Journal of Computational Physics **172,** 543–571 (2001)
doi:10.1006/jcph.2001.6836, available online at http://www.idealibrary.com on IDEAL®

Implicit Monte Carlo Diffusion—An Acceleration Method for Monte Carlo Time-Dependent Radiative Transfer Simulations

N. A. Gentile

University of California, Lawrence Livermore National Laboratory, Livermore, California 94550
E-mail: gentile1@llnl.gov

Received September 24, 2000; revised April 30, 2001

Implicit Monte Carlo (IMC) is often employed to numerically simulate radiative transfer. In problems with regions that are characterized by a small mean free path, IMC can take a prohibitive amount of time, because many particle steps must be simulated to advance the particle through the time step. Problems containing regions with a small mean free path can frequently be accurately simulated much more quickly by employing the diffusion equation as an approximation. However, the diffusion approximation is not accurate in regions of the problem where the mean free path is large.

We present a method for accelerating time-dependent Monte Carlo radiative transfer calculations by using a discretization of the diffusion equation to calculate probabilities that are used to advance particles in regions with small mean free paths. The method is demonstrated on problems with one-and two-dimensional orthogonal grids. It results in decreases in run time of more than an order of magnitude on these problems, while producing answers with accuracy comparable to pure IMC simulations. We call the method Implicit Monte Carlo Diffusion, which we abbreviate IMD.

1. INTRODUCTION

The time-dependent transport equation for gray photons in the absence of scattering is [1]

$$\frac{1}{c}\frac{\partial I}{\partial t} + \hat{\Omega}\cdot\nabla I = -\sigma_a I + \frac{c\sigma_a a T^4}{4\pi}, \tag{1}$$

where c is the speed of light, σ_a is the macroscopic absorption cross section in inverse length units, and T is the matter temperature. The transport equation is coupled to the material

0021-9991/01 $35.00

energy balance equation [1]

$$\frac{\partial E_m}{\partial t} = \rho c_v \frac{\partial T}{\partial t} = \sigma_a \int I \, \mathrm{d}\Omega - c\sigma_a a T^4. \tag{2}$$

Here, E_m is the matter energy density in units of energy per volume, ρ is the mass density, and c_v is the specific heat capacity in units of energy per mass per temperature. These equations can be solved by a Monte Carlo method described in [2]. The method discretizes the problem on a mesh. Each zone has a temperature and an absorption cross section. Particles representing photons are created in the zones at the beginning of each time step according to the emission term in the transport equation. Then the photons are followed through the zones, which heats them according to the absorption term in Eq. (1). The temperatures are updated at the end of the time step, using Eq. (2), and the process is repeated.

This method becomes unstable when time steps of the order of

$$\Delta t = \frac{\rho c_v}{a T^3 c \sigma_a} \tag{3}$$

are taken [3]. This instability occurs when the matter and radiation fields exchange an amount of energy comparable to the amount of energy necessary to change the matter temperature a nonnegligible amount in one time step. If the matter is only able to absorb energy during a time step, but is not able to reradiate, as in the algorithm in [2], then instabilities will occur. The inability of the matter to reradiate the energy it absorbs from the radiation during a time step is related to the fact that the temperature in the emission term of the transport equation is calculated using the temperature at the beginning of the time step.

A method for solving the photon transport equation with improved stability when large time steps are taken was provided by Fleck and Cummings [3]. The method was dubbed Implicit Monte Carlo, usually abbreviated IMC. IMC works by using the matter energy balance equation to estimate the future matter temperature, and using this estimate in the transport equation. This substitution has the effect of reducing the absorption opacity in the transport equation by a factor of

$$f = \frac{1}{1 + \beta c \Delta t \sigma_a} \tag{4}$$

and adding an equal amount of thermally redistributed isotropic scattering. Here $\beta = 4aT^3/\rho c_v$. This change allows the calculation to be run with much larger time steps before instabilities arise [4].

The factor f is small when photons are being absorbed and quickly reemitted by the matter. Problems in which this occurs are said to exhibit tight coupling between the radiation and matter. IMC replaces the absorption and rapid reemission occurring in tightly coupled problems with isotropic scattering. This scattering is usually referred to as the effective scattering, to distinguish it from physical scattering. The effective scattering $\sigma_s = (1 - f)\sigma_a$.

