JOURNAL OF THERMOPHYSICS AND HEAT TRANSFER
Vol. 13, No. 4, October-December 1999

Coupled Ordinates Method for Multigrid Acceleration of Radiation Calculations

S. R. Mathur*
Fluent, Inc., Lebanon, New Hampshire 03766-1442
and
J. Y. Murthy†
Carnegie-Mellon University, Pittsburgh, Pennsylvania 15213

The finite volume and discrete ordinates methods are known to converge increasingly slowly as the optical thickness is increased. This is a result of the sequential nature of the solution procedure; the equations for energy and the directional intensities are solved one by one, assuming prevailing values for other variables. The coupled ordinates method is proposed whereby the discrete energy and intensity equations at each cell are solved simultaneously, assuming spatial neighbors to be known. The point-coupled procedure is used as a relaxation sweep in a multigrid scheme. The formulation of coarse-level discrete equations admits arbitrary nonconvex polyhedra, making the scheme suitable for use with arbitrary unstructured polyhedral meshes. The scheme is shown to substantially accelerate convergence over a range of optical thicknesses.

Nomenclature

$\boldsymbol{A}$	= area vector
I^B	= blackbody intensity
I^l	= intensity in direction l
k	= conductivity
N_θ, N_ϕ	= polar and azimuthal discretization
S_h	= volumetric heat source in energy equation
s	= coordinate along ray
$\boldsymbol{s}$	= ray direction vector
T	= temperature
$\boldsymbol{V}$	= velocity vector
x_j	= coordinate direction
$\Delta\mathcal{V}$	= volume of control volume
θ	= polar angle
κ	= absorption coefficient
σ_s	= scattering coefficient
Φ	= scattering phase function
ϕ	= azimuthal angle
Ω	= solid angle

Introduction

OVER the past few years, the finite volume has been gaining increasing acceptance in the heat transfer community as a robust and accurate method for simulating radiation heat transfer. Proposed originally by Raithby and Chui,[1] and Chui and Raithby,[2] and extended and modified by a number of researchers,[3,4] the finite volume method and its close relative, the discrete ordinates method,[5] today are being used to compute radiation transfer in gas turbine combustion, furnace heat transfer, packed beds, automotive glass manufacture, optical fiber processing, and other areas. These applications span a wide range of optical thicknesses. In the manufacture of automotive windshields, for example, the glass is optically thick in the near infrared region but nearly transparent otherwise. It is necessary to ensure that the finite volume method performs optimally over this range of parameters.

Received 20 November 1998; presented as Paper 99-0872 at the 37th Aerospace Sciences Meeting, Reno, NV, 11–14 January 1999; revision received 1 April 1999; accepted for publication 9 April 1999.

*Senior Member of Technical Staff. Member AIAA.
†Associate Professor, Department of Mechanical Engineering. Member AIAA.

Though issues of accuracy have been addressed quite thoroughly in the published literature, relatively little attention has been paid to the convergence behavior of the finite volume method. The finite volume and discrete ordinates methods are known to converge increasingly slowly as the optical thickness is increased[5]; at optical thicknesses greater than 10, these methods may take many hundreds of iterations more than at optical thickness of less than unity. The primary culprit is the sequential nature of the solution algorithm. Typically, for participating radiation, the energy equation and the intensity equations in the different directions are solved sequentially, using prevailing values of the other variables to resolve coupling and nonlinearity. When the optical thickness is large, the energy equation encounters large source terms involving the directional intensities; similarly, a large blackbody emission term appears in the equations for the directional intensities. A sequential solution procedure makes for loose coupling between equations and causes oscillation and slow convergence. Similar problems with interequation coupling arise for large scattering coefficients and at reflecting, Fresnel, and rotationally periodic boundaries; here, the cause of slow convergence is the strong coupling between intensities in different directions.

There have been some attempts at convergence acceleration for the discrete ordinates method in neutron transport.[6] Chui and Raithby[7] proposed a multiplicative correction of the average intensity in the context of the finite volume scheme. Fiveland and Jessee[8] proposed and evaluated a number of acceleration strategies in the discrete ordinates context. These included the successive over-relaxation method, the mesh rebalance method, and the synthetic acceleration method. The mesh rebalance method was very similar to the method of Chui and Raithby[7] and was found to perform the best, but its performance deteriorated as the mesh-based optical thickness decreased. The method was modified to perform rebalance on a coarser mesh than that for the actual solution to keep the rebalance mesh optical thickness greater than unity; with this modification it appeared to perform well. These methods continue to use a sequential solution procedure, with the solution of additional variables to accelerate convergence.

In this paper, we present a convergence acceleration procedure called the coupled ordinates method (COMET) based on the multigrid scheme.[9] In keeping with the multigrid idea, the solution is carried out on a sequence of nested meshes formed by agglomerating the cells of the finest mesh. A V cycle is used to prolongate and restrict residuals between mesh levels. The energy and intensity equations are discretized on the finest mesh using the unstructured mesh

method described by Murthy and Mathur.[4] The discrete equations on the coarser meshes are derived by a mixed geometric/additive-correction strategy designed to allow the method to be used with arbitrary unstructured polyhedral meshes. At each mesh level, cells are visited one by one. At each cell, the energy and intensity corrections are solved in a point-coupled fashion by the inversion of a local matrix, with values at spatial neighbors being assumed known at prevailing values. Similar procedures are implemented at boundaries. The method is tested for three radiation problems involving absorption, emission, and scattering and is shown to substantially improve convergence rates.

