Received April 1998
Revised December 1998
Accepted January 1999

Parallelization of the finite volume method for radiation heat transfer

P.J. Coelho and J. Gonçalves
Instituto Superior Técnico, Technical University of Lisbon, Mechanical Engineering Department, Lisbon, Portugal

Keywords *Finite volume, Heat transfer, Parallel computing*

Abstract *The finite volume method for radiative heat transfer calculations has been parallelized using two strategies, the angular domain decomposition and the spatial domain decomposition. In the first case each processor performs the calculations for the whole domain and for a subset of control angles, while in the second case each processor deals with all the control angles but only treats a spatial subdomain. The method is applied to three-dimensional rectangular enclosures containing a grey emitting-absorbing medium. The results obtained show that the number of iterations required to achieve convergence is independent of the number of processors in the angular decomposition strategy, but increases with the number of processors in the domain decomposition method. As a consequence, higher parallel efficiencies are obtained in the first case. The influence of the angular discretization, grid size and absorption coefficient of the medium on the parallel performance is also investigated.*

Nomenclature

A_x, A_y, A_z	= Area of the cell faces of a control volume normal to the x, y and z directions, respectively
$D_{i,cx}$, $D_{i,cy}$, $D_{i,cz}$	= Quantities defined in equation (3)
G	= Incident radiation
I	= Radiation intensity
I_{x_e}, I_{y_e} I_{z_e}	= Outgoing radiation intensities at cell faces x=constant, y=constant and z=constant, respectively
I_{x_i}, I_{y_i}, I_{z_i}	= Incoming radiation intensities at cell faces x=constant, y=constant and z=constant, respectively
$\vec{i}$ $\vec{j}$ $\vec{k}$	= Unit vectors along x, y and z directions, respectively
$\vec{n}$	= Normal unit vector
n_{iter}	= Number of iterations required to achieve convergence
N_x, N_y, N_z	= Number of grid nodes along x, y and z directions, respectively
N_θ, N_φ	= Number of discrete polar and azimuthal angles per octant, respectively
p	= Number of processors
p_x, p_y, p_z	= Number of processors along x, y and z directions, respectively
q	= Heat flux
s	= Distance travelled by a beam
$\vec{s}_i$	= Unit vector along the direction i travelled by a beam
S	= Speedup
t	= Time
T	= Temperature
V	= Volume
γ	= Weighting coefficient
$\Delta\Omega$	= Solid angle
ε	= Emissivity; efficiency
θ	= Polar angle
κ	= Absorption coefficient
σ	= Stefan-Boltzmann constant
φ	= Azimuthal angle
Subscripts	
b	= Blackbody
c	= Communication
cp	= Computation
e	= Execution
i	= Direction

This work has been financially supported by the Commission of the European Communities under the ESPRIT project 8114, HP-PIPES – "High performance parallel computing for process engineering simulation". The authors are indebted to the Computing Centre of the University of Minho, Portugal, for the permission to use their parallel computers.

International Journal of Numerical Methods for Heat & Fluid Flow, Vol. 9 No. 4, 1999, pp. 388-404.

iter	= Iteration	θ	= Polar angle
p	= Number of processors	φ	= Azimuthal angle
P	= Grid node	*Superscripts*	
w	= Wall	$\Rightarrow$	= Vector

Introduction

Many different models are presently available to perform radiative heat transfer calculations. However, radiation is often coupled to fluid flow problems and this restricts the choice of the radiation model. One of the most popular methods for solving fluid flow problems is the finite volume method. A finite volume method has also been developed for predicting radiant heat transfer[1] and aimed at a straightforward coupling with fluid flow solvers in complex domains.

The computational requirements of the radiative heat transfer calculations may be quite high if, for example, the radiative properties of the medium need to be determined, spectral effects are important or complex geometries are involved. These requirements greatly increase if a fluid flow problem is also present. Parallel computers may significantly reduce the computing time. Although a lot of work has been done on the parallelization of fluid flow solvers, only little attention has been given to the parallelization of radiation models[2], despite the impact that parallel computing may have[3]. A brief review of the work done on the parallelization of radiation models has been presented in [4].

In order to exploit high-performance computing capabilities the present paper addresses the parallelization of the finite volume method. This follows previous work of the authors on the parallelization of the discrete transfer[5] and discrete ordinates[4,6] methods. The coupling with computational fluid dynamics codes is not investigated here. However, keeping in mind that most fluid flow solvers are parallelized using the domain decomposition method, two different parallelization techniques for the finite volume radiation method are implemented are compared. In the angular decomposition parallelization (ADP) method each processor performs the calculations for the whole domain but only treats a certain number of solid angles which results from splitting all the solid angles among the processors. In the spatial domain decomposition parallelization (DDP) method the spatial domain is split into subdomains and each processor performs the calculations for one subdomain and for all the solid angles.

A short description of the finite volume method is given in the next section, which is followed by the description of the parallel implementation. Then, the results obtained are presented and discussed, and the paper ends with a summary of the main conclusions.

The finite volume model

A short description of the finite volume method is presented here to facilitate the explanation of the parallelization methods. A complete description of the method may be found elsewhere[1,7]. The finite volume method is a numerical

technique which is applied in this work to the solution of the radiative transfer equation. This equation may be written as follows for the direction s in the case of emitting-absorbing and non-scattering grey media:

$$\frac{dI}{ds} = -\kappa I + \kappa I_b \tag{1}$$

In the finite volume method the spatial domain is divided into control volumes and the angular space 4π at any spatial location is divided into discrete non-overlapping solid angles, also known as control angles. The radiation intensity is assumed constant inside a control volume and a solid angle. A typical control volume is depicted in Figure 1. Although the method is applied here to Cartesian geometries, it may be applied to arbitrary geometries using curvilinear coordinates[1]. Equation (1) is integrated over a control volume and a control angle, and the divergence theorem is applied to the term on the left hand side of the equation. The variation of the radiation direction in a control angle is accounted for. The application of the divergence theorem to the left hand side of equation (1) yields six terms which represent the inflow and outflow of radiative energy across the six control volume faces. The net flow of radiative energy is equal to the difference between the emitted and the absorbed radiative energy, i.e. the right hand side of equation (1). The following relationship between the volume average intensity, I_{P_i}, and the radiation intensities entering (subscript i) and leaving (subscript e) a control volume is obtained[7]:

$$I_{P_i} \frac{\kappa I_b V \gamma \Delta\Omega_i + D_{i,cx} A_x I_{x_i,i} + D_{i,cy} A_y I_{y_i,i} + D_{i,cz} A_z I_{z_i i}}{\kappa V \gamma \Delta\Omega_i + D_{i,cx} A_x + D_{i,cy} A_y + D_{i,cz} A_z} \tag{2}$$

where

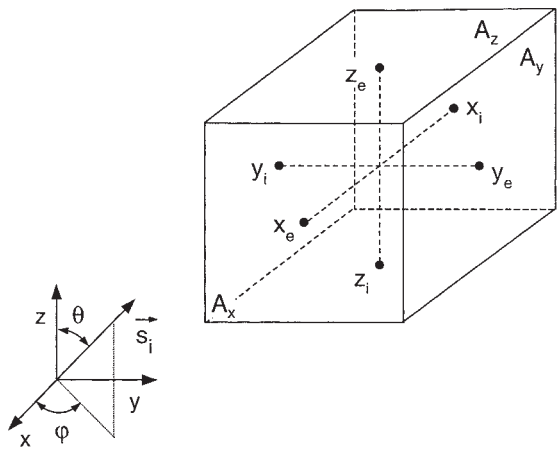

Figure 1.
Typical control volume

$$D_{i,cx} = \int_{\Delta\Omega_i} \vec{s}_i \cdot \vec{i} d\Omega_i \quad , \quad D_{i,cy} = \int_{\Delta\Omega_i} \vec{s}_i \cdot \vec{j} d\Omega_i \quad , \quad D_{i,cz} = \int_{\Delta\Omega_i} \vec{s}_i \cdot \vec{k} d\Omega_i \tag{3}$$

and

$$\Delta\Omega_i = \int_{\Delta\varphi_i} \int_{\Delta\theta_i} \sin\theta d\theta d\varphi \tag{4}$$

The parameter γ relates the incoming and the outgoing radiation intensities to the volume average intensity, and it was set to one in the present work. This means that the step discretization scheme was used[8]. The boundary conditions at x=constant surfaces may be written as

$$I_i = \varepsilon_w I_{bw} + (i - \varepsilon_w) \frac{q_x^-}{\pi} \qquad \text{at} \qquad x = x_{min} \tag{5a}$$

$$I_i = \varepsilon_w I_{bw} + (1 - \varepsilon_w) \frac{q_x^+}{\pi} \qquad \text{at} \qquad x = x_{max} \tag{5b}$$

and the incident heat fluxes are determined by

$$q_x^- = \sum_{\substack{j \\ (\vec{s}_j \cdot \vec{n} < 0)}} I_j \int_{\Delta\Omega_j} |\vec{s}_j \cdot \vec{i}| d\Omega_j \tag{6a}$$

$$q_x^+ = \sum_{\substack{j \\ (\vec{s}_j \cdot \vec{n} > 0)}} I_j \int_{\Delta\Omega_j} |\vec{s}_j \cdot \vec{i}| d\Omega_j \tag{6b}$$

Similar boundary conditions may be written for the other boundaries.

If the temperature field is not known, it must be determined from the simultaneous solution of the energy conservation and the radiative transfer equations. The energy conservation equation may be written as follows, provided that radiation is the dominant mode of heat transfer, the others being negligible:

$$\nabla . q = \kappa(4\sigma T^4 - G) \tag{7}$$

where $\nabla . q$ is the divergence of the radiative heat flux and G is the incident radiation given by

$$G = \int_0^{4\pi} I d\Omega \tag{8}$$

The sequential solution algorithm starts from a guess of the surface radiosities, and a guess of the medium temperature field if it is not prescribed, and

proceeds iteratively as follows. In each iteration, and for each solid angle, the numerical solution is carried out starting from a control volume at one of the corners of the computational domain. The corner is selected according to the sign of the direction cosines of the directions limited by the solid angle under consideration, in such a way that the radiation intensities entering the control volume at that corner are available from the boundary conditions. At that control volume equation (2) is used to compute I_P, and the step scheme is applied to obtain the radiation intensities leaving the control volume. This scheme allows the discrete set of algebraic equations for the direction under consideration to be solved in a straightforward way. In fact, according to equation (2) the radiation intensity at a control volume is only a function of the radiation intensities at the upstream cell faces. Therefore, no iterative solver is needed for the solution of the set of equations for each direction, although the solution algorithm is iterative due to the need to update the boundary conditions. The solution algorithm proceeds visiting all the control volumes to compute their volume average radiation intensities. For example, if the direction cosines of the direction under consideration are all positive, three nested loops sweeping over the positive x, y and z directions are used to compute the cell average radiation intensities according to equation (2). The process is repeated for all the directions (solid angles). Then, the incident heat fluxes on the boundaries are computed from equation (6) and the radiation intensities leaving the boundaries are updated using the boundary conditions (equation 5). If the temperature field is not known, it is also updated via equation (7), and using the incident radiation intensity obtained from equation (8). The iterative procedure continues until the convergence criterion has been achieved, e.g. until the normalized difference between the incident heat fluxes in two successive iterations decreases below a prescribed tolerance.

Parallelization of the finite volume method

Angular decomposition parallelization (ADP)

In the ADP the solid angles resultant from the discretization of the angular space at any spatial location are split among the processors. Each processor performs the calculations for the whole domain, but deals only with a subset of solid angles. The solution algorithm in each processor may be summarized as follows:

(1) Define the problem data. Set the incident heat fluxes on the boundary and the incident radiation in every control volume to zero, and set the iteration counter to one.

