mole fractions of the gaseous species
B	right-hand-side vector of the linear matrix system	ϵ	total emissivity
C_2	constant in the Planck's function	η_f	incremental fixed problem size efficiency
f_v	soot volume fraction	η_s	scaled problem size efficiency
G	incident radiation, W/m^2	κ	absorptive index
k	absorption coefficient, m^{-1}	σ	Stefan-Boltzmann constant
L	geometric distance, m	σ_s	scattering coefficient, m^{-1}
n	real part of the complex index of refraction	Subscripts	
N_p	number of processors	b	blackbody
\mathbf{q}	radiative heat flux vector, W/m^2	i	gray gas component
\mathbf{r}	location vector	P	node at which the incident radiation is being solved
t_g	wall clock time	R	index that runs over all neighboring surface nodes
T	temperature, K	w	wall
V	volume of a grid element, m^3		

feasibility of employing it for performing nongray calculations in LES calculations of pool fires. The P-1 approximation reduces the RTE to a relatively simple partial differential equation which is conceptually simple to solve numerically at each nongray interval [1]. Morvan et al. [2, 3] have demonstrated numerical simulations of pool fires by employing the P-1 radiation model to treat radiative transfer. Their simulations showed the fraction of the total heat released that was radiated away to compare favorably with experimental observations.

Obtaining solutions to the discretized transport equations consumes the maximum time in most radiation algorithms. Therefore, the speed and parallelizability of an algorithm is often determined by the speed and parallelizability of the solver employed in the solution procedure. Hence, it is extremely beneficial to interface the radiation algorithm to libraries of robust, scalable, nonlinear and linear solvers such as the Portable Extensible Toolkit for Scientific Computation (PETSc) [4] and High Performance Preconditioners (HYPRE) [5]. This provides the developer with easy access to a suite of direct and iterative solvers and preconditioners that can be suitably selected depending on the problem being solved. We have previously employed the solvers and preconditioners in PETSc to solve the system of linear equations arising from a finite-volume discretization of the discrete ordinates method [6]. The multigrid options in HYPRE are employed in this study to solve the symmetric matrices that result from a finite-volume discretization of the P-1 transport equation. Multigrid methods are among the best available methods for solving the elliptic transport equation [7].

Timing and scaling studies are first performed using the various multigrid options available in HYPRE to determine the best solver/preconditioner combination for the intended application. As part of a validation methodology, the weighted-sum-of-gray-gases model (WSGGM) [8] is coupled with the P-1 radiation calculations and the computed radiative transfer variables are compared with the

benchmark results available in the open literature. All simulations reported in this work were carried out on a distributed-memory Linux cluster at the University of Utah that has 128 dual processor nodes.

P-1 APPROXIMATION

The transport equation for the P-1 approximation determines the distribution of the incident radiation (G) throughout the domain. If k represents the absorption coefficient, σ_s the scattering coefficient, C the linear-anisotropic phase function, and T the temperature, then the differential equation governing the P-1 approximation can be written as [9]

$$\nabla(\Gamma\nabla G) - kG + 4k\sigma T^4 = 0 \quad (1)$$

where

$$\Gamma = \frac{1}{3(k + \sigma_s) - C\sigma_s} \quad (2)$$

If ε_w and T_w are the emissivities and temperature of the boundaries, then the boundary condition associated with the above equation is

$$q_w = -\frac{\varepsilon_w}{2(2 - \varepsilon_w)}(4\sigma T_w^4 - G) \quad (3)$$

If the absorption coefficient and temperature within the domain are specified or can be obtained along with the boundary properties, then Eq. (1) can be solved iteratively for the irradiation (G) throughout the domain.

The radiative flux vector and the radiative flux divergence at a position vector \mathbf{r} inside the domain can be calculated as [9]

$$\mathbf{q}(\mathbf{r}) = -\Gamma\nabla G \quad (4)$$

$$\nabla \cdot \mathbf{q}(\mathbf{r}) = k(\mathbf{r})[4\sigma T(\mathbf{r})^4 - G(\mathbf{r})] \quad (5)$$

Multiplying both sides of Eq. (1) by $dV = dx dy dz$ and integrating over the volume elements results in the following equation for the face-centered fluxes and the cell-centered incident radiation:

$$\begin{aligned} & \left[\left(\Gamma \frac{\partial G}{\partial x} \right)_e - \left(\Gamma \frac{\partial G}{\partial x} \right)_w \right] dy dz + \left[\left(\Gamma \frac{\partial G}{\partial y} \right)_n - \left(\Gamma \frac{\partial G}{\partial y} \right)_s \right] dx dz \\ & + \left[\left(\Gamma \frac{\partial G}{\partial z} \right)_t - \left(\Gamma \frac{\partial G}{\partial z} \right)_b \right] dx dy - k_P G_P dV + 4k_P \sigma T_P^4 dV = 0 \end{aligned} \quad (6)$$

In Eq. (6), the lowercase subscripts represent face-centered values and the uppercase subscripts represent cell-centered values. The subscript P corresponds to the node at

which Eq. (1) is being approximated. The lowercase subscripts correspond to the interfaces of the control volume under consideration with its neighboring cell to the east (e), west (w), north (n), south (s), top (t), and bottom (b), respectively. A central differencing scheme is employed to evaluate the fluxes (the derivatives). The medium properties (Γ , k , and T) are all cell-centered variables. Γ at a control-volume face is determined as a harmonic mean of the Γ of the cells adjoining that face. Equation (6) can then be represented in the form [10]

$$A_P G_P + \sum_R A_R G_R = B_P \quad (7)$$

where the subscript P corresponds to the node at which Eq. (1) is being approximated, the index R runs over all the neighboring surface nodes in Eq. (6), A_R corresponds to the product of the geometric properties and medium properties, and B_P contains all the terms that are known. Equation (7) can be written in matrix form as

$$AG = B \quad (8)$$

where, A corresponds to a square coefficient matrix which is sparse and symmetric.

If the control volume under consideration is adjacent to a boundary cell(s), then the corresponding flux(s) in Eq. (6) must be replaced by the boundary condition, Eq. (3).