When the scattering, either physical or effective, is large, then the mean free path of photons can be much smaller than a typical dimension of the zones in the discretized mesh. IMC particles take many steps in these zones. Each simulated photon path in an IMC calculation is about equally expensive; thus, simulations with a large scattering cross section

can be very time consuming to calculate. The end result of applying the IMC algorithm is to enable the use of much larger time steps, at the cost of making highly absorbing, tightly coupled problems as expensive to run as highly scattering ones.

When the factor f is small, most of the absorption opacity is replaced by an effective scattering opacity which is isotropic. If the mean free path for this effective scattering is small, then the solution of the transport equation will be well approximated by the solution of the diffusion equation [5].

The diffusion equation describes the time development of the radiation energy density, which is the zeroth angular moment of the intensity I;

$$\frac{\partial E}{\partial t} + \nabla \cdot F = c\sigma_a a T^4 - c\sigma_a E, \tag{5}$$

where E, the radiation energy density, is defined by

$$E = \frac{1}{c} \int I \, \mathrm{d}\Omega, \tag{6}$$

and the flux F, the first angular moment of I is defined by

$$F = \int \hat{\Omega} I \, \mathrm{d}\Omega. \tag{7}$$

To allow us to calculate E from this equation, we must define F in terms of E. This is usually done by using Fick's law,

$$F = -cD\nabla E, \tag{8}$$

where $D = 1/(3\sigma)$ is the diffusion coefficient.

The diffusion equation can be derived from the transport equation by taking angular moments of the transport equation and taking the limit as the factor

$$\epsilon = \frac{\sigma^{-1}}{L} = \frac{\textit{mean free path}}{\textit{characteristic length of the flux}} \tag{9}$$

becomes small [5]. The solution of the diffusion equation is an accurate approximation to the first angular moment of the solution of the transport equation when the intensity describes photons with a nearly isotropic angular distribution which is slowly varying in space and time. Thus, the diffusion equation can provide an accurate approximation for highly scattering problems where IMC is prohibitively expensive.

The numerical solution of the diffusion equation is usually more rapid than the numerical solution of the transport equation in situations where the diffusion approximation is applicable [1]. Since IMC is expensive where diffusion is accurate, solution techniques have been developed that employ IMC in the parts of the problem with a small effective scattering, and some form of diffusion in the parts of the problem with a large effective scattering [6]. These are referred to as hybrid methods.

A hybrid technique involves solving the diffusion equation on some regions of the grid and using IMC on other regions. The IMC simulation provides a flux that is used as a boundary condition for the solution of the diffusion equation, which usually requires a

matrix inversion. The flux of energy out of the diffusion region is turned into particles used by the IMC in the next time step.

This paper presents a new hybrid scheme, which is based on a Monte Carlo solution of the diffusion equation. This method uses the matrix resulting from discretizing the diffusion equation to derive probabilities for particles to deposit energy, reach census, or jump to another zone. Since this method involves a Monte Carlo solution that is similar to IMC, and uses the same stabilization technique as IMC, we have dubbed it Implicit Monte Carlo Diffusion, which we abbreviate IMD.

The IMD particle can jump to a new zone in one step, rather than taking many IMC steps. Our method in effect rolls many expensive IMC steps into one very cheap IMD step. Thus the calculation proceeds much more rapidly in the diffusion region than it would if we employed IMC there. Particles can freely cross the boundary between the diffusion region and the IMC region, and so the two methods are easy to couple together.

In the following sections, we develop the IMD algorithm and show how to make a hybrid IMC/IMD method. In Section 2, we describe the discretization of the diffusion equation. In Section 3, we show how to solve the matrix equation obtained by this discretization by a Monte Carlo method which resembles IMC. In Section 4, we describe how to use the Monte Carlo diffusion method with IMC in a hybrid method. In Section 5, we compare and contrast this hybrid method to other hybrid methods. In Section 6, we apply this hybrid method to various gray opacity test problems. We show that it can be considerably faster than IMC alone on problems with very opaque regions, while yielding a similar result.