(2) Loop over all the solid angles assigned to a processor, and for each one of these angles loop over all the control volumes; perform the following operations for each control volume:

- Get the incoming radiation intensities at the cell faces from the upstream control volumes or from the boundary conditions, as appropriate.

- Calculate the grid node radiation intensity via equation (2).
- Calculate the outgoing radiation intensities at the cell faces.

(3) Calculate partial values of the heat fluxes on the walls using equation (6) with the summation restricted to the solid angles handled by the processor.

(4) Calculate the total incident heat fluxes on the walls, accounting for all the solid angles. This is accomplished by means of data exchange among the processors according to a binary tree network.

(5) If a volumetric heat source is prescribed, calculate a partial value of the incident heat flux in each control volume using equation (8) with the summation restricted to the solid angles treated by the processor. Then, compute the total values, accounting for all the angles, using a binary tree for communication purposes, and update the temperature field.

(6) Apply the boundary conditions to update the radiation intensities leaving the boundaries.

(7) Check if the convergence criteria are satisfied. If not, increase the iteration counter by one and return to step (2).

Communications among the processors are involved in steps (4) and (5) of the solution algorithm outlined above. In both cases the communications are global, i.e., each processor needs to share information with all the others. In particular, a total value, e.g., the total incident heat flux on the boundaries needs to be computed by adding the partial contribution available at every processor, and then this total value must be broadcast such that all the processors know it. Briefly, in the case of 2^n processors, n steps are needed to obtain the total value in one processor. In the first step the even processors send their data to the neighbouring odd processors. So, the relevant data are now spread among only 2^{n-1} processors. The remaining steps are similar. After n steps, one of the processors is able to compute the total value. A reverse procedure consisting again of n steps is performed to enable the broadcast of the total value. This process is easily generalized to the case where the number of processors is not a power of two. Further details may be found elsewhere[4].

If the number of discrete solid angles is a multiple integer of the number of processors there is a perfect load balance among the processors. Otherwise, load imbalance problems will occur associated with the difference between the number of solid angles assigned to each processor. Only the former case was considered in the present work.

Spatial domain decomposition parallelization (DDP)

In the DDP the computational domain is divided into subdomains and each subdomain is assigned to a different processor which treats all the solid angles resultant from the discretization of the angular space. Although the subdomains do not overlap, there is a buffer of halo points added to their

boundaries, including the virtual boundaries, i.e. the boundaries between neighbouring processors. This buffer is used to simplify the exchange of data, namely the radiation intensities, at the virtual boundaries between neighbouring processors, as explained below.

The solution algorithm in each processor consists of the following steps:

(1) As in ADP.

(2) Loop over all the solid angles, and for each one of them loop over the control volumes assigned to the processor; perform the following operations for each control volume:

- Get the incoming radiation intensities at the cell faces from the upstream control volumes or from the boundary conditions, as appropriate. If an upstream control volume lies in a subdomain assigned to a different processor the required radiation intensity lies in the halo region.
- Next two steps as in ADP.

(3) Calculate the total incident heat fluxes on the walls.

(4) Calculate the incident radiation in the control volumes assigned to the processor. If a radiative heat source is prescribed update the temperature of the medium.

(5) Exchange the radiation intensities along the virtual boundaries between neighbouring processors, as schematically shown in Figure 2.

(6) As in ADP.

(7) As in ADP.

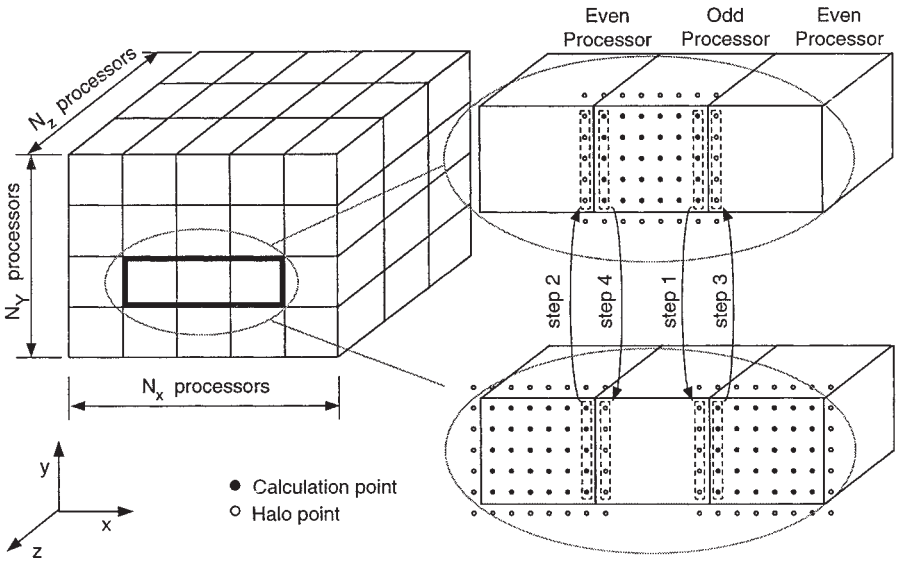

Figure 2.
Schematic of the halo data exchange between neighbouring processors

Local communication between neighbouring processors is involved in step (5) of the solution algorithm. This data exchange is illustrated in Figure 2 for the x direction, and consists of four steps. In step 1 the radiation intensities at the control volumes adjacent to the east boundary of odd processors are copied to the halo region at the west boundary of the neighbouring even processors. In step 2 the radiation intensities at the control volumes adjacent to the east boundary of even processors are copied to the halo region at the west boundary of the neighbouring odd processors. Therefore, the data on the east boundary of all the processors are copied in the first two steps. Steps 3 and 4 are similar and refer to the copy of the data on the west boundary of all the processors to the halo region. Then, the y and z directions are treated similarly. Additional details may be found in [4].