The WSGGM is employed here as a nongray gas model with the P-1 approximation. The temperature-independent absorption coefficients (k_i), the emissivity and absorptivity polynomial coefficients (b_ε , c_α) for the temperature-dependent emissivity, and absorptivity weighting factors (a_ε , a_α) for the i th nongray gas are obtained from Smith et al. [8] for three gray gases and one ‘‘clear gas’’ ($i = 0$). Equation (1) in a nonscattering medium for the i th gray gas may be written as

$$\nabla \left(\frac{1}{3k_i} \nabla G_i \right) - k_i G_i + 4k_i a_{\varepsilon,i}(T) \sigma T^4 = 0 \quad (9)$$

with the boundary conditions

$$q_{w,i} = - \frac{\varepsilon_w}{2(2 - \varepsilon_w)} (4a_{\alpha,i}(T, T_w) \sigma T_w^4 - G_i) \quad (10)$$

The radiative fluxes and the radiative flux divergence for the WSGGM are evaluated as

$$\mathbf{q}(\mathbf{r}) = \sum_i - \frac{1}{3k_i(\mathbf{r})} \nabla G_i(\mathbf{r}) \quad (11)$$

$$\nabla \cdot \mathbf{q}(\mathbf{r}) = \sum_i k_i(\mathbf{r}) \{ 4a_{\varepsilon,i}[T(\mathbf{r})] \sigma T(\mathbf{r})^4 - G_i(\mathbf{r}) \} \quad (12)$$

ALGORITHM DESCRIPTION

We employ HYPRE [5] to efficiently solve the system of linear equations in Eq. (8). HYPRE 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 coefficients of the matrices and vectors are first computed on a global mesh and then passed on to HYPRE. HYPRE then computes and distributes the matrices and vectors among all the processors involved in the simulation. Global indices are used for accessing the distributed data structures also. HYPRE thus offers the benefits of being a general-purpose solver that can easily be interfaced to any model component. For instance, Yee [11] has employed HYPRE to solve the pressure equation in a computational fluid dynamics (CFD) code.

MULTIGRID

Numerical analysis shows that the convergence rates of all basic iterative methods slow down due to the presence of low-frequency (smooth) components in the error [12]. Multigrid methods restore and improve the performance of the basic relaxation schemes by exploiting discretizations with different mesh sizes. The basic idea of multigrid is that the components associated with the smooth modes of error are mapped naturally into high-frequency modes on a coarser mesh. These high-frequency components are then eliminated by the relaxation schemes. This process is repeated recursively using a hierarchy of meshes [13].

The key components of a multigrid algorithm are the smoothing procedure, the coarsening strategy, the coarse-grid operator, transfer operators, and the cycle type [7, 14]. The smoothing property is the ability of relaxation methods such as Gauss-Siedel and Jacobi-type iterative methods to reduce the high-frequency error components. The coarsening strategy involves doubling the mesh size in each direction (standard coarsening) or in only one or two directions (semicoarsening). Semicoarsening is particularly useful in anisotropic problems and has been employed in the multigrid options for structured grids available in HYPRE [5]. The coarse-grid operator defines the problem on the coarse grid. The common choices are either to use the direct analog of the fine-grid problem on the coarse grid or to use the Galerkin coarse-grid operator [7, 14]. The transfer operators are vectors that map vectors between the fine and coarse grids. Restriction operators map fine-grid functions to coarse-grid functions, whereas the interpolation functions map coarse-grid functions to fine-grid functions. A standard multigrid cycle consists of presmoothing, coarse-grid correction step, and postsmoothing.

An important fact which underlies much of the power of the multigrid methods is that the smoothing rate for the basic relaxation schemes is small and independent of the grid spacing. Therefore, the number of iterations remains constant with increased spatial resolution, making multigrid as well as its parallel implementation scalable. However, due to various complications that occur in real-life problems such as anisotropies, the high-frequency components of the error may still not be eliminated by straightforward multigrid techniques. In such situations, multigrid may be employed as a preconditioner in combination with Krylov subspace methods such as the conjugate gradient (CG) method to provide more robustness and address a

larger class of problems [14]. Employing multigrid as a preconditioner to Krylov iterative methods is discussed in [7, 14]. In this work the semicoarsening multigrid (SMG) and the PFMG algorithms in HYPRE are employed as a solver or as a preconditioner to Krylov subspace methods [5].

DESCRIPTION OF CASES

Case 1 consists of a model radiation problem selected to compare the time to converge to a given residual norm (10^{-6}) for various solver and preconditioner options in HYPRE. This problem, which was first introduced by Hsu and Farmer [15], consists of an isothermal unit cube with cold, black walls. The interior of the cube consists of a gray, nonscattering, absorbing/emitting material. A more detailed description of the problem can be found in [15]. Scaling and timing studies have previously been performed for this model problem using parallel implementations of the discrete ordinates method [6, 16].

A homogeneous medium model problem (case 2) was selected next to investigate the effects of geometric anisotropy on the solve times of various solver and preconditioner options in HYPRE. The problem consists of two cold (0 K), black, infinitely parallel plates that confine between them a 1-m-thick isothermal layer of a hypothetical gas at 1,000 K and at a pressure of 1 atm. The absorption coefficient of the gas was set at 1 m^{-1} . This 1-D problem was modeled by employing a 3-D domain taking the geometric length and breadth to be 50 times the height. The four end walls were modeled as black surfaces at 0 K. For the timing studies, cases 1 and 2 were run at a grid of resolution 75^3 and times to converge to a residual norm of 10^{-6} were then compared for different solvers and preconditioner options.