Governing Equations

Radiative Transfer

The radiative transfer equation (RTE) for a gray absorbing, emitting, and scattering medium in the direction $\boldsymbol{s}$ may be written as[10]

$$\nabla \cdot [I(\boldsymbol{s})\boldsymbol{s}] = -(\kappa + \sigma_s) I(\boldsymbol{s}) + B(\boldsymbol{s}) \tag{1}$$

where

$$B(\boldsymbol{s}) = \kappa I^B + \frac{\sigma_s}{4\pi} \int_{4\pi} I(\boldsymbol{s}') \Phi(\boldsymbol{s}', \boldsymbol{s})\, \mathrm{d}\Omega' \tag{2}$$

For simplicity we consider only isotropic scattering so that $\Phi = 1$.

Energy Equation

For problems in which radiative equilibrium does not exist, it is necessary to solve the energy equation, which can be written as

$$\frac{\partial}{\partial t}(\rho E) + \nabla \cdot (\rho \boldsymbol{V} E) = \nabla \cdot (k \nabla T) + \nabla \cdot (\tau \cdot \boldsymbol{V}) - \nabla \cdot (p\boldsymbol{V}) + S_r + S_h \tag{3}$$

Here, k is the thermal conductivity, τ is the stress tensor, p is the pressure, E is the total energy

$$E = e(T) + (\boldsymbol{V} \cdot \boldsymbol{V}/2) \tag{4}$$

and e is the internal energy. The radiative source term S_r is given by

$$S_r = \kappa \int_{4\pi} [I(\boldsymbol{s}) - I^B]\, \mathrm{d}\Omega \tag{5}$$

S_h contains all other volumetric energy sources.

Boundary Conditions

For the purposes of this paper, all boundaries are assumed to be gray diffuse, though this is not necessary. Under this assumption, the boundary intensity I_b for all outgoing directions ($\boldsymbol{s} \cdot \hat{\boldsymbol{e}}_{\boldsymbol{n}} < 0$) is given by

$$I_b = \frac{(1-\varepsilon)}{\pi} \int_{\boldsymbol{s} \cdot \hat{\boldsymbol{e}}_{\boldsymbol{n}} > 0} I(\boldsymbol{s}) \boldsymbol{s} \cdot \hat{\boldsymbol{e}}_{\boldsymbol{n}}\, \mathrm{d}\Omega + \varepsilon I_b^B \tag{6}$$

where ε is the boundary emissivity. The unit vector $\hat{\boldsymbol{e}}_{\boldsymbol{n}}$ is the surface normal pointing out of the domain.

Discretization of RTE

Details of the spatial and angular discretization have been presented elsewhere[4]; we present only a brief summary here. The spatial domain is discretized into N arbitrary unstructured convex polyhedral cells. The angular space 4π at any spatial location is discretized into discrete nonoverlapping control angles ω_l, the centroids of which are denoted by the direction vector $\boldsymbol{s}_l$ and the polar and azimuthal angles θ_l and ϕ_l. Each octant is discretized into $N_\theta \times N_\phi$ solid angles, each with a constant angular extent of $\Delta\theta$ and $\Delta\phi$ to yield a total of $L (\equiv 8 \times N_\theta \times N_\phi)$ discrete directions. All unknowns are stored at cell and boundary face centroids.

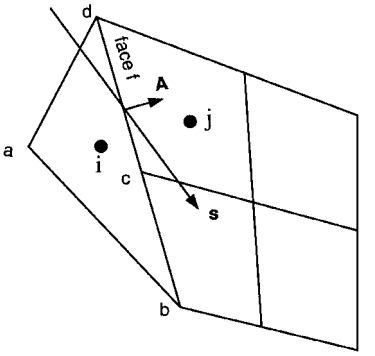

Fig. 1 Control volume.

For each direction l, Eq. (1) is integrated over the the control volume i in Fig. 1 and the solid angle ω_l to yield discrete equations of the form

$$\sum_f J_f^l I_f^l |\boldsymbol{A}| = \left[-(\kappa + \sigma_s)\, I_i^l + \kappa I_i^B\right] \omega_l \Delta\mathcal{V}_i + \omega_l \Delta\mathcal{V}_i \frac{\sigma_s}{4\pi} \sum_{k=1}^{L} \left(I_i^k \omega_k\right) \tag{7}$$

where $\Delta\mathcal{V}_i$ is the volume of the cell i and $J_f^l I_f^l$ is the radiative flux leaving the face f. Using a step scheme, the intensity at an interior face is obtained by upwinding the intensity from the appropriate cell. For example, if the entire solid angle is outgoing at the face f in Fig. 1, we obtain

$$J_f^l = \hat{\boldsymbol{e}}_{\boldsymbol{n}} \cdot \int_{\Delta\theta_l} \int_{\Delta\phi_l} \boldsymbol{s} \sin\theta\, \mathrm{d}\theta\, \mathrm{d}\phi \tag{8}$$

$$I_f^l = I_i^l \tag{9}$$

In general, however, the control angle boundaries do not coincide with the control volume face, and control angle overhang results.[4] In this case, the radiative flux at a face has the general form

$$J_f^l I_f^l = J_{f,\mathrm{out}}^l I_i^l + J_{f,\mathrm{in}}^l I_j^l \tag{10}$$

where I_j^l is the intensity in the neighbor cell j. Here out and in are defined with respect to the cell i in Fig. 1. Murthy and Mathur[4] describe a pixelation approach for computing the coefficients $J_{f,\mathrm{out}}^l$ and $J_{f,\mathrm{in}}^l$ in the presence of control angle overhang. Essentially, it involves determination of the incoming and outgoing fractions of the control angle by dividing the solid angle into smaller control angles for the purposes of numerical integration.