Load balancing problems do not occur provided that the number of grid nodes along each direction is a multiple integer of the number of processors assigned to that direction. This is the case considered in the present work. If this is not the case, load imbalance will occur exactly as in fluid flow problems parallelized using domain decomposition.

Results and discussion

The calculations were performed using a Parsytec MC3-DE with 112 nodes with the transputer T805. The code is written in FORTRAN and the communication among the processors is achieved by means of calls to subroutines of the message passing library of Parsytec. These calls may be easily replaced by calls to subroutines of more widely available software, such as PVM or MPI. In all the studied cases the number of solid angles is a multiple of the number of processors in the ADP, while N_x/p_x, N_y/p_y and N_z/p_z are integers in the DDP. Since the processors are dedicated, the computational load is the same in every processor, and load balancing problems are avoided.

The parallel performance is measured by means of the efficiency ε, and speedup S, which are defined as

$$\varepsilon = \frac{S}{p} \tag{9}$$

$$S = \frac{t_1}{t_p} \tag{10}$$

where t_1 and t_p are the wall clock execution times on one and p processors, respectively. In the problems considered here the influence of the number of processors, angular discretization, grid size and absorption coefficient of the medium on the efficiency and speedup was investigated.

Two test cases are studied below. In the first case the temperature of the medium is prescribed, and so only the radiative transfer equation (1) needs to

be solved. In the second case a volumetric heat source is prescribed instead of the temperature field. Therefore, the energy equation (7) is solved simultaneously with equation (1).

Test case 1

A three-dimensional rectangular furnace ($6\times2\times2$ m^3) of the International Flame Research Foundation was studied in the first test case. The gas temperature field was prescribed according to available measurements[9] and the absorption coefficient of the medium was set equal to $\kappa = 0.2m^{-1}$. The temperature and the emissivity of the boundaries are 320K and 0.86 for the bottom surface, and 1090K and 0.70 for the other surfaces, respectively, as schematically shown in Figure 3a. The standard calculations were performed using a Cartesian uniform grid with $36\times12\times12$ control volumes and an angular discretization with $N_\theta = 5$, $N_\varphi = 2$, where N_θ and N_φ are the number of discrete polar and azimuthal angles per octant, respectively.

This furnace has previously been studied by other researchers using several different models, namely the zone method[9], the discrete ordinates method[10] and the discrete transfer method[11]. Figure 3b shows the predicted incident heat fluxes on the top and bottom surfaces along the section y=1m. These results were computed using a sequential version of the code. It can be seen that the predictions obtained using the finite volume method are in close agreement with the zone method results[9]. In the following parallel calculations it was systematically verified that the numerical solution is independent of the number of processors and coincides with the solution computed using the sequential code. In the results shown below the influence of the number of processors, angular discretization, grid size and absorption coefficient was studied by varying one of these four parameters while the others were kept constant.

The influence of the number of processors is shown in Figure 4. The efficiency decreases with the number of processors, p, as expected. However, this decrease is quite fast for the DDP, yielding $\varepsilon = 53.3$ per cent and $\varepsilon = 12.7$ per cent using four and 72 processors, respectively, and comparatively slow for the ADP, where $\varepsilon = 98.3$ per cent and $\varepsilon = 44.1$ per cent using four and 80 processors, respectively. These evolutions may be explained by the number of iterations required to achieve convergence, n_{iter}, and the ratio of the communication to the execution time, t_c/t_e.

In the ADP n_{iter} is independent of p, while in the DDP there is a marked increase of n_{iter} with the increase of p. In fact, in the DDP the calculations performed in a processor during an iteration require data from the halo points, namely the radiation intensities computed during the previous iteration. Hence, the boundary radiation intensities cannot travel beyond the physical or virtual boundaries of a processor during an iteration. Therefore, several iterations are needed to allow boundary data to spread over the whole domain. On the contrary, both in the sequential algorithm and in ADP the boundary radiation intensities travel through the whole domain in one iteration. If the increase of

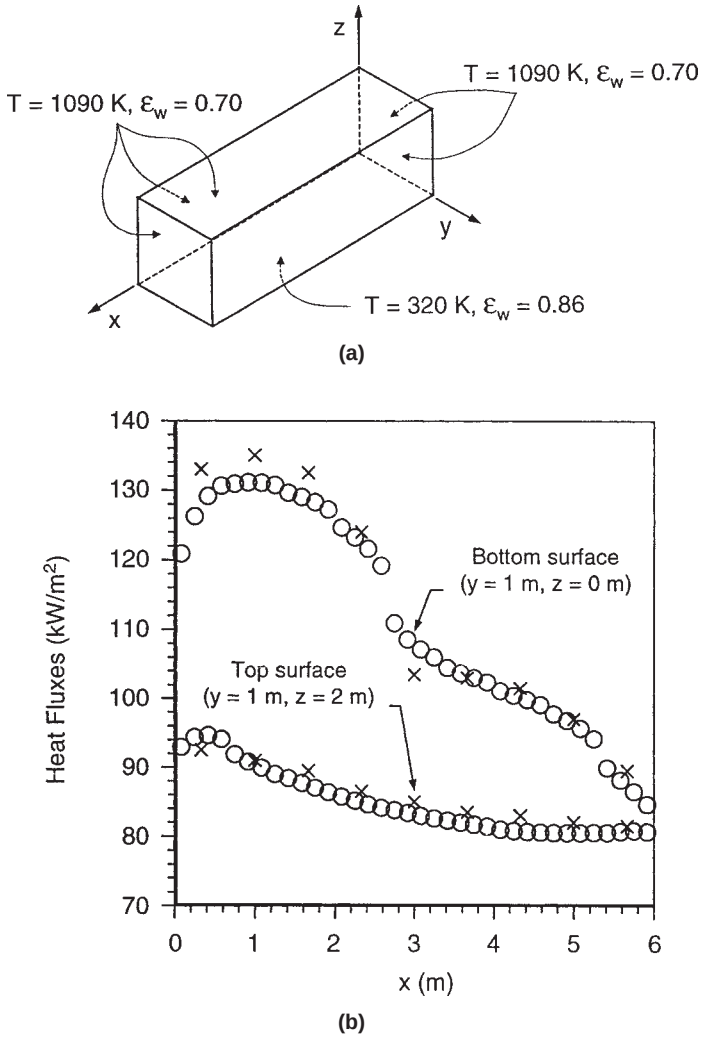

Figure 3.
Rectangular furnace with a prescribed temperature field. (a) Schematic of the furnace. (b) Predicted incident heat fluxes (O: finite volume method; X: zone method[9])