Nongray radiation calculations were performed next with the P-1 radiation model (cases 3 and 4). The Smith et al. [8] WSGGM was used in conjunction with the P-1 radiation model and the computed results compared against numerical benchmarks or predictions from more accurate radiation models. The geometry for cases 3 and 4 is a rectangular enclosure of dimensions $2 \text{ m} \times 2 \text{ m} \times 4 \text{ m}$ with the surrounding walls being black at 300 K. For case 3, the medium within the enclosure is assumed to be a mixture of 10 mole% CO_2 , 20 mole% H_2O , and 70 mole% N_2 . The gas temperature is nonuniform but symmetrical about the centerline of the enclosure and is specified as

$$T = (T_c - T_e)f\left(\frac{r}{R}\right) + T_e \quad [\text{K}]$$

Here, T_c is the gas temperature along the centerline of the enclosure, which increases linearly from 400 K at the inlet ($z = 0$) to a maximum value of 1,800 K at $z = 0.375 \text{ m}$ and then decreases linearly to 800 K at the exit ($z = 4 \text{ m}$). T_e is the exit temperature at $z = 4 \text{ m}$. Inside the enclosure a circular region or a cylinder around the centerline can be imagined, where $R = 1 \text{ m}$ is the radius of the cylinder. The variation of gas temperature inside the circular region of the cross-section enclosure is given by the equation

$$f\left(\frac{r}{R}\right) = 1 - 3\left(\frac{r}{R}\right)^2 + 2\left(\frac{r}{R}\right)^3$$

where r is the distance from the centerline of the enclosure. The gas temperatures outside the circular region, i.e., between the cylinder with radius $R = 1$ m and the walls of the rectangular enclosure, were defined as uniform and equal to the exit temperature, T_e .

For this case, Liu [17] has reported highly accurate numerical solutions of non-gray radiative heat transfer by solving the exact narrow-band averaged RTE using a ray-tracing method. Case 4 consists of a newer calculation published recently by Trivic [18], who employed the Smith et al. WSGGM [8] with the finite-volume radiation model. The geometry and the temperature distribution for case 4 is the same as for case 3. The gas mixture for case 4 consists of a uniform mixture of 10 mole% CO₂, 10 mole% H₂O, and 80 mole% N₂. Cases 3 and 4 were both solved on a grid of resolution $41 \times 41 \times 80$. Cases 5 and 6 were aimed at investigating the importance of soot loading in these calculations. Case 5 consists of an isothermal homogeneous mixture of thickness 1 m consisting of 20 mole% H₂O, 10 mole% CO₂, 3 mole% CO, and the rest N₂ along with a soot volume fraction of 10^{-7} at a temperature of 1,000 K between cold (0 K), black plates. The predicted radiative source term distributions resulting from line-by-line (LBL) calculations for this model problem have been reported by Solovjov and Webb [19]. Case 6 consists of a one-dimensional layer of nonuniform gas/soot mixture with the following distributions of temperatures and species:

$$T(x) = 4,000x(L - x) + 800 \quad [\text{K}]$$

$$Y_{\text{CO}_2}(x) = 0.4x(L - x) + 0.06$$

$$Y_{\text{H}_2\text{O}}(x) = 2Y_{\text{CO}_2}(x)$$

$$f_v(x) = [40x(L - x) + 6] \times 10^{-7}$$

Here, T is the temperature, L is the distance between the plates, Y represents the mole fractions of the components, and f_v is the soot volume fraction. The medium variations for this case are characteristic of nonpremixed diffusion flames produced by an inner fuel core inflow, surrounded by an outer oxidant stream [20]. For cases 5 and 6 the soot absorption coefficient is defined as [1, 21]

$$k_s = \frac{3.72f_v C_0 T}{C_2} \quad [\text{m}^{-1}] \quad (13)$$

where

$$C_0 = \frac{36\pi n \kappa}{(n^2 - \kappa^2 + 2)^2 + 4n^2 \kappa^2} \quad C_2 = 1.4388 \text{ cm K}; \quad n = 1.85; \quad \kappa = 0.22$$

The soot absorption coefficient is added to the absorption coefficient of the nongray gas. C_2 , n , and κ are respectively a constant in the Planck's function, real part of the complex index of refraction, and the absorptive index. The predictions

for the radiative source term distributions across this layer using different property models are available [19, 20]. The one-dimensional cases 5 and 6 were both run on grids of resolution 75^3 . Cases 5 and 6 were modeled employing a 3-D domain, taking the geometric length and breadth to be 50 times the height and modeling the four end walls as black surfaces at 0 K. The surrounding surfaces for all the cases mentioned above are black. The walls are also cold (≤ 300 K) for all the cases except case 6. Hence, emissivity weighting factors (a_e) were employed in the boundary conditions for the cold walls [cf. Eq. (10)] of cases 3–5. The errors associated with the use of emissivity weighting factors in the boundary conditions for cold walls have been previously shown to be small [22]. For case 6, however, two calculations were performed, one employing the emissivity weighting factors and the other employing absorptivity weighting factors (evaluated at a mean gas temperature of 1,467 K) in the boundary condition. For all the cases mentioned above, we also ensured that an overall radiative energy balance was being obeyed over the entire domain as a further verification step. According to this, the total radiative heat source in a radiating medium is always equal to the total radiative heat flux through its boundaries [23]. For all the cases in this article, the residual norm was set at 10^{-6} [5].

RESULTS AND DISCUSSION

The solver run times of various iterative methods and preconditioners or the smoothing procedure (indicated within brackets) during the solution of cases 1 and 2 is shown in Figure 1. Two smoothing options were employed with the PFMG solver, “w-Jacobi” (weighted Jacobi) and “RBGS” (red-black Gauss-Siedel). A set of calculations were also performed for cases 1 and 2 employing the discrete ordinates method (also interfaced to HYPRE libraries) with the SSD_{2a} quadrature set