At a boundary face, the intensity I_b^l of radiation from the boundary to cell i replaces I_j^l in Eq. (10). At gray diffuse boundaries, I_b^l is independent of l and is given by

$$I_b = \frac{(1-\varepsilon)}{\pi} \sum_{k=1}^{L} J_{b,\mathrm{out}}^k I_i^k + \varepsilon I_b^B \tag{11}$$

Here the first term is the reflected portion of energy incident on the boundary. At symmetry or mirror boundaries, the intensity I_b^l in a direction l outgoing from the boundary is obtained from the cell intensity in the reflection direction l':

$$I_b^l = I_b^{l'} \tag{12}$$

Using the step scheme, we may write

$$I_b^l = I_i^{l'} \tag{13}$$

An important step in the discretization process is the linearization of the T^4 terms that result from emission in the volume as well as at the boundary. Following Patankar,[11] we write

$$T^4 = T^{*4} + \left.\frac{\partial T^4}{\partial T}\right|_* (T - T^*)$$

$$= 4T^{*3}T - 3T^{*4} \tag{14}$$

where T^* is the current iterate of T.

Assembling the various components, we obtain a set of linear equations of the following form at each cell i for each direction l:

$$\sum_{j=1}^{N} \mu_{ij}^l I_j^l + \sum_{k=1}^{L} \eta_i^{lk} I_i^k + \alpha_i^l I_i^l + \beta_i^l T_i = \gamma_i^l \sum_{k=1}^{L} I_i^k \omega_k + S_i^l + S_i^B \tag{15}$$

where

$$\alpha_i^l = (\kappa + \sigma_s)\omega_l \Delta\mathcal{V}_i, \qquad \beta_i^l = (-4\kappa\sigma/\pi)T_i^{*3}\omega_l \Delta\mathcal{V}_i$$

$$\gamma_i^l = (\sigma_s/4\pi)\omega_l \Delta\mathcal{V}_i, \qquad S_i^l = (-3\kappa\sigma/\pi)T_i^{*4}\omega_l \Delta\mathcal{V}_i$$

The coefficients μ_{ij}^l in Eq. (15) result from the discretization of the $\nabla \cdot (I\boldsymbol{s})$ term in Eq. (1) and the use of Eq. (10). They denote the influence of neighboring cell intensities in the l direction on the intensity in cell i in the same direction. The coefficients η_i^{lk} denote the influence of other directions on direction l for the cell i due to diffuse reflections at any boundary faces [Eq. (11) or (12)]. This term is zero for interior cells and for boundary cells bounded by black walls. The $\alpha_i^l I_i^l$ term results from absorption and out scattering, whereas the $\beta_i^l T_i$ and S_i^l terms result from the linearization of the emission term. The term containing γ_i^l results from the discretization of the in-scattering term and contains the contributions of other directions to the current direction l. The term S_i^B contains wall emission terms contained in Eq. (11).

Discretization of Energy Equation

The energy equation is also integrated over the control volume i to yield discrete equation of the form

$$\sum_{j=1}^{N} \mu_{ij}^T T_j - \beta_i^T T_i = \alpha_i^T \sum_{k=1}^{L} I_i^k \omega_k - S_i^T + S_i^h \tag{16}$$

where

$$\alpha_i^T = \kappa \Delta\mathcal{V}_i, \qquad \beta_i^T = 16\kappa\sigma T_i^{*3}\Delta\mathcal{V}_i, \qquad S_i^T = 12\kappa\sigma T_i^{*4}\Delta\mathcal{V}_i$$

The coefficients μ_{ij}^T and the source term S_i^h arise from discretization of the convection and diffusion terms as well as nonradiative source terms. The terms $\beta_i^T T_i$ and S_i^T result from linearizing the T^4 dependence of the I^B term in Eq. (5). The term containing α_i^T results from discretizing the $I(\boldsymbol{s})$ integral in Eq. (5).

Solution

Equations (15) and (16) constitute a set of coupled, nominally linear equations in $N(L+1)$ unknowns, viz., the L intensities I_i^l and temperature T_i at the N cells. In principle, this set could be solved simultaneously using any linear equation solver. However, the large amount of storage required for the coefficient matrix makes it impractical to use this approach for any reasonable grid size. For example, using sparse matrix storage, a typical hexahedral mesh with six neighbors per cell would require $7N(L+1)^2$ words. Even for a relatively coarse angular discretization of $N_\theta = N_\phi = 2$ this represents approximately a hundredfold increase in memory for a typical flow and heat transfer code. Therefore, in practice the equations are always solved in a sequential manner.

Sequential Solution Procedure

The sequential approach is the most commonly used approach for solving the discrete equations (see, for example, Murthy and Mathur[4]). In this approach, only the spatial coupling of intensity in any given direction is treated implicitly. The influence of all other directions as well as temperature is included explicitly via source terms, evaluated using the latest available values. This leads to a linear system of the form

$$\boldsymbol{M}^l \boldsymbol{x}^l + \boldsymbol{b}^l = 0 \tag{17}$$

where $\boldsymbol{x}^l$ is a vector of intensities at all cells in the direction l and $\boldsymbol{M}^l$ is a $N \times N$ matrix with coefficients:

$$M_{ij}^l = \begin{cases} \mu_{ij}^l, & i \neq j \\ \mu_{ii}^l + \alpha_i^l, & i = j \end{cases} \tag{18}$$

Because μ_{ij} are nonzero only for neighboring cells, $\boldsymbol{M}$ is a very sparse matrix requiring, for example, $7N$ words of storage for a hexahedral mesh. The discretization procedure described earlier ensures that the matrix is diagonally dominant, and therefore, the system can be solved using iterative methods. Because the different directions are decoupled from each other (and from the temperature) in the M_{ij}^l matrix, the overall procedure reduces to visiting each direction (and the temperature equation) sequentially, solving for the variable over the entire domain. An algebraic multigrid method is used for the solution of the discrete equations for each variable, as described next. The explicit terms, evaluated using prevailing values, are updated in an outer iteration.