2. DISCRETIZING THE DIFFUSION EQUATION

We begin by considering the diffusion equation in Cartesian coordinates in a one-dimensional slab geometry, as derived in [7]. Our development will rely heavily on this derivation. The diffusion equation under these assumptions is [7]

$$\frac{1}{c}\frac{\partial E}{\partial t} - \frac{\partial}{\partial x}\left[D\frac{\partial E}{\partial x}\right] = \sigma_a a T^4 - \sigma_a E. \tag{10}$$

We discretize this equation as in [7]. We take E to be a zone-centered variable, and use backward Euler time differencing. The result is

$$\frac{1}{c}\frac{E_j^{n+1} - E_j^n}{\Delta t} - \frac{F_{j+1/2}^{n+1} - F_{j-1/2}^{n+1}}{\Delta x_j} = \sigma_{a\,j}\left(a T_j^{n4} - E_j^{n+1}\right). \tag{11}$$

Here j is a zone index, and $j + (-)1/2$ indicates the face in the increasing (decreasing) x direction.

In zones with neighbors, we can get F at the edges by discretizing Fick's law as

$$F_{j+1/2}^{n+1} = cD_{j+1/2}\frac{E_{j+1}^{n+1} - E_j^{n+1}}{\Delta x_{j+1/2}}, \tag{12}$$

where

$$D_{j+1/2} = 2\Delta x_{j+1/2}\left[\frac{D_j D_{j+1}}{D_j \Delta x_{j+1} + D_{j+1}\Delta x_j}\right] \tag{13}$$

and

$$\Delta x_{j+1/2} = (\Delta x_j + \Delta x_{j+1})/2. \tag{14}$$

(As noted in [7], the harmonic average given in Eq. (13) will yield a small value for the face diffusion coefficient in cases where either zone has a small diffusion coefficient. This can occur in problems with temperature-dependent opacities that are large in cold matter. A small diffusion coefficient leads to a small flux via Eq. (12); thus, the diffusion of heat into cold, opaque material may be unphysically retarded. In [7], a means of dealing with this problem is outlined: essentially, the values of D employed in face calculations such as Eq. (13) are calculated using a common temperature derived from the temperatures of the neighboring zones. In this paper, this issue does not arise, because we have run only simulations with temperature-independent opacities.)

For the two edges at the ends of the problem, we get F from the boundary condition given in [7],

$$F = \frac{c}{4}\left[E + 2D\hat{n}\cdot\nabla E\right], \tag{15}$$

where $\hat{n}$ is the outward normal at the edge.

(We should note an issue, described in [7], that arises in the calculation of the diffusion coefficient D. This is the use of a flux limiter to prevent superluminal energy transport in regions with a small opacity. Several flux limiters and their effects are considered in [8]. Flux limiters are not used in the IMD method, because IMC is used in the places where the flux limiter would be employed in a diffusion approximation. This is discussed further in the next section.)

The source term in Eq. (11) depends on the current matter temperature T_j^n, rather than the future matter temperature T_j^{n+1}, which would make the time differencing fully implicit. This can affect the quality of the solution, as discussed in [7]. This dependence on the current matter temperature rather than the future matter temperature is what led to the time step restriction given by Eq. (3) in the Monte Carlo solution of the transport equation.

This problem is addressed in [7] by iterating on the matter temperature, but that requires multiple matrix inversions per time step. Instead, we choose to solve it in the same manner employed to stabilize the IMC algorithm. We will use the matter temperature equation, Eq. (2), to get an estimate of the future matter temperature, and use that estimate for T in the $\sigma_a a T^4$ term in Eq. (10). Our development will closely follow that in [3], and has the same purpose—to allow us to use larger time steps in our simulation without encountering instability.