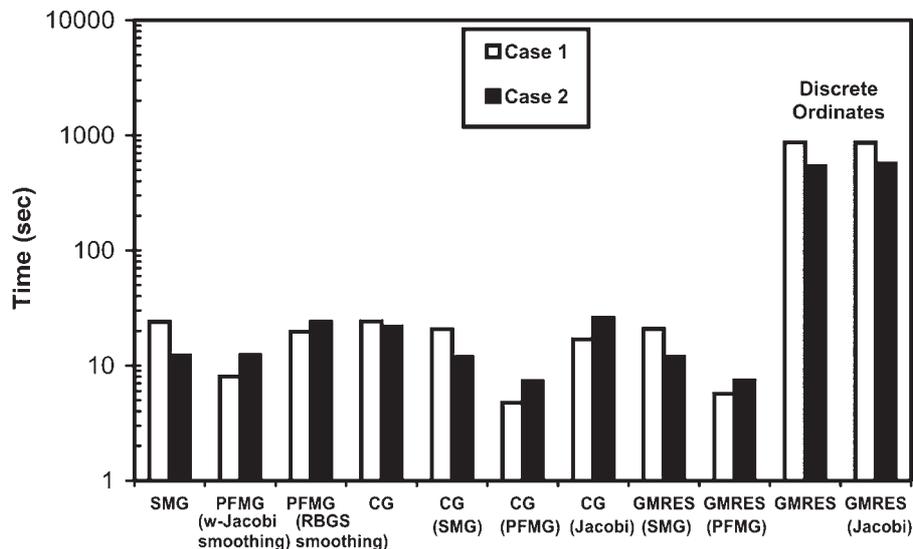


Figure 1. Comparison of solver run times for cases 1 and 2 run at a spatial resolution of 75^3 .

[24] for comparing the relative speeds of these methods. This requires the solution of 48 linear systems of equations compared to just one for the P-1 approximation [cf. Eq. (1)]. In combustion applications, the discrete ordinates method is typically employed to solve the RTE numerically in 24 to 48 directions for obtaining the best combination of accuracy and numerical efficiency [25]. The finite-volume discretization of the discrete ordinates method results in sparse, nonsymmetric matrices [6]. Hence, the generalized minimal residual method (GMRES) option in HYPRE was employed to solve the nonsymmetric matrix systems. We notice that employing PFMG (with weighted Jacobi smoothing) as a preconditioner to the conjugate gradient (CG) method emerges as the fastest solver/preconditioner combination for both cases. Employing this combination with the P-1 radiation model is seen to be more than two orders of magnitude faster than the discrete ordinates method and more than two times faster than employing conjugate gradients alone to solve the P-1 radiation model. Employing the parallel PFMG preconditioner to accelerate the convergence of conjugate gradients was demonstrated by Ashby and Falgout [14] during numerical simulations of groundwater flow through heterogeneous porous media. Along the lines of the study by Ashby and Falgout [14], numerical experiments were next performed to compare the number of iterations to converge to a given residual norm for the conjugate gradient method without any preconditioning (CG), two-step Jacobi preconditioning to the CG method (J2CG), PFMG (with weighted Jacobi smoothing) preconditioning to the CG method (MGCG), and PFMG (with weighted Jacobi smoothing) (referred to from hereon as simply “PFMG”) for case 1. Table 1 shows the variations in the number of iterations taken by the four methods to converge to a residual norm of 10^{-6} with increase in spatial resolution. We notice that the number of iterations to achieve convergence is nearly independent of the problem size in the case of MGCG and PFMG. This independence of the number of iterations for convergence on the grid resolution is true for both standalone multigrid as well as multigrid preconditioned conjugate gradients. The convergence of Krylov iterative methods such as conjugate gradients and preconditioned Krylov iterative methods for solving systems arising from discretized partial differential equations tends to slow down as the system gets larger [13]. This is reflected in Table 1 by the increase in number of iterations to convergence for the CG and J2CG iterative methods.

The parallel performance of these methods was investigated next by examining the scaled problem size efficiencies. Here, the size of the local on-processor grid remains constant and therefore the size of the distributed global grid increases with increase in processor number. Theoretically, in scaled problem size problems, as each

Table 1. Number of iterations to converge to a residual norm of 10^{-6} norm for case 1 run at different spatial resolutions

No. of iterations	CG	J2CG	MGCG	PFMG
15^3	53	21	11	28
35^3	142	51	10	21
55^3	235	80	12	27
75^3	329	111	11	22

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

$$\eta_s = \frac{t_g(1)}{t_g(N_p)} \quad (14)$$

where $t_g(N_p)$ is the wall clock time taken to perform the entire calculation except the problem specification (specifying boundary conditions) using N_p processors. The problem specifications take only a small fraction of the total solve time and hence have 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. The scaled problem size efficiencies of the four solvers (CG, J2CG, MGCG, and PFMG) for the solution of case 1 have been tabulated in Table 2. Scaled problem sizes of 50^3 and 75^3 are considered. First, we notice degradations in the scaled problem efficiencies for all solvers. The multigrid options (MGCG, PFMG), however, perform better than the other solver options. The degradation in the parallel efficiencies of the Krylov methods (CG, J2CG) as well as the multigrid methods may be attributed to the communication between the processors. Though the number of iterations in multigrid algorithms remains nearly constant with increase in problem size (cf. Table 1), the number of communication calls in the semi-coarsening strategy employed in the multigrid algorithms increase as the size of the problem increases [14]. Ashby and Falgout [14] mention the increased communication overhead to be the only impediment to the perfect scalability of the multigrid algorithm. Second, the scaled problem efficiencies increase with increase in problem size. The efficiencies increase with increase in problem size as the size of the local on-processor matrix size is large enough to be solved efficiently in parallel. However, it is important to notice that for a scaled problem size of 75^3 , an efficiency of 12–14% at 64 processors corresponds to only 7 to 8 times increase in the solve time compared to a single processor case [cf. Eq. (14)] for solving a problem which is 64 times larger. This corresponds to very useful speed-ups in massively parallel simulations that employ large computational grids.