Algebraic Multigrid Method

In the algebraic multigrid (AMG) method, a hierarchy of coarse equation sets is constructed by grouping a number of fine-level discrete equations. Residuals from a fine-level relaxation sweep are restricted to form the source terms for the coarser level correction equations. The solution from the coarser equations in turn is prolonged to provide corrections at the finer level. This effective use of different grid sizes permits the reduction of errors at all wavelengths using relatively simple smoothing operators. A detailed description may be found in Ref. 9.

The agglomeration procedure used in the present work is to visit each ungrouped fine-level cell and group it with n of its adjacent neighboring ungrouped cells for which the coefficient M_{ij} is the largest.[12] Best performance is usually obtained for $n = 2$, and this value was used for all calculations presented here. The coefficients for the coarse-level matrix $\bar{\boldsymbol{M}}$ are obtained by summing up coefficients of the fine-level equations:

$$\bar{M}_{IJ} = \sum_{i \in G_I} \sum_{j \in G_J} M_{ij} \tag{19}$$

where G_I is the set of fine-level cells that belong to the coarse-level group I. This results in a system of equations of the same form as the fine level, i.e., Eq. (17), with $(1/n)$th the number of unknowns:

$$\bar{\boldsymbol{M}}\bar{\boldsymbol{x}} + \bar{\boldsymbol{b}} = 0 \tag{20}$$

where the source vector at the coarse level has components given by

$$\bar{b}_I = -\sum_{i \in G_I} r_i \tag{21}$$

and r_i is the residual in the fine-level equation at the current iteration

$$r_i = M_{ij} x_j^{*l} + b_i \tag{22}$$

The process is repeated recursively until no further coarsening is possible. We use a V cycle to visit the grid levels. The solution at any level is obtained by a Gauss–Seidel iterative scheme and is used

to correct the current iterate at the next finer level. Thus, for all $i \in G_I$,

$$x_i = x_i^* + \bar{x}_I \tag{23}$$

To summarize, one global iteration of the sequential solution procedure consists of the following steps:

1) Update all outgoing direction intensities at boundary faces using Eq. (11) using the prevailing values (or initial guesses) of I^l and T.

2) For each direction l, construct the linear system [Eq. (17)] using the latest values for temperature and intensities in other directions and solve the system to a specified tolerance using several cycles of the algebraic multigrid procedure.

3) Similarly construct and solve the linear system [Eq. (16)] for temperature.

4) Repeat until convergence.

The additive correction multigrid scheme is often used as a linear solver for unstructured mesh methods[13] because coarse-level discrete equations can be constructed by algebraic manipulation of fine-level equations, rather than by direct discretization on the coarse mesh (a procedure we shall refer to as *geometric* multigrid). For unstructured meshes, coarse-mesh cells are typically arbitrary, multifaceted nonconvex polyhedra. Direct discretization may lead to a loss of diagonal dominance in the evaluation of first derivatives required for diffusive fluxes.

COMET

Several iterations are required in the sequential solution procedure because of both the lagged treatment of intensities in other directions as well as the nonlinear terms in the energy equation. The rate of convergence of the sequential approach decreases as the magnitude of the explicitly included source terms increases. This can occur in many situations. In case of large absorption coefficients, the emission terms in the intensity equations [Eq. (15)] and the absorption and emission terms in the energy equation dominate the transport terms. Similarly, a large scattering coefficient couples the intensity in different directions at a point more strongly than to the intensities in spatial neighbors. Coupling between different directions is also introduced due to reflection and refraction at boundaries, as well as due to geometric factors such as axisymmetry and periodicity.

As discussed earlier, a coupled solution of all of the equations at all of the cells is virtually impossible. Instead, we update all of the intensities and temperature at a cell together. To this end, we write Eqs. (15) and (16) at a cell i as

$$\boldsymbol{P}_i \boldsymbol{q}_i + \boldsymbol{r}_i = 0 \tag{24}$$

where the solution vector $\boldsymbol{q}_i$ is given by

$$\boldsymbol{q}_i = \begin{bmatrix} I_i^1 \\ I_i^2 \\ \vdots \\ I_i^L \\ T_i \end{bmatrix}$$

$\boldsymbol{P}_i$ is an $(L+1) \times (L+1)$ matrix given by

$$\begin{bmatrix} M_{ii}^1 + \eta_i^{11} + \gamma_i^1 \omega_1 & \eta_i^{12} + \gamma_i^1 \omega_2 & \cdots & \beta_i^1 \\ \eta_i^{21} + \gamma_i^2 \omega_1 & M_{ii}^2 + \eta_i^{22} + \gamma_i^2 \omega_2 & \cdots & \beta_i^2 \\ \vdots & & & \\ -\alpha_i^T \omega_1 & -\alpha_i^T \omega_2 & \cdots & M_{ii}^T \end{bmatrix}$$

The vector $\boldsymbol{r}_i$ is given by

$$\boldsymbol{r}_i = \begin{bmatrix} \sum_{j=1, i \neq j}^{N} \mu_{ij}^l I_j^l - S_i^1 - S_i^B \\ \vdots \\ \sum_{j=1, i \neq j}^{N} \mu_{ij}^T T_j + S_i^T - S_i^h \end{bmatrix}$$

Optimization of Point Inversion

The matrix $\boldsymbol{P}_i$ is fully populated and is not always diagonally dominant. In general, therefore, Eq. (24) is best solved by direct methods. However, in many situations, the matrix $\boldsymbol{P}_i$ has a specific structure that can be exploited to optimize the direct solution and to reduce its cost.