We begin by defining $E_r \equiv aT^4$ and using the chain rule to relate the time derivative of E_r to that of E_m:

$$\frac{\partial E_r}{\partial t} = \frac{\partial E_r}{\partial E_m}\frac{\partial E_m}{\partial t}. \tag{16}$$

We use the chain rule again to get an expression for $\partial E_r/\partial E_m$, which is

$$\frac{\partial E_r}{\partial E_m} = \frac{\rho c_v \partial T/\partial t}{4aT^3 \partial T/\partial t} = \frac{\rho c_v}{4aT^3} = \beta, \tag{17}$$

where β is defined as in Eq. (4). Substituting these last two results, and the definition of E given by Eq. (6), in Eq. (2), we get an expression for the time derivative of E_r:

$$\frac{\partial E_r}{\partial t} = \beta c\sigma_a E - \beta c\sigma_a E_r. \tag{18}$$

Next, we difference Eq. (18) by backward Euler to get an expression for the future value of E_r:

$$E_r^{n+1} = E_r^n + \Delta t \left[\beta c\sigma_a E - \beta c\sigma_a E_r^{n+1}\right]. \tag{19}$$

Collecting terms containing E_r^{n+1} yields

$$E_r^{n+1}(1 + \beta c\Delta t\sigma_a) = E_r^n + \beta c\Delta t\sigma_a E, \tag{20}$$

and solving this for E_r^{n+1} results in

$$E_r^{n+1} = f E_r^n + (1 - f)E, \tag{21}$$

where f is the same factor, defined by Eq. (4), that is employed in IMC.

Finally, using the definition of E_r as aT^4, we use E_r^{n+1} as the estimate for aT^4 in Eq. (10). The result of this substitution is

$$\frac{1}{c}\frac{\partial E}{\partial t} - \frac{\partial}{\partial x}\left[D\frac{\partial E}{\partial x}\right] = f\sigma_a aT^4 - f\sigma_a E, \tag{22}$$

where T is regarded as the t^n value. As in IMC, this transformation results in effective isotropic scattering, which shows up in the diffusion coefficient. In the presence of scattering, D is defined in terms of the sum of the absorption and scattering opacities, so its value remains unchanged.

This transformed diffusion equation could also have been derived by starting with the transport equation as modified by IMC, integrating over angle, and applying Fick's law. This derivation would parallel the one in [5] with a factor of f modifying σ_a everywhere and a scattering opacity of $(1 - f)\sigma_a$ in the transport equation.

We now difference Eq. (22) in time. The result is Eq. (11), with the addition of factors of f multiplying the source terms on the right-hand side. Using the definitions of the flux from Eq. (12) and collecting coefficients results in

$$\left[1 + c\Delta t f_j\sigma_{a\,j} + c\frac{\Delta t}{\Delta x_j}\frac{D_{j-1/2}}{\Delta x_{j-1/2}} + c\frac{\Delta t}{\Delta x_j}\frac{D_{j+1/2}}{\Delta x_{j+1/2}}\right] E_j^{n+1} - c\frac{\Delta t}{\Delta x_j}\frac{D_{j-1/2}}{\Delta x_{j-1/2}}E_{j-1}^{n+1}$$
$$- c\frac{\Delta t}{\Delta x_j}\frac{D_{j+1/2}}{\Delta x_{j+1/2}}E_{j+1}^{n+1} = c\Delta t f\sigma_{a\,j} aT_j^{n4} + E_j^n. \tag{23}$$

At the $x = 0$ boundary, we get a similar equation relating E_1^{n+1} and E_2^{n+1} by using a difference formulation of Eq. (15) for $F_{-1/2}^{n+1}$ and inserting the result into Eq. (11). The result is

$$\left[1 + c\Delta t f_1\sigma_{a\,1} + c\frac{\Delta t}{\Delta x_1}\frac{3/2D_1\sigma_{a\,1}}{1 + 3\Delta x_1\sigma_{a\,1}/4} + c\frac{\Delta t}{\Delta x_j}\frac{D_{3/2}}{\Delta x_{3/2}}\right] E_1^{n+1} - c\frac{\Delta t}{\Delta x_1}\frac{D_{3/2}}{\Delta x_{3/2}}E_2^{n+1}$$
$$= c\Delta t f\sigma_a aT_1^{n4} + E_1^n + 4\left[\frac{3/2D_1\sigma_{a\,1}}{1 + 3\Delta x_1\sigma_{a\,1}/4}\right]\frac{F_0}{c\Delta x_1}. \tag{24}$$

In this equation, F_0 is the flux through the boundary at $x = 0$, which adds energy into the source term in zone 1. A similar equation is obtained at the $x = x_{max}$ end of the problem.