In accordance with Burns and Christon [16], we also define incremental fixed problem size efficiency as

$$\eta_f = \frac{2t_g(2)}{N_p t_g(N_p)} \quad (15)$$

Table 2. Scaled problem size efficiencies for case 1 for two different problem sizes per processor

N_p	Problem size (50^3)				Problem size (75^3)			
	MGCG	PFMG	J2CG	CG	MGCG	PFMG	J2CG	CG
1	100	100	100	100	100	100	100	100
8	27	32	21	17	35	32	22	22
27	11	19	4	4	23	24	8	9
64	6	10	2	2	12	14	5	4

The incremental fixed problem size efficiencies of the MGCG solver were examined next as a function of the medium properties. The domain geometry and temperature distribution within the medium and at the boundaries selected for this study were assumed to be the same as in cases 3 and 4. The spatial resolution was set at $81 \times 81 \times 160$. The absorption coefficient (k), scattering coefficient (σ_s), and the wall emissivity (ϵ) were then varied. The corresponding times to solution and incremental fixed problem size efficiencies are reported in Table 3. We notice that the parallel performance is not severely affected by the radiative properties or the boundary conditions of the problem. This is in contrast to previous studies [26–28]. The use of a global solver such as CG with a robust preconditioner such as PFMG 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. However, there is degradation in the parallel performance due to increased communication between processors for large processor numbers. This is particularly noticeable for $N_p > 8$, when the size of the local on-processor matrix size becomes too small to be solved efficiently in parallel.

Figures 2*a* and 2*b* show the distributions of the computed radiative source terms and heat flux vectors, respectively, for case 3 along different lines in the enclosure. Also shown are the benchmark results of Liu [17] employing the statistical narrow-band (SNB) model and discrete ordinates (D.O) calculations previously carried out in [29] for this model problem employing the coefficients from the Smith et al. [8] WSGGM. First, we notice that both D.O as well as the P-1 radiation models qualitatively capture the trends in the source term and heat flux vectors. However, the errors associated with the predictions are greater with the P-1 radiation model. Figures 3*a* and 3*b* show the distributions of the source terms and heat flux vectors for case 4 along different lines in the enclosure. Also shown are the corresponding D.O calculations from [29] and finite-volume (F.V) calculations from Trivic [18]. The F.V calculations, however, were performed with five gray gases as opposed to four gray gases employed in this study, thereby not enabling a direct comparison among the F.V, D.O, and P-1 approximations. Due to the similarities between the D.O and the F.V methods, the results from these two methods are close to each other. The distributions of the radiative source terms in the gas–soot mixtures of case 5 and case 6 are shown in Figures 4*a* and 4*b*, respectively. Also shown in

Table 3. Solve times and incremental fixed problem size efficiencies for MGCG for different medium properties and boundary conditions

K (m^{-1})	0.1	10.0	0.5	0.5	0.1	10.0	0.5	0.5
σ (m^{-1})	0	0	0.5	0.5	0	0	0.5	0.5
ϵ	1.0	1.0	1.0	0.7	1.0	1.0	1.0	0.7
N_p	Solve time (s)				Efficiencies			
2	12.7	10.8	13.3	15.4	100	100	100	100
4	9.4	9.2	10.7	9.4	68	59	62	82
8	5.0	4.9	5.8	6.0	63	55	57	64
27	7.7	7.1	8.7	7.6	12	11	11	15
64	10.7	9.8	15.3	14.8	4	3	3	3

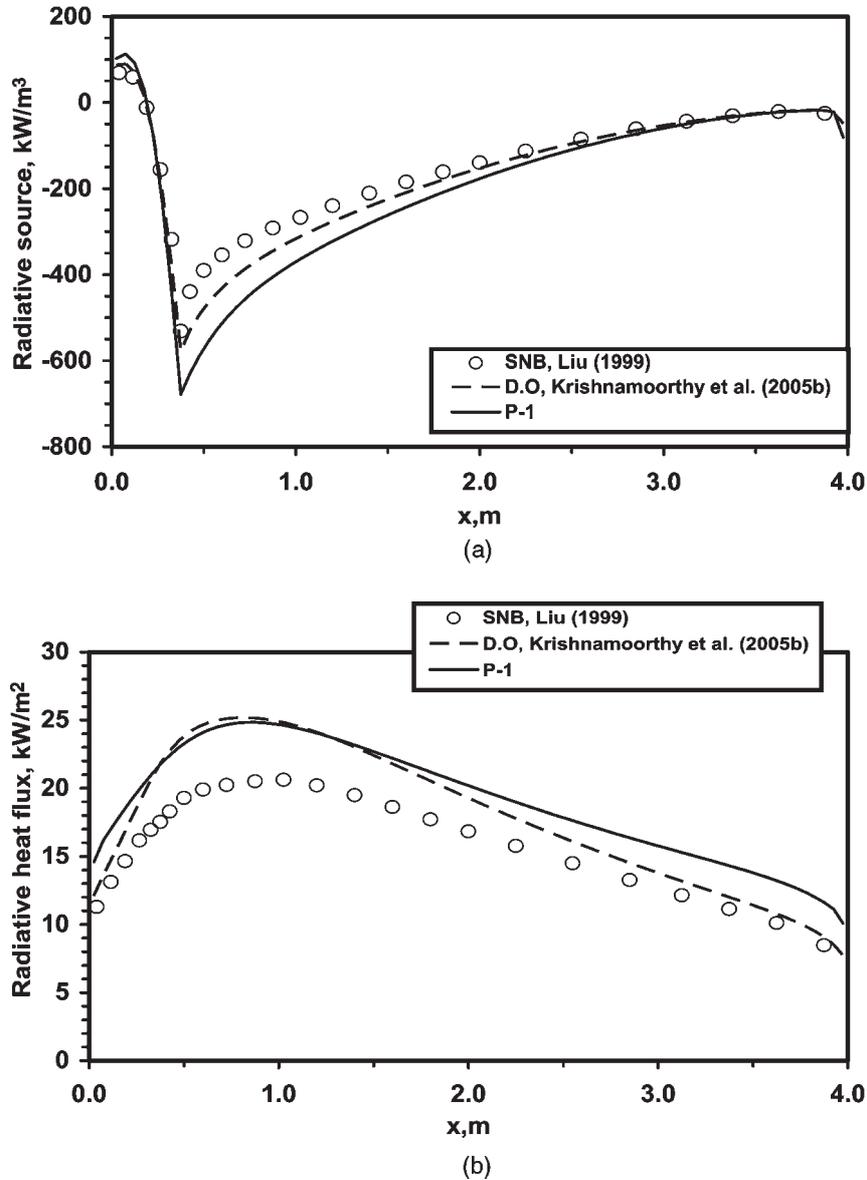


Figure 2. Predictions of: (a) source term distributions along the centerline of the enclosure for case 3; (b) wall heat flux distributions along (2 m, 1 m, z) for case 3.