n_{iter}, i.e. the decrease of the convergence rate, were the only reason for the decrease of the efficiency with p, then the efficiency per iteration would be 100 per cent. However, the communication time implies the decrease of the efficiency per iteration with p. This decrease is faster for the ADP as discussed below.

As an example, let us suppose that p is doubled. Then, in the ADP the computation time ($t_{cp} = t_e - t_c$) is reduced to approximately one half because each processor deals only with one half of the solid angles. However, t_c increases because each processor needs to broadcast data to twice as many processors as before. As a consequence, t_c/t_e also increases yielding the observed decrease of the efficiency per iteration. If the DDP is used, t_{cp} is again

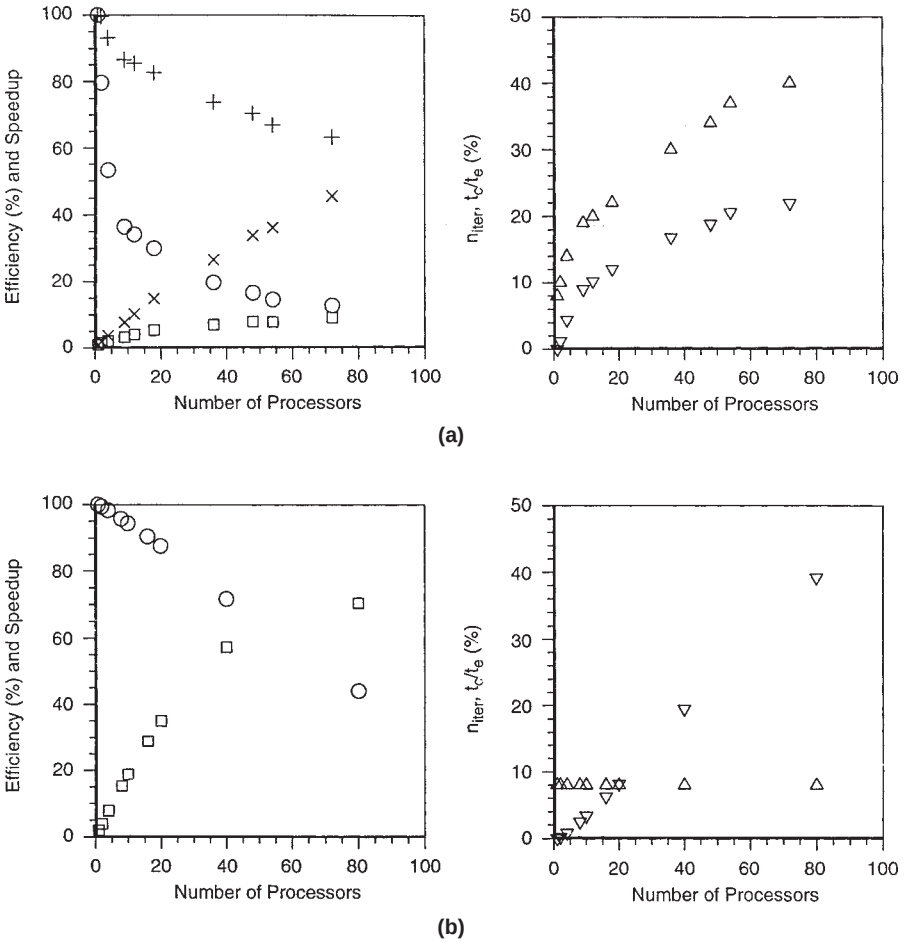

Figure 4.
Influence of the number of processors on the efficiency, speedup, number of iterations and ratio of communication to execution time for test case 1 (O: efficiency; +: efficiency per iteration; □: speedup; X: speedup per iteration; Δ: n_{iter}; ∇: t_c/t_e (%)).
(a) DDP. (b) ADP

reduced to approximately one half because the number of grid nodes assigned to each processor becomes one half of the initial number. However, t_c also decreases because the number of halo points decreases. Nevertheless, the ratio t_c/t_e still increases, as shown in Figure 4, but at a slower rate than in the ADP. In fact, this ratio is approximately proportional to the ratio of the number of halo points to the number of grid nodes assigned to a processor, which increases with p. Consequently, the efficiency per iteration decreases with p.

The significant decrease of the efficiency with p observed in the ADP for p = 40 and p = 80 may be somewhat surprising since the algorithm is highly parallelizable. The reason for that decrease is the small amount of calculations that need to be performed. The execution time is 8.62 and 7.01s for p = 40 and p = 80, respectively. Hence, the communication time becomes a significant part of the execution time, justifying the relatively small efficiencies obtained.

Table I summarizes the results obtained when the angular discretization, the grid size or the absorption coefficient of the medium are varied. If the angular discretization is modified, n_{iter} does not change. Moreover, in the DDP the ratio t_c/t_e remains also approximately constant, because both t_{cp} and t_c increase linearly with the number of solid angles. Therefore, the efficiency and the speedup are independent of the angular discretization. In the ADP t_c is independent of the number of solid angles, but t_{cp} increases with the refinement of the angular discretization. This yields a decrease of t_c/t_e with p, and the consequent increase of the efficiency and speedup.