The equation defining E_j^{n+1} in terms of the neighboring values E_{j+1}^{n+1} and E_{j-1}^{n+1} is a matrix equation $Ax = b$, with the E_j^{n+1} the components of the unknown x and the source terms the components of b. The matrix A is tridiagonal and can easily be solved by standard techniques [9]. In the next section, we outline a Monte Carlo technique for solving it.

In a manner similar to that applied above, we can get a matrix equation for a zone-centered discretization of the diffusion equation in cylindrical coordinates in a two-dimensional axially symmetric geometry on an orthogonal mesh. Employing the usual five-point differencing scheme results in a similar set of equations which has five off-diagonal bands.

In both the Cartesian one-dimensional and the orthogonal cylindrical case, the matter energy density satisfies the same equation:

$$\frac{\partial E_m}{\partial t} = f\sigma_a \int I \, \mathrm{d}\Omega - cf\sigma_a a T^4. \tag{25}$$

This is the same equation satisfied by the matter energy density in the IMC formulation.

The change in the matter energy density given by Eq. (25) is the difference between the energy thermally radiated by the matter and the energy absorbed from the radiation field. Often this equation is solved by introducing the heat capacity and writing

$$\rho c_v \frac{\partial T}{\partial t} = f\sigma_a \int I \, \mathrm{d}\Omega - cf\sigma_a a T^4, \tag{26}$$

which can be differenced and solved for the temperature if c_v is assumed constant. This difference equation only conserves energy if c_v is actually constant. We prefer to difference Eq. (25) as

$$E_{m\,j}^{n+1} = E_{m\,j}^{n} + E_{absorbed\,j} - \Delta t c f_j \sigma_{a\,j} a T_j^{n4}, \tag{27}$$

where $E_{absorbed\,j}$ is the amount of energy absorbed by the matter from the radiation field and is obtained from the solution of the diffusion equation.

Using Eq. (27), $E_{m\,j}^{n+1}$ is solved for the new matter energy at t^{n+1}, which can be numerically inverted using the equation of state to obtain the new matter temperature T^{n+1}.

3. SOLVING THE DISCRETIZED DIFFUSION EQUATION BY A MONTE CARLO TECHNIQUE

The matrix equations arising from the discretizations of the one-dimensional Cartesian and the two-dimensional orthogonal cylindrical diffusion equations are similar. In both cases, the diagonal element consists of the sum of the following terms: unity, arising from the time derivative term; $c\Delta t f \sigma_a$, arising from the absorption term; and several terms, one for each neighboring zone, involving the diffusion coefficient and geometric factors, which arise from the flux term. The off-diagonal elements are the negatives of these diffusion coefficient terms. Both matrices are diagonally dominant, and symmetric, and hence are symmetric positive definite.

The source terms are the same in both cases, consisting of the sum of the old energy density in the zone E_j^n and a source term depending on the temperature, $c\Delta t f \sigma_a a T^4$.

Here, we derive a Monte Carlo solution technique that is applicable to both the one-dimensional and the two-dimensional discretized diffusion equations (and is in fact applicable to any symmetric matrix equation in which the sign of the diagonal is the opposite of the sign of the off-diagonal terms.) Our derivation employs Eq. (23) above, but the generalization to the two-dimensional case, and other matrixes, is clear.