Figures 4a and 4b are predictions previously made [29] employing the discrete ordinates radiation model employing the coefficients from the Smith et al. [8] WSGGM. For case 5, both the P-1 radiation model and the discrete ordinates method with the WSGGM come very close to predicting the source term distribution at the center of the medium against the line-by-line (LBL) calculations of Solovjov and Webb [19]. For case 6 (which closely represents the conditions in a nonpremixed diffusion flame)

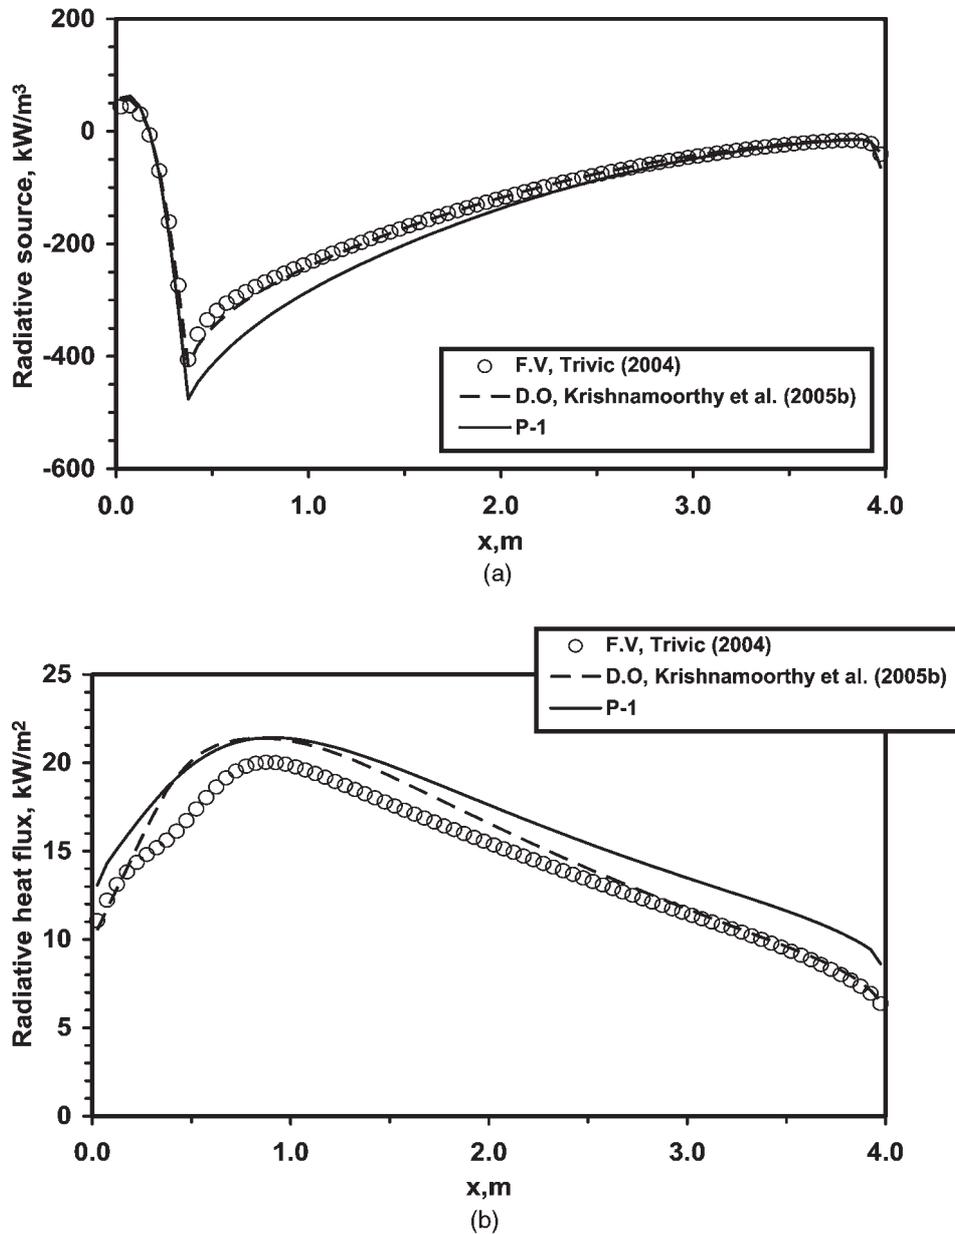


Figure 3. Predictions of: (a) source term distributions along the centerline of the enclosure for case 4; (b) wall heat flux distributions along (2 m, 1 m, z) for case 4.

we notice that employing the WSGGM with the P-1 radiation model and discrete ordinates method adequately captures the trends in the radiative source term distributions. The results, however, deviate from the “benchmark” statistical narrow-band model (SNB) predictions of Bressloff [20], which were generated using RADCAL and the