If the spatial discretization, i.e., the grid size, is modified, n_{iter} does not change regardless of the parallelization method. However, t_c/t_e decreases with the grid size for both methods, and so both the efficiency and the speedup increase. The reason for the decrease of t_c/t_e with the grid size is the following. In the ADP t_{cp} is proportional to the total number of control volumes and t_c is proportional to the number of control volumes adjacent to the boundary. Therefore, if the grid size is doubled in all directions, it is expected that t_c/t_{cp} decreases to approximately one half. If t_c is small compared to t_{cp}, then t_c/t_e will also decrease to approximately one half. This is confirmed by the results presented in Table I. In the DDP the same reasoning can be applied at the processor level, i.e. the evolution of t_c/t_e with the grid size is similar for both parallelization methods.

If the absorption coefficient of the medium increases, n_{iter} becomes smaller (see Table I). This may be explained by equation (2). The emission term in that equation is known if temperature field is prescribed, as in the present test case. This term increases with κ and becomes larger compared with the other terms of the equation which change during the course of the iterative procedure. Therefore, the convergence rate tends to increases with κ. In both parallelization methods t_c/t_e is independent of the absorption coefficient, as expected. Therefore, the efficiency and the speedup depend only on the

Variable	Parallelization method	p	Variable value	ε (%)	S	n_{iter}	t_c/t_e (%)
Angular discretization	DDP	9	$N_\theta=3$, $N_\varphi=1$	36.6	3.3	19	9.2
			$N_\theta=5$, $N_\varphi=2$	36.5	3.3	19	9.0
	ADP	8	$N_\theta=3$, $N_\varphi=1$	87.6	7.0	8	7.5
			$N_\theta=5$, $N_\varphi=2$	95.8	7.7	8	2.5
Grid Size	DDP	9	$N_x=N_y/3=N_z=6$	31.6	2.8	19	15.6
			$N_x=N_y/3=N_z=12$	36.5	3.3	19	9.0
	ADP	8	$N_x=N_y/3=N_z=6$	92.3	7.4	8	4.9
			$N_x=N_y/3=N_z=12$	95.8	7.7	8	2.5
Absorption coefficient (m^{-1})	DDP	9	0.2	36.5	3.3	19	3.7
			10	67.3	6.1	9	3.7
	ADP	8	0.2	95.8	7.7	8	2.5
			10	95.8	7.7	5	2.5

Table I. Influence of the angular discretization, grid size and absorption coefficient of the medium on the parallel performance for test case 1

convergence rate. If the ADP is used, n_{iter} is independent of both p and κ and so neither the efficiency nor the speedup change with κ. However, if the DDP is used n_{iter} increases with p. Hence, the ratio of n_{iter} using p processors to n_{iter} using one processor approaches the unity as κ increases. This yields an increase of both ε and S with p, while ε and S per iteration remain constant.

Test case 2

In this test case a rectangular furnace ($4\times2\times2m^3$) is also studied, but a volumetric heat source (5 kW/m^3) is prescribed rather than the temperature field. The temperature and the emissivity of the walls are 1200 K and 0.85 for x = 0m, 400 K and 0.70 for x = 4m, and 900 K and 0.70 for the remaining boundaries, respectively. The standard calculations were carried out using a Cartesian uniform grid with $30\times15\times15$ control volumes, $N_\theta = 5$, $N_\varphi = 2$ and $\kappa = 0.5m^{-1}$. This problem was studied by other authors using the zone and the spherical harmonics methods[12], and the discrete transfer and discrete ordinates methods[13]. The present results are in good agreement with those computed by the other methods, as demonstrated in Figure 5. This figure shows several computed gas temperature profiles and the absolute value of the net heat fluxes at the symmetry plane of the firing and exit ends of the furnace. The results of the finite volume method are compared with those reported in [12] for the zone method. In the following calculations it was systematically checked that the results are independent of the number of processors.

The influence of the number of processors on the parallel performance is shown in Figure 6, and exhibits the same trends observed for test case 1. In the DDP the efficiency is a little better than in test case 1, because n_{iter} did not increase with the increase of p as much as before. If the ADP is employed the efficiency drops fast with the increase of p. In this case $\varepsilon = 17.7\%$ for p = 80. In fact, although n_{iter} is independent of p, the ratio t_c/t_e increases very fast with p. When p = 80, t_c is 73 per cent of the total wall clock time. This is related to the need to broadcast the incident radiation in every control volume (step 5 of the solution algorithm described before). This global operation is not explicitly available as part of the communication software of the Parsytec, and needs to be implemented according to a binary tree algorithm.

The influence of the angular discretization, grid size and absorption coefficient of the medium on the parallel performance is shown in Table II. The role of the angular discretization is the same in both test cases and it will not be further discussed. The grid size also has a similar influence on both cases if the DDP is used. However, if the ADP is used the ratio t_c/t_e decreases very slowly with the grid size. In fact, the communication time comprises two contributions. One of them is the time needed to transfer the halo data which is proportional to the number of control volumes on the boundary of a processor. The other one is the time required to transfer the incident radiation in each control volume which is proportional to the number of control volumes assigned to a processor. If $N_x=N_y=N_z=N$, then the first contribution is proportional to $6N^2$ and the