Consider radiation in an absorbing, emitting, and isotropically scattering medium. For cells that do not involve boundary faces, the coefficients η_i^{lk} are zero. We introduce the net irradiation I^a defined as

$$I_i^a = \sum_{k=1}^{L} I_i^k \omega_k \tag{25}$$

We can write Eq. (24) as

$$\tilde{\boldsymbol{P}}_i \tilde{\boldsymbol{q}}_i + \tilde{\boldsymbol{r}}_i = 0 \tag{26}$$

where $\tilde{\boldsymbol{q}}$ is the augmented vector

$$\boldsymbol{q}_i = \begin{bmatrix} I_i^1 \\ I_i^2 \\ \vdots \\ I_i^L \\ T_i \\ I_i^a \end{bmatrix}$$

and $\tilde{\boldsymbol{P}}_i$ is an $(L+2) \times (L+2)$ matrix given by

$$\begin{bmatrix} M_{ii}^1 & 0 & \cdots & 0 & \beta_i^1 & \gamma_i^1 \\ 0 & M_{ii}^2 & \cdots & 0 & \beta_i^2 & \gamma_i^2 \\ \vdots & & & & & \\ 0 & 0 & \cdots & M_{ii}^L & \beta_i^L & \gamma_i^L \\ 0 & 0 & \cdots & 0 & M_{ii}^T & -\alpha_i^T \\ \omega_1 & \omega_2 & \cdots & \omega_L & 0 & -1 \end{bmatrix}$$

Matrix $\tilde{\boldsymbol{P}}_i$ is easily upper triangularized. The solution can be obtained by back substitution in $\mathcal{O}(N)$ operations instead of the $\mathcal{O}(N^2)$ operations typically required for a general matrix. Because most of the cells in any grid do not have boundary faces, this represents a considerable reduction in overall computational time.

Similar optimized solution strategies can also be devised for other situations such as linear anisotropic scattering or axisymmetry where the equations at interior cells are coupled but the matrix has a regular structure. For our purposes, we use the special linear solution procedure described earlier for interior cells but revert to using the general lower-upper decomposition method for cells that have faces on reflective boundaries.

One global sweep of the point-coupled solution consists of visiting each cell in the domain and updating its intensities and temperature using the latest available values of neighboring cell intensities and temperatures. Several such sweeps are necessary for convergence. Although this procedure efficiently handles the coupling of the RTE and energy equation at a point, it is slow in transmitting boundary condition information to the interior. As in a Gauss-Seidel scheme, the computational stencil is local; consequently, the method can reduce short wavelength errors efficiently, but not long wavelength errors. To accelerate overall convergence, it is desirable to incorporate a multigrid procedure. This is discussed next.

Coupled Multigrid Procedure

The principle behind the coupled multigrid procedure used in the present work is similar to the AMG method described earlier for the solution of a linear system. However, in the point-coupling context, the AMG procedure requires the coefficients of all equations to be available simultaneously for agglomeration. It is not feasible to store the entire coefficient matrix for all cells because the number of discrete directions is large. Consequently, coarse-level matrices

cannot be constructed algebraically in the manner of Eq. (19). Instead, we use a combination of geometric and AMG techniques, as described next.

For the energy equation, an AMG idea is used to derive coarse-level coefficients. We compute and store the coefficient matrix $\boldsymbol{M}^T$ at the beginning of every iteration. The coarse levels are then constructed from this matrix by summing up the coefficients of the fine levels, as in the sequential solution procedure. This is necessary because the agglomeration of unstructured fine-level meshes leads to nonconvex polyhedra at coarse levels. If diffusion terms are evaluated directly on these polyhedra, using the procedures outlined by Mathur and Murthy,[13] a loss of diagonal dominance is possible. An algebraic summation in the manner of Eq. (19) avoids this problem. Also, agglomeration based on the coefficients of the energy equation permits the agglomeration to change every iteration to capture optimally the direction of information transfer.

Unlike the energy equation, the discretization of the RTE at coarse levels presents no special problems since no diffusion terms are involved; the procedures outlined in previous sections can be used without modification on arbitrary polyhedra even if they are not convex. Consequently, coarse-level coefficients for the RTE are computed directly on the coarse level mesh. Because these computations are relatively inexpensive, coefficients are recomputed for each coarse-level cell when it is visited.

As an illustration, consider the coarse-level cell I shown in Fig. 2 and its coarse-level neighbor J. The dashed lines show the boundaries of the finest level grid. The coefficient $\bar{M}^T_{IJ}$ relating cells I and J is computed algebraically from Eq. (19), but the coefficients $\bar{M}^l_{IJ}$ are obtained by discretization, i.e., by using Eq. (10) over faces c–d and d–e. Further optimization is possible by agglomerating the faces at the coarse levels. Thus, for the cell I in Fig. 2, $\bar{M}^l_{IJ}$ can be obtained by using the face c–e rather than the faces c–d and d–e separately. This results in considerable savings in computational time at coarse levels, where the number of faces bordering a cell can be very large.

The coupled multigrid procedure requires additional memory over the sequential procedure for the storage of intensity corrections and residuals for all directions at coarse levels. For a coarsening ratio n of 2, this represents approximately an increase of $2L$ words per cell. The matrices $\boldsymbol{M}^T$ and $\bar{\boldsymbol{M}}^T$ are required in both the sequential and point-coupled procedures and do not imply any extra storage in the coupled scheme.

One iteration of the coupled multigrid solution procedure consists of the following steps:

1) Discretize the energy equation to obtain the linear system Eq. (17).

2) Create the hierarchy of coarse levels based on $\boldsymbol{M}^T$ and compute the coarse-level matrices $\bar{\boldsymbol{M}}^T$.