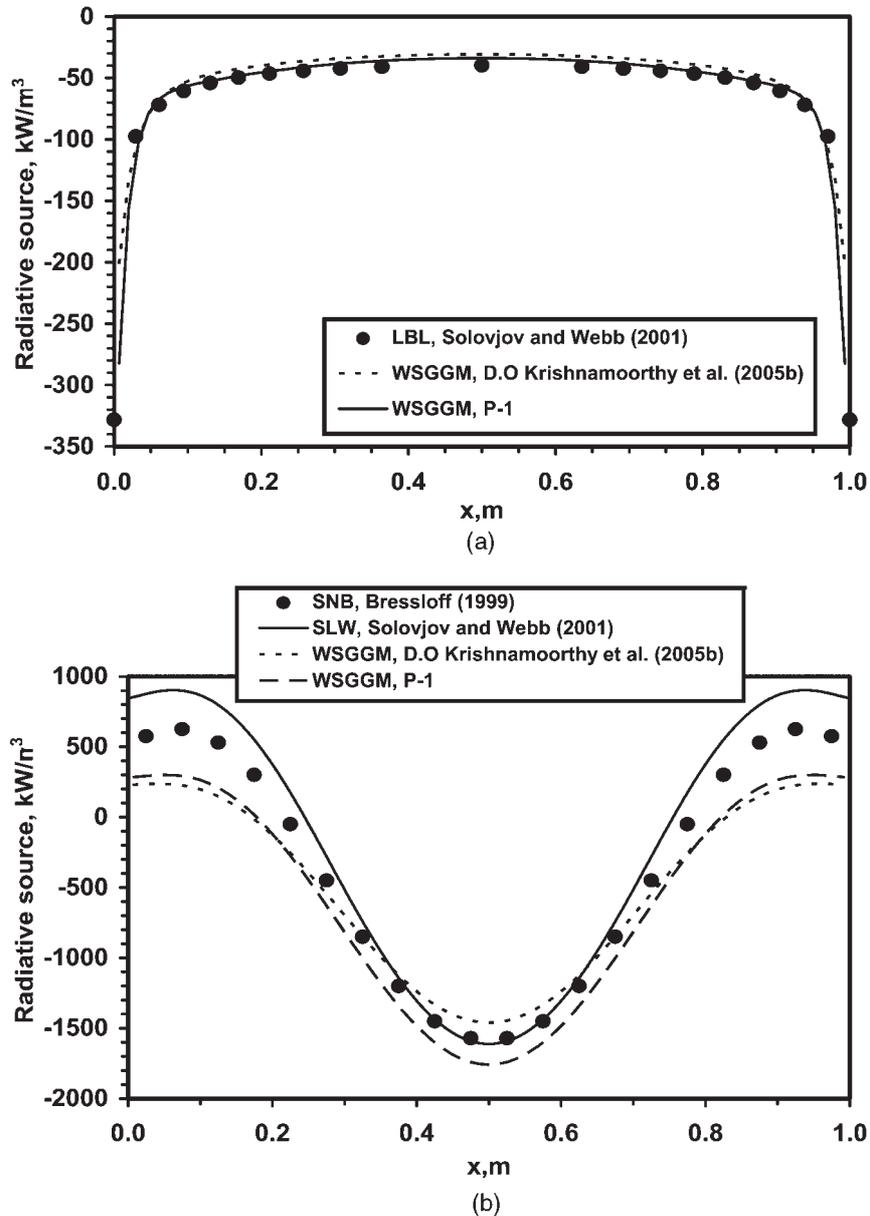


Figure 4. Predicted source term distributions for: (a) case 5; (b) case 6.

spectral-line weighted-sum-of-gray gases (SLW) model predictions of Solovjov and Webb [19]. An additional calculation was performed for case 6 using the WSGGM but this time employing the absorptivity weighting factors in the boundary conditions evaluated at an average medium temperature of 1,467 K. However, the computed values (not shown in Figure 4b) did not differ significantly from the corresponding

profile obtained with the emissivity weighting factors in the boundary conditions of the WSGGM which is shown in Figure 4b. Emission and absorption in case 6 are dominated by the soot continuum. However, the model for the soot properties employed by us with the WSGGM [cf. Eq. (13)] is different from the ones employed by Bressloff [20] and by Solovjov and Webb [19]. Second, the inaccurate representation of gas absorptivities by the WSGGM may be the reason for the differences in the predictions of the radiative source term, especially near the colder regions of the domain (near the walls). This has also been observed in previous studies [22]. Hence, the discrepancies in the predictions of the various studies may be attributed to differences in the solution methods for the RTE, different spectral databases employed in generating the various model parameters, and/or inaccuracies in the models itself. From cases 3–6 we see that employing the WSGGM with the P-1 radiation model results in only a moderate loss of accuracy in the predictions of radiative heat fluxes and the source terms when compared to using the discrete ordinates method with the WSGGM. This is likely due to the near-isotropic intensity distributions in the medium for cases 3–6. The P-1 radiation model is usually accurate in media with near-isotropic intensity distributions [1]. The directional intensity at any given point inside the medium is due to radiation originating from the boundaries and radiation emanating from the medium (due to emission and in-scattering). Intensity from within a nonscattering medium generally shows a very slow directional dependence, as emission results in an isotropic intensity source [1]. Cases 3–5 had boundaries at a uniform surface temperature much lower than that of the emitting medium. Hence the boundary conditions are not expected to significantly affect the intensity distributions or contribute to anisotropic intensities within the medium. Although the boundaries of case 6 were at a uniform temperature of 800 K, the medium was also optically thick, so the boundaries are not expected to cause any anisotropy in the intensity distributions about a location in the medium. Cases 3–6 were also run employing the discrete ordinates method at different angular resolutions (not shown in Figures 2–4). There was very little difference in the results, which further suggested a near-isotropic intensity distribution about a location in the media. The radiation component developed here is intended for use in numerical simulations of open pool fires where radiation within the medium is dominated by emission, and scattering is neglected. For the radiation boundary conditions in this application we assume that the enclosure wall is a black surface at ambient temperature. Hence, the boundaries are not expected to contribute to any anisotropy in the intensities and near-isotropic distributions of intensities are expected to prevail within the fire. However, the P-1 approximation is known to be inaccurate when an optically thin medium (for instance, air) acts as a radiation barrier between a hot emitting medium (fire) and a cold boundary (wall) [1]. This may be due to the fact that the intensity about a location is no longer isotropic as we move away from the flame, due to the presence of a strong emitting source on one side and a cold medium on the other. Hence, the P-1 approximation may lose its accuracy in predicting the radiative heat fluxes to distant objects away from the fire. However, it is important to note that if solid/liquid interface energy balances are not being performed, it is only through the distributions of the radiative source terms $[-\nabla \cdot \mathbf{q}(\mathbf{r})]$ that radiation is coupled to other physical processes associated with the fire (through the source term of the fluid energy balance equation). As seen from cases 3–6, the P-1 approximation is expected to predict the distributions of the radiative source

terms within the fire accurately, thereby providing adequate coupling between radiation and the hydrodynamics. Hence, the computational savings associated with the use of P-1 radiation model will prove to be particularly valuable in transient simulations of pool fires that are performed employing very fine grids to validate the hydrodynamic models.