If we take the matrix equation defined by Eq. (23) and solve it for E_j^{n+1}, we obtain the following relation which defines the radiation energy density in zone j in terms of the radiation energy density of neighboring zones and the energy of the source in zone j as

$$E_j^{n+1} = E_{j-1}^{n+1}\hat{f}_{j-1}^{+} + E_{j+1}^{n+1}\hat{f}_{j+1}^{-} + E_{source\ j}^{n}/\hat{d}_j, \tag{28}$$

where we have made several definitions. The source energy in zone j is defined as

$$E_{source\ j}^{n} \equiv E_j^n + c\Delta t f_j \sigma_{a\ j} a T_j^{n4}; \tag{29}$$

the diagonal term of the matrix, $\hat{d}_j$, which is the coefficient of E_j^{n+1} in Eq. (23) is defined as

$$\hat{d}_j \equiv \left[1 + c\Delta t f_j \sigma_{a\ j} + c\frac{\Delta t}{\Delta x_j}\frac{D_{j-1/2}}{\Delta x_{j-1/2}} + c\frac{\Delta t}{\Delta x_j}\frac{D_{j+1/2}}{\Delta x_{j+1/2}}\right]; \tag{30}$$

and the off-diagonal terms are defined as

$$\hat{f}_j^{+} = c\frac{\Delta t}{\Delta x_j}\frac{D_{j+1/2}}{\Delta x_{j+1/2}}\Big/\hat{d}_j, \tag{31}$$

with $\hat{f}_j^{-}$ defined similarly. We also define $E_{total} = \sum_j E_{source\ j}^{n}$. The symmetry of the matrix is expressed by the fact that $\hat{f}_j^{+} = \hat{f}_{j+1}^{-}$ and $\hat{f}_j^{-} = \hat{f}_{j-1}^{+}$.

We now describe a Monte Carlo technique that can be used to solve Eq. (28) above, which is equivalent to solving the matrix equation Eq. (23) from which it was derived. Then we will show that the estimate for E_j^{n+1} obtained by this technique approaches (in the limit of many particles) the solution of Eq. (28) and hence of Eq. (23).

The Monte Carlo solution technique requires that the following probabilities be defined. The probability that a particle created in zone j will jump to zone $j+1$ is

$$P_j^{+} \equiv \hat{f}_j^{+}/\hat{d}_j. \tag{32}$$

The probability that a particle created in zone j will jump to zone $j-1$ is

$$P_j^{-} \equiv \hat{f}_j^{-}/\hat{d}_j. \tag{33}$$

The probability that the energy of a particle will be tallied into the array representing the solution of E_j^{n+1} is

$$P_j^{c} \equiv 1/\hat{d}_j. \tag{34}$$

The probability that the energy of a particle will be tallied into the array representing the solution of $E_{m\ j}^{n+1}$, the matter energy density, is

$$P_j^{a} \equiv c\sigma_a \Delta t/\hat{d}_j = 1 - P_j^{+} - P_j^{-} - P_j^{c}. \tag{35}$$

The expressions defined in Eqs. (32)–(35) all satisfy the requirement that they be greater than or equal to zero and less than or equal to unity, and they all add to unity. This can be easily demonstrated. Since the diffusion coefficient, Δx, and the Δt are all positive, $\hat{f}$ is positive. Since the factor f defined by Eq. (4) and the absorption opacity σ are positive, every term in $\hat{d}$ is positive. The probabilities defined by Eqs. (32)–(35) are ratios of positive numbers and so are greater than zero. All are ratios of some fraction of $\hat{d}$ to $\hat{d}$, and so they are all less than or equal to unity. Summing them yields unity. Thus the terms defined by Eqs. (32)–(35) meet the requirements that probabilities must satisfy.

The Monte Carlo solution technique begins with the creation of a total number N of particles distributed among the zones of the mesh. The number of particles created in zone j is $N_j = N E^n_{source\ j}/E_{total}$. Each particle created in zone j has energy $E_{particle} = E_{source\ j}/N_j$. The sum of $E_{particle}$ over all particles and zones is E_{total}, so energy is strictly conserved during creation of the particles.

Each particle now undergoes one of four processes. It can jump to zone $j+1$ with a probability P^+_j, or it can jump to zone $j-1$ with a probability P^-_j. Its energy can be tallied into E^{n+1}_j with a probability P^c_j. We refer to this event as census, because it will be seen to be analogous to the census event in the IMC algorithm. Or the energy of the particle can be tallied into another variable $E_{absorbed\ j}$ with probability P^a_j. $E_{absorbed\ j}$ will be seen to be the energy absorbed by the matter in zone j. The determination of which of the four alternatives occurs is done by comparing a random number to the probabilities.