3) Perform one V cycle consisting of the following:

a) Perform ν_1 relaxation sweeps at the finest level. This requires visiting each cell and constructing and solving Eq. (24) or (26).

b) Compute the residuals at the finest level. This also requires visiting each cell and constructing Eq. (24) or (26).

c) Inject the residuals to the next coarse level in the manner of Eq. (21).

d) Recursively repeat steps 3a–3c for each coarse level.

e) Using the solution at the coarsest level, correct the solution at the next finer level in the manner of Eq. (23). Perform ν_2 relaxation sweeps.

f) Repeat step 3e recursively for each finer level.

4) Repeat step 3 until the desired reduction of residuals is obtained.

With the coupled multigrid procedure, multiple iterations are required only because of nonlinearities in the energy equation. Because all of the coupling between the intensities is handled implicitly, solutions for problems that do not involve coupling with the energy equation can be obtained in one global iteration.

For problems involving radiative equilibrium, it is possible to choose T^4 as the temperature variable, rather than the temperature T. For given heat flux and given temperature boundary conditions, a point-coupled procedure would yield convergence in one iteration. Most practical problems, however, require the solution of the energy equation because transport due to conduction and convection plays an important role. Consequently, we use T as the dependent variable.

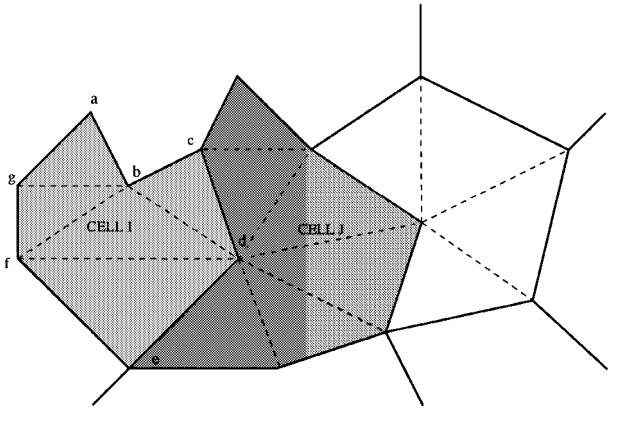

Fig. 2 Coarse-level grid.

Results

We evaluate the performance of COMET for problems involving radiation and compare it with that of sequential solution procedure. Because the issues of accuracy and spatial and angular discretization, as well control-angle overhang, have been addressed elsewhere,[4] we focus our attention purely on the convergence characteristics of the two approaches. Because the underlying discretization procedure is the same in both cases and only the solution path is different, the converged results from the two approaches are nearly identical, as expected; individual intensities are within 0.01% of each other. In the results that follow, the solution is judged to have converged when the scaled residuals for all of the equations have reduced to 10^{-6}. At any global iteration, each of the linear equation set is solved to a tolerance of 0.1 reduction in the residual. All of the timings reported are obtained on a Sun Ultra Sparcstation 1.

Radiation in Quadrilateral Enclosure

The objective of this problem is to compare the two approaches for participating radiation in an absorbing and emitting medium. In this case, the intensity equations and the energy equation are mutually coupled. We use a quadrilateral cavity, shown in Fig. 3. All walls are black and have specified temperatures. The medium has a Planck number ($Pl = k/4L\sigma T_c^3$) of 1.4×10^{-5}, where k is the thermal conductivity, L is the length of the bottom wall, and T_c is its temperature. There is no scattering. The optical thickness κL is varied from 0 to $100L$.

A baseline unstructured grid of 640 triangles, shown in Fig. 4, is used; the problem is repeated on two finer grids of 2600 and

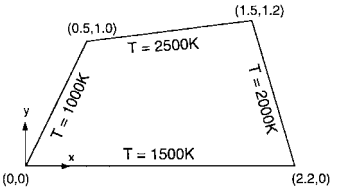

Fig. 3 Radiation in quadrilateral enclosure: schematic.

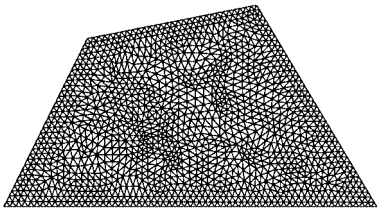

Fig. 4 Radiation in quadrilateral enclosure: baseline mesh.

Table 1 Radiation in a quadrilateral cavity: $Pl = 1.4 \times 10^{-5}$ and $N_\theta = N_\phi = 2$

	Sequential		COMET	
κL	CPU, s	Iterations	CPU, s	Iterations
		650 Cells		
0	2.41	11	3.14	10
$0.1L$	3.14	12	2.65	5
$1.0L$	5.78	25	2.73	5
$10.0L$	45.98	247	3.38	5
$100.0L$	775.61	4127	5.65	5
		2,600 Cells		
0	13.78	10	15.32	13
$0.1L$	21.23	11	13.24	5
$1.0L$	30.31	25	11.89	5
$10.0L$	186.80	262	11.83	5
$100.0L$	3655.93	5952	15.04	5
		10,400 Cells		
0	47.27	9	71.99	15
$0.1L$	86.42	12	62.92	5
$1.0L$	155.81	24	74.66	5
$10.0L$	998.80	271	49.19	5
$100.0L$	$>1e5$	>6000	89.55	5

Table 2 Radiation in a quadrilateral cavity: $Pl = 1.4 \times 10^{-5}$ and $N_\theta = N_\phi = 4$

	Sequential		COMET	
κL	CPU, s	Iterations	CPU, s	Iterations
		650 Cells		
$0.1L$	11.25	12	12.4	5
$1.0L$	18.66	25	12.82	5
$10.0L$	144.06	239	16.45	5
		2,600 Cells		
$0.1L$	98.53	12	81.85	5
$1.0L$	112.46	25	73.23	5
$10.0L$	690.99	252	72.65	5
		10,400 Cells		
$0.1L$	488.57	12	376.99	5
$1.0L$	714.97	24	344.63	5
$10.0L$	3209.37	266	278.73	5

10,400 cells respectively, obtained by isotropic division of the baseline grid. An angular discretization of $N_\theta = N_\phi = 2$ is used. Control-angle overhang occurs in the ϕ direction because some of the walls are not orthogonal, and therefore, a 1×10 pixelation is used at all boundaries. The CPU time and iterations required for solution are shown in Table 1 for the two approaches.