CONCLUSIONS

A domain decomposition paradigm was applied to a radiation component that solves the transport equation of the P-1 radiation model. Robust, parallel solver components developed by third parties were employed to solve the sparse, symmetric matrices that result during the solution procedure. This provided us with a wide range of solver and preconditioner options. By a proper choice of solver and multi-grid preconditioner, the P-1 radiation model was seen to be more than two orders of magnitude faster than the discrete ordinates method. Parallel multigrid solvers and preconditioners were found to scale well, as the number of iterations to convergence is nearly independent of the problem size. The parallel performance did not depend strongly on the radiative properties of the medium or the boundary conditions. The communications overhead between the processors caused degradations in the parallel efficiencies, but useful speed-ups can be obtained in massively parallel combustion simulations that employ large computational grids.

Nongray calculations for model problems performed by employing the WSGGM to the discrete ordinates method and the P-1 radiation model seem to indicate that the P-1 radiation model may provide large savings in computational time in fire simulations, with only moderate losses in accuracies.

REFERENCES

1. M. F. Modest, *Radiative Heat Transfer*, 2nd ed., Academic Press, New York, 2003.
2. D. Morvan, B. Porterie, J. C. Loraud, and M. Larini, Numerical Simulation of a Methane/Air Radiating Turbulent Diffusion Flame, *Int. J. Numer. Meth. Heat Fluid Flow*, vol. 10, pp. 196–227, 2000.
3. D. Morvan, B. Porterie, J. C. Loraud, and M. Larini, A Numerical Investigation of Cross Wind Effects on a Turbulent Buoyant Diffusion Flame, *Combustion Sci. Technol.*, vol. 164, pp. 1–35, 2001.
4. S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith, The Portable Extensible Toolkit for Scientific Computing (PETSc) version 2.8, www.mcs.anl.gov/petsc/petsc.html, 2000.
5. E. Chow, A. J. Cleary, and R. D. Falgout, *Design of the Hypre Preconditioner Library*, Lawrence Livermore National Laboratory, UCRL-JC-132025, Livermore, CA, 1998.
6. G. Krishnamoorthy, R. Rawat, and P. J. Smith, Parallel Computations of Radiative Heat Transfer Using the Discrete Ordinates Method, *Numer. Heat Transfer B*, vol. 47, pp. 19–38, 2005a.
7. U. Trottenberg, C. Oosterlee, and A. Schuller, *Multigrid*, Academic Press, Los Angeles, 2001.
8. T. F. Smith, Z. F. Shen, and J. N. Friedman, Evaluation of Coefficients for the Weighted Sum of Gray Gases Model, *ASME J. Heat Transfer*, vol. 104, pp. 602–608, 1982.
9. FLUENT/UNS & RAMPANT 4.2—User's Guide Volume 2, Lebanon, NH, 1996.
10. S. V. Patankar, *Numerical Heat Transfer and Fluid Flow*, Hemisphere, New York, 1980.

11. W. K. Yee, Validation Science for Large Scale Pool Fire Simulation, Ph.D. thesis, University of Utah, Salt Lake City, UT, 2004.
12. W. L. Briggs, *A Multigrid Tutorial*, SIAM, Philadelphia, 1987.
13. Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd ed., SIAM, Philadelphia, 2003.
14. S. F. Ashby and R. D. Falgout, A Parallel Multigrid Preconditioned Conjugate Gradient Algorithm for Groundwater Flow Simulations, *Nuclear Sci. Eng.*, vol. 124, pp. 145–159, 1996.
15. 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.
16. 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.
17. F. Liu, Numerical Solutions of Three-Dimensional Non-Grey Gas Radiative Heat Transfer Using the Statistical Narrow-Band Model, *J. Heat Transfer*, vol. 121, pp. 200–203, 1999.
18. D. N. Trivic, Modeling of 3-D Non-Gray Gases Radiation by Coupling the Finite Volume Method with Weighted Sum of Gray Gases Model, *Int. J. Heat Mass Transfer*, vol. 47, pp. 1367–1382, 2004.
19. V. P. Solovjov and B. W. Webb, An Efficient Method for Modeling Radiative Transfer in Multicomponent Gas Mixtures with Soot, *ASME J. Heat Transfer*, vol. 123, pp. 450–457, 2001.
20. N. W. Bressloff, The Influence of Soot Loading on Weighted Sum of Grey Gases Solutions to the Radiative Transfer Equation across Mixtures of Gases and Soot, *Int. J. Heat Mass Transfer*, vol. 42, pp. 3469–3480, 1999.
21. M. J. Yu, S. W. Baek, and J. H. Park, An Extension of the Weighted Sum of Gray Gases Non-Gray Radiation Model to a Two Phase Mixture of Non-Gray Gas with Particles, *Int. J. Heat Mass Transfer*, vol. 43, pp. 1699–1713, 2000.
22. A. Soufiani and E. Djavdan, A Comparison between Weighted Sum of Gray Gases and Statistical Narrow-Band Radiation Models for Combustion Applications, *Combustion & Flame*, vol. 97, pp. 240–250, 1994.
23. G. Li and M. F. Modest, A Method to Accelerate Convergence and to Preserve Radiative Energy Balance in Solving the P_1 Equation by Iterative Methods, *ASME J. Heat Transfer*, vol. 124, pp. 580–582, 2002.
24. 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.
25. B. R. Adams, Computational Evaluation of Mechanisms Affecting Radiation in Gas- and Coal-Fired Industrial Furnaces, Ph.D. thesis, University of Utah, Salt Lake City, UT, 1993.
26. J. Gonçalves and P. J. Coelho, Parallelization of the Discrete Ordinates Method, *Numer. Heat Transfer B*, vol. 32, pp. 151–173, 1997.
27. 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.
28. 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.
29. G. Krishnamoorthy, R. Rawat, and P. J. Smith, Parallel Computations of Nongray Radiative Heat Transfer, *Numer. Heat Transfer B*, vol. 48, pp. 191–211, 2005b, 2005.