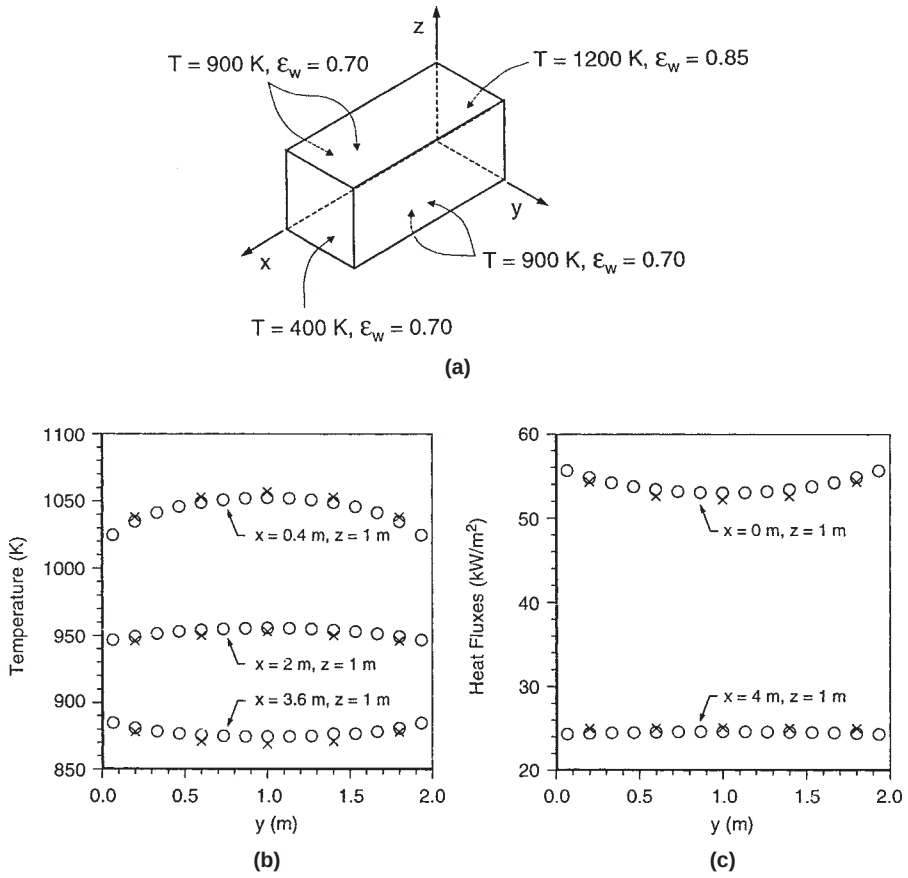

Figure 5.
Rectangular furnace with a prescribed volumetric heat source. (a) Schematic of the furnace. (b) Predicted gas temperature profiles (O: finite volume method; X: zone method[12]). (c) Predicted absolute value of the net heat fluxes (O: finite volume method; X: zone method[12])

second one is proportional to N^3. Since t_{cp} is also proportional to N^3 and t_{cp} is much larger than t_c for p = 8, i.e. the standard case, it follows that t_c/t_e decreases slowly with N, as shown in Table II. Since n_{iter} is independent of p, then the efficiency increases slowly with the grid size.

The number of iterations required to achieve convergence increases markedly with the absorption coefficient of the medium, contrary to the evolution reported in test case 1. However, in the present case the temperature field is not given and, therefore, I_b needs to be computed during the iterative procedure using the energy equation. Hence, the term of equation (2) associated with I_b becomes dominant as κ increases, yielding a corresponding increase of n_{iter}. The ratio t_c/t_e is not influenced by κ. In the ADP the efficiency and the speedup are independent of κ, as already observed in test case 1, because n_{iter} is independent of p for a fixed κ. In the DDP n_{iter} increases with p, but the ratio of n_{iter} using p processors to n_{iter} using one processor decreases with κ. This justifies the observed increase of ε and S with κ.

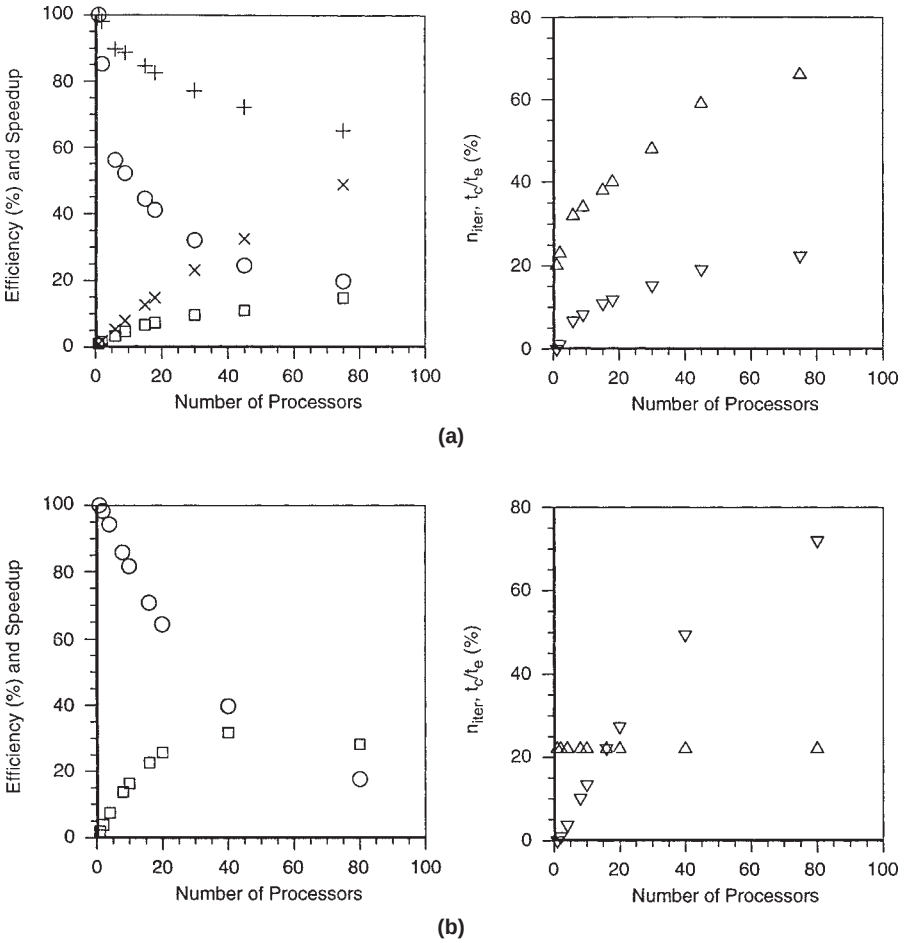

Figure 6.
Influence of the number of processors on the efficiency, speedup, number of iterations and ratio of communication to execution time for test case 2 (O: efficiency; +: efficiency per iteration; □: speedup; x: speedup per iteration; Δ: n_{iter}; ∇: t_c/t_e (%)).
(a) DDP. (b) ADP