If the particle jumps to a new zone, we continue the process, using the four probabilities for the new zone. If the particle tallies into either E^{n+1}_j or $E_{absorbed\ j}$, or leaves the problem through the ends, we are finished advancing it. A particle leaves the problem when it is in the first zone, $j = 1$, and the particle is determined to jump to zone $j-1$, or it is in the last zone $j = j_{max}$ and it is determined to jump to zone $j+1$. When all the particles are either tallied or have left the problem, we are through with the time step.

When we are through with the time step, we have computed two numbers for each zone by the Monte Carlo process: E^{n+1}_j and $E_{absorbed\ j}$. E^{n+1}_j is the estimate for the new radiation energy density in the zone, and is used as the initial condition for the new time step. $E_{absorbed\ j}$ is the amount of energy absorbed by the matter in the zone. The net energy change of the matter in the zone is thus $E_{absorbed\ j} - c\Delta t f_j \sigma_{a\ j} a T^{n4}_j$. This net energy change is used along with the equation of state to get the new temperature T^{n+1}_j of the material. This gives us all the information we need to repeat the process and do another time step.

We now show that the expected value of E^{n+1}_j, which we denote $\langle E^{n+1}_j \rangle$, obtained from this algorithm satisfies Eq. (28).

The expected value $\langle E^{n+1}_j \rangle$ satisfies

$$\left\langle E^{n+1}_j \right\rangle = P^c_j \lim_{N\to\infty} E_{particle} N_j, \tag{36}$$

where N_j is the number of particles that pass through the zone j. A particle in zone j must either be born in zone j or jump there from another zone. This implies that

$$\lim_{N\to\infty} N_j = P^-_{j+1} \lim_{N\to\infty} N_{j+1} + P^+_{j-1} \lim_{N\to\infty} N_{j-1} + N_{born},$$

where

$$N_{born} = N_{total} E_{source\ j}/E_{total} = E_{source\ j}/E_{particle}.$$

Using these relations, we obtain

$$\left\langle E_j^{n+1} \right\rangle = P_j^c P_{j+1}^- \lim_{N\to\infty} E_{particle} N_{j+1} + P_j^c P_{j-1}^+ \lim_{N\to\infty} E_{particle} N_{j-1} + P_j^c E_{source\ j}. \quad (37)$$

Using the symmetry of the matrix, the terms involving the products of probabilities satisfy

$$P_{j-1}^+ = \frac{\hat{f}_{j-1}^+}{\hat{d}_{j-1}} = \frac{\hat{f}_j^-}{\hat{d}_j} \frac{\hat{d}_j}{\hat{d}_{j-1}} = P_j^- P_j^c / P_{j-1}^c, \quad (38)$$

which leads to $P_{j-1}^+ P_j^c = P_j^- P_j^c$. Similarly, $P_{j+1}^- P_j^c = P_{j+1}^c P_j^+$.

Using these relations, and the definition of $P_j^c = 1/\hat{d}_j$, we find that

$$\left\langle E_j^{n+1} \right\rangle = P_j^- \lim_{N\to\infty} P_{j+1}^c E_{particle} N_{j+1} + P_j^+ \lim_{N\to\infty} P_{j-1}^c E_{particle} N_{j-1} + E_{source\ j}/\hat{d}_j. \quad (39)$$

Using the definition of the expected value, this becomes

$$\left\langle E_j^{n+1} \right\rangle = P_j^- \left\langle E_{j+1}^{n+1} \right\rangle + P_j^+ \left\langle E_{j-1}^{n+1} \right\rangle + E_{source\ j}/\hat{d}_j. \quad (40)$$

The expected values for the energy produced by the algorithm satisfy the difference equation Eq. (28) we wish E_j^{n+1} to satisfy. Thus, we can take the results for E_j^{n+1} we get from the Monte Carlo simulation, which are obtained with a finite number of particles, as an approximate solution of the equation. Since the E_j^{n+1} represent the photon energy in zone j at time $n+1$, the tally of a photon into E_j^{n+1} is analogous to an IMC particle reaching census. On the next time step, the values for E_j^{n+1} become the E_j^n used in the source term.