At $\kappa L = 0$, the equations are independent of one another, and therefore, there is no advantage to using COMET. We include the results for this limit only to provide an idea of the overhead associated with recomputing the coefficients at each relaxation sweep. For moderate optical thickness, the coupled approach takes 15–60% less time. The speedup for higher optical thicknesses is dramatic; factors of $\mathcal{O}(10^3)$ are obtained. In this limit, the speedup improves with increased grid size. Similar speedup behavior is observed for a finer angular discretization of $N_\theta = N_\phi = 4$ and at a higher Planck number of 1.4×10^{-3}, results for which are given in Tables 2 and 3, respectively.

Another important feature of the coupled approach is that for all of the participating cases the number of iterations required for convergence remains the same with increasing optical thickness. This is because of the direct solution of the point coupling of intensity and temperature at each cell. The number of iterations remains the same for all grid sizes as well; this indicates the efficacy of the multigrid cycling.

Isotropic Scattering in a Square Cavity

In this problem, the objective is to study the performance for a isotropically scattering medium. The problem geometry consists of a square box of side L with all but the bottom walls at $T = 0$. The bottom wall is black and set at $T = 1000$. The medium is neither absorbing nor emitting so that the radiation problem is decoupled from the thermal problem. All of the intensity equations, however, are coupled to all of the other equations through the scattering term. This problem was also used by Fiveland and Jessee[8] to examine various acceleration schemes for the discrete ordinates method.

Table 3 Radiation in a quadrilateral cavity: $Pl = 1.4 \times 10^{-3}$ and $N_\theta = N_\phi = 2$

	Sequential		COMET	
κL	CPU, s	Iterations	CPU, s	Iterations
		650 Cells		
$0.1L$	3.17	11	3.10	5
$1.0L$	5.14	22	2.61	5
$10.0L$	46.28	244	3.44	5
		2,600 Cells		
$0.1L$	21.27	11	17.02	7
$1.0L$	27.20	22	15.52	7
$10.0L$	186.12	254	11.83	5
		10,400 Cells		
$0.1L$	131.42	11	83.40	8
$1.0L$	119.92	19	78.34	8
$10.0L$	947.28	271	66.40	7

Table 4 Scattering in a square cavity

	Sequential		COMET	
$\sigma_s L$	CPU, s	Iterations	CPU, s	Iterations
		10 × 10 Cells		
0.1	1.28	18	0.53	5
1.0	2.46	35	0.57	5
10.0	20.44	297	0.93	5
100.0	239.77	3385	1.15	5
		20 × 20 Cells		
0.1	3.03	15	1.68	5
1.0	5.65	35	2.05	5
10.0	45.77	346	5.66	5
100.0	679.07	5080	10.04	5
		40 × 40 Cells		
0.1	13.46	15	6.45	5
1.0	22.46	32	10.01	5
10.0	159.71	371	35.13	5
100.0	>2500	>6000	93.81	5
		80 × 80 Cells		
0.1	58.28	14	24.31	5
1.0	100.90	29	41.44	5
10.0	659.78	371	157.13	5
100.0	$>1e5$	>6000	593.63	5

We use a sequence of structured quadrilateral grids ranging from 10×10 to 80×80 cells. The angular discretization is $N_\theta = N_\phi = 2$; no pixelation is necessary because of the orthogonal geometry. The problem is run for $\sigma_s L$ ranging from 0.1 to 100. CPU times and number of iterations for convergence are presented in Table 4.

The coupled approach results in significant speedup for the entire range of scattering coefficients studied. As in the absorption case, it increases dramatically with increased $\sigma_s L$ as well as with grid size. The number of iterations again remain constant over the entire range. These observations are in contrast to the results obtained by Fiveland and Jesse[8] using the mesh-rebalance method. They observed a strong correlation between the mesh-based optical thickness and the efficacy of the acceleration scheme and had to construct coarse-level grids so as to obtain local optical thickness of $\mathcal{O}(1)$. The current approach does not show such a dependence and can be used with arbitrary grid sizes.

Radiation Between Parallel Plates

To evaluate the effectiveness of COMET for cases where coupling between directions is caused by geometric factors, we consider the problem of radiation between two parallel plates. The geometry is shown in Fig. 5 and consists of two diffuse black walls of length L separated by a distance d. The lateral boundaries are specularly reflecting mirrors. Although the problem has a trivial one-dimensional

Table 5 Radiation between parallel plates

	Sequential		COMET	
Grid size	CPU, s	Iterations	CPU, s	Iterations
		$N_\theta = N_\phi = 2$		
10×10	3.36	62	0.54	4
20×20	5.85	49	0.75	4
40×40	16.98	41	0.96	2
80×80	60.26	36	3.47	2
		$N_\theta = N_\phi = 4$		
10×10	17.49	114	3.96	4
20×20	36.20	91	5.81	4
40×40	114.20	78	8.31	2
80×80	401.26	69	18.28	2

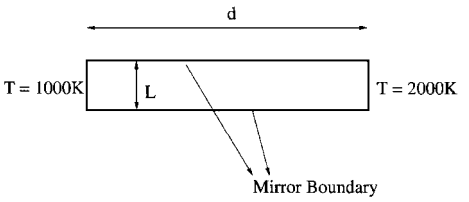

Fig. 5 Radiation between parallel plates: schematic.

solution that is independent of the aspect ratio L/d, the convergence behavior of the numerical procedure does depend on the aspect ratio as well as the angular discretization. As the aspect ratio is decreased, the number of cells that have coefficients connecting them directly to the Dirichlet conditions at the two walls decreases. Most of the cells in the middle of the domain get information only from rays that originate from the mirror boundaries. A sequential solution procedure, thus, requires several iterations to propagate the boundary intensities to all of the interior cells. The coupled procedure is expected to improve this because the interdirectional coupling of intensities is treated implicitly. Note that, for this simple configuration, both the sequential and coupled solution procedures will produce the exact heat flux solution on the two walls for any angular or spatial discretization.