Table II.
Influence of the angular discretization, grid size and absorption coefficient of the medium on the parallel performance for test case 2

Variable	Parallelization method	p	Variable value	ε (%)	S	n_{iter}	t_c/t_e (%)
Angular discretization	DDP	9	$N_\theta=3$, $N_\varphi=1$	50.7	4.6	35	8.3
			$N_\theta=5$, $N_\varphi=2$	52.1	4.7	34	8.3
	ADP	8	$N_\theta=3$, $N_\varphi=1$	65.1	5.2	22	25.1
			$N_\theta=5$, $N_\varphi=2$	85.9	6.9	22	10.1
Grid size	DDP	9	$N_x=N_y/2=N_z=9$	48.9	4.4	33	12.4
			$N_x=N_y/2=N_z=15$	52.1	4.7	34	8.2
	ADP	8	$N_{x/2}=N_y=N_z=9$	84.1	6.7	22	11.1
			$N_{x/2}=N_y=N_z=15$	85.9	6.9	22	10.1
Absorption coefficient (m^{-1})	DDP	9	0.1	39.0	3.5	25	8.3
			10	83.2	7.5	213	8.3
	ADP	8	0.1	85.9	6.9	13	10.0
			10	85.9	6.9	252	10.1

Conclusions

The finite volume method for radiative heat transfer calculations was parallelized using either angular domain or spatial domain decomposition strategies. The method was applied to three-dimensional rectangular enclosures containing an emitting-absorbing medium. The following conclusions may be drawn from the work carried out:

- If the DDP is used both n_{iter} and t_c/t_e, increase significantly with p. Consequently, the efficiency drops fast with the increase of p. If the ADP is used the convergence rate is independent of p. However, the ratio t_c/t_e increases with p, especially if a volumetric heat source is prescribed, and t_c is a significant fraction of t_e. Therefore, the efficiency also exhibits a fast drop with the increase of p.
- The number of iterations required to achieve convergence is independent of the angular and spatial discretizations. However, it is strongly influenced by the absorption coefficient of the medium. If the temperature of the medium is prescribed, n_{iter} decreases with the increase of κ, and the reverse occurs if a volumetric source is prescribed.
- The ratio t_c/t_e is independent of the angular discretization in the DDP, but decreases with the refinement of the angular discreatization in the ADP. That ratio decreases with the spatial grid refinement, but the influence is very small if the ADP is used and a volumetric heat source is prescribed. The absorption coefficient of the medium does not influence the ratio t_c/t_e.
- As a consequence of the influences referred to in the points above, the efficiency is independent of the angular discretization in the DDP, but improves with the angular refinement in the ADP. The efficiency is also improved when the grid is refined for both parallelization methods. The efficiency increases with the absorption coefficient of the medium if the DDP is employed, but it does not depend on the absorption coefficient if the ADP is used.

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, 1990, pp. 415-23.
2. Gritzo, L.A., Skocypec, R.D. and Tong, T.W., "The use of high-performance computing to solve participating media radiative heat transfer problems – results of an NSF workshop", *Sandia Report*, SAND95-0225, 1995.
3. Howell, J.R., "Thermal radiation in participating media: the past, the present, and some possible futures", *Journal of Heat Transfer*, Vol. 110, 1998, pp. 1220-9.
4. Gonçalves, J. and Coelho, P.J., "Parallelization of the discrete ordinates method", *Numerical Heat Transfer, Part B: Fundamentals*, Vol. 32, 1997, pp. 151-73.
5. Novo, P.J., Coelho, P.J. and Carvalho, M.G., "Parallelization of the discrete transfer method: two different approaches", *ASME HTD*, Vol. 235, 1996, pp. 45-54.

6. Coelho, P.J., Gonçalves, J. and Novo, P., "Parallelization of the discrete ordinates method: two different approaches", *Vector and Parallel Processing – VECPAR'96*, Springer-Verlag, Lecture Notes in Computer Science, 1215, 1997, pp. 222-35.
7. Chai, J.C., Lee, H.S. and Patankar, S.V., "Finite volume method for radiation heat transfer", *Journal of Thermophysics and Heat Transfer*, Vol. 8, 1994, pp. 419-25.
8. Chai, J.C., Patankar, S.V. and Lee, H.S., "Evaluation of spatial differencing practices for the discrete-ordinates method", *J. Thermophysics and Heat Transfer*, Vol. 8, 1994, pp. 140-4.
9. Hyde, D.J. and Truelove, J.S., "The discrete ordinates approximation for multi-dimensional radiant heat transfer in furnaces", AERE R-8502, AERE Harwell, UK, 1977.
10. Jamaluddin, A.S. and Smith, P.J., "Predicting radiative transfer in rectangular enclosures using the discrete ordinates method", *Combustion Science and Technology*, Vol. 59, 1988, pp. 321-40.
11. Carvalho, M.G., Farias, T. and Fontes, P., "Predicting radiative heat transfer in absorbing, emitting and scattering media using the discrete transfer method", *ASME FED*, Vol. 160, 1991, pp. 17-26.
12. Mengüç, M.P. and Viskanta, R., "Radiative transfer in three-dimensional rectangular enclosures containing inhomogeneous anisotropically scattering media", *J. Quant. Spectrosc. Radiative Transfer*, Vol. 33, 1985, pp. 533-49.
13. Coelho, P.J., Gonçalves, J.M. and Carvalho, M.G., "A comparative study of radiation models for coupled fluid flow/heat transfer problems", *Numerical Methods in Thermal Problems*, Vol. IX, Part 1, 1995, pp. 378-89.