We compute the solution for aspect ratio of 0.1. Because the preceding test cases have already examined the performance for cases involving absorption, emission, and scattering and the interest here is to isolate the geometric effects, the medium is set to be nonparticipating. Calculations are performed for three sets of grids and two angular discretizations. Performance figures are given in Table 5. The coupled procedure is found to converge in two to four iterations for all of the cases. As expected, the sequential procedure requires a large number of iterations; the number increases with the angular discretization. Speedup factors of 10–20 in overall solution time are obtained with the coupled procedure. One interesting feature of the results in Table 5 is that the number of iterations is found to decrease with increasing grid sizes. (The overall solution time is, of course, higher because of the larger number of cells involved.) This is probably a result of the larger number of global sweeps that occur as the grid size increases because there are greater number of coarse levels.

Conclusions

COMET has been devised to accelerate the convergence of the finite volume scheme for radiative heat transfer. Care has been taken to formulate a scheme usable with arbitrary unstructured polyhedra both at the fine as well as the coarse levels. Substantial acceleration, of over a factor of 100, has been obtained for optical thicknesses greater than 10. Acceleration has been obtained even for relatively low optical thicknesses. There is a performance penalty for completely transparent media; however, because the total CPU time for these cases is small in any case, the penalty should not matter in most practical problems. Consequently, COMET may be used as a general purpose solver for the entire range of optical thicknesses.

Though not explored in this paper, the procedure may find use in a variety of radiation problems involving interdirectional coupling. These include semitransparent media where, because of refraction, energy is transferred from one direction to another. Sequential procedures treat this transfer explicitly, resulting in slow convergence. Similarly, in axisymmetric geometries, interdirectional coupling occurs as a result of angular redistribution terms. Similar explicit coupling occurs at rotationally periodic boundaries when sequential solvers are used. The application of this procedure to these situations may be worth exploring.

Acknowledgments

We wish to acknowledge the use of the Fluent, Inc., solver FLUENT/UNS, and its mesh generators PreBFC and TGrid, in this work.

References

[1]Raithby, G. D., and Chui, E. H., "A Finite-Volume Method for Predicting Radiant Heat Transfer in Enclosures with Participating Media," *Journal of Heat Transfer*, Vol. 112, May 1990, pp. 415–423.

[2]Chui, E. H., and Raithby, G. D., "Computation of Radiant Heat Transfer on a Non-Orthogonal Mesh Using the Finite-Volume Method," *Numerical Heat Transfer*, Vol. 23, No. 3, 1993, pp. 269–288.

[3]Chai, J., Parthasarathy, G., Patankar, S., and Lee, H., "A Finite-Volume Radiation Heat Transfer Procedure for Irregular Geometries," AIAA Paper 94-2095, June 1994.

[4]Murthy, J., and Mathur, S., "Finite Volume Method for Radiative Heat Trnasfer Using Unstructured Meshes," *Journal of Thermophysics and Heat Transfer*, Vol. 12, No. 3, 1998, pp. 313–321.

[5]Fiveland, W. A., and Jessee, J. P., "Finite Element Formulation of the Discrete Ordinates Method for Multidimensional Geometries," *Journal of Thermophysics and Heat Transfer*, Vol. 8, No. 3, 1994, pp. 426–433.

[6]Lewis, E., and Miller, W., *Computational Methods for Neutron Transport*, New York, Wiley, 1984.

[7]Chui, E. H., and Raithby, G. D., "Implicit Solution Scheme to Improve Convergence Rate in Radiative Transfer Problems," *Numerical Heat Transfer*, Vol. 22, No. 3, 1992, pp. 251–272.

[8]Fiveland, W., and Jessee, J., "Acceleration Schemes for the Discrete Ordinates Method," *1995 National Heat Transfer Conference–Volume 13*, edited by Y. Bayazitoglu, D. Kaminski, and P. D. Jones, HTD-Vol. 315, American Society of Mechanical Engineers, New York, 1995, pp. 11–19.

[9]Hutchinson, B. R., and Raithby, G. D., "A Multigrid Method Based on the Additive Correction Strategy," *Numerical Heat Transfer*, Vol. 9, 1986, pp. 511–537.

[10]Modest, M. F., *Radiative Heat Transfer*, Series in Mechanical Engineering, McGraw–Hill, New York, 1993, pp. 295–322.

[11]Patankar, S. V., *Numerical Heat Transfer and Fluid Flow*, McGraw–Hill New York, 1980, pp. 48, 49.

[12]Lonsdale, R. D., "An Algebraic Multigrid Scheme for Solving the Navier–Stokes Equations on Unstructured Meshes," *Proceedings of the 7th International Conference on Numerical Methods on Laminar and Turbulent Flow*, edited by C. Taylor, J. H. Chen, and G. Homsy, Pineridge, Swansea, Wales, UK, 1991, pp. 1432–1442.

[13]Mathur, S., and Murthy, J., "A Pressure Based Method for Unstructured Meshes," *Numerical Heat Transfer*, Vol. 31, No. 2, 1997, pp. 195–216.