A similar calculation shows that Eq. (24) can be manipulated to give probabilities that hold in zone 1, and that these give rise to expectation values that satisfy Eq. (24). The probabilities for zone 1 are derived by dividing the equation by the diagonal term, just as they are in the case of Eq. (23). In particular, the value of P_1^-, the probability that the particle leaves the problem though the boundary at $x = 0$, comes from the term in the denominator that arises from the boundary term in Eq. (24):

$$P_1^- = c \frac{\Delta t}{\Delta x_1} \frac{3/2 D_1 \sigma_{a\ 1}}{1 + 3\Delta x_1 \sigma_{a\ 1}/4} \Big/ \hat{d}_1. \quad (41)$$

A similar term arises at the $x = x_{max}$ boundary. We tally the energy of particles that leave the problem into the variable $E_{escaped}$.

The IMD particles collectively carry E_{total}, and they deposit energy in a conservative way. This implies $\sum_j E_{absorbed\ j} = E_{total} - \sum_j E_j^{n+1} - E_{escaped}$. $E_{absorbed\ j}$, the energy that is not tallied into census, must therefore be the amount of energy that was absorbed by the matter in each zone. This energy is used to solve the matter energy balance Eq. (27) to obtain the new matter temperature T_j^{n+1}.

The algorithm in the two-dimensional orthogonal case is exactly the same, except that there are probabilities of jumping to four neighboring zones rather than two. The probabilities of jumping in the two-dimensional case are given by dividing the flux by the diagonal term, just as in the one-dimensional case. P^c and P^a are given by the same expressions as in the one-dimensional case.

Applying the Monte Carlo solution technique above to the matrix equations produces a very simple, easy to code algorithm with a very physical interpretation, and which resembles IMC very closely.

The Monte Carlo diffusion procedure begins by starting some number of particles in each zone. The weight of each particle in a zone is the value of the source term in that zone divided by the number of particles in that zone. The source term is $E_j^n + c\Delta t f_j \sigma_{a\ j} a T_j^{n4}$. E_j^n is the radiation energy at time n, which corresponds to the photons in census in an IMC calculation, while $c\Delta t f_j \sigma_{a\ j} a T_j^{n4}$ is the energy radiated from the matter in that zone, just as in an IMC calculation. So our particles have weights, interpretable as the energies that IMC particles generated in the zones would have had. We will refer to these particles as IMD particles.

IMD particles are advanced by drawing a random number and comparing it to the probabilities derived from the matrix. The probability of reaching census, i.e., contributing to E_j^{n+1}, is $1/\hat{d}_j$. If this event occurs, the particle's energy is tallied into the variable holding E_j^{n+1}. The probability of absorption is $c\Delta t f_j \sigma_{a\ j}$. If this event occurs, the particle's energy is tallied into the variable holding the amount of energy absorbed by the matter in that zone from the radiation field. The probability of jumping to a new zone k is $\hat{f}_{j,k}/\hat{d}_j$, the ratio of the fringe term associated with zone k to the diagonal of the current zone j. If this event occurs, the particle is moved to zone k, a new random number is drawn, and the particle is advanced again using the probabilities from zone k.

As a variance reduction technique, we can use a path length estimator for the census and absorption events, rather than the last event estimator described above. In the last event estimator, we tally all of the energy of the particle into census or absorption if that is the event selected, and stop advancing the particle. In the path length estimator, we tally an amount equal to the probability of census and absorption multiplied by the particle's energy on each step. Then we subtract the tallied energy from the particle's energy. The particle continues to move indefinitely with a decreasing energy. To prevent wasting computational resources on particles with small weights, the particle is terminated when the weight reaches some small fraction of the initial weight. The result of using this variance reduction method is smoother results for the same number of particles. We have used .01 of the initial energy as the termination energy in the calculations described in Section 6.

IMD, as a Monte Carlo method, produces a solution with statistical noise in it. This can impact the solution of discretized diffusion equations with flux limiters.

As discussed in [7] and [8], flux limiters use some approximation of the slope of the energy to modify the value of diffusion coefficient D when $|\nabla E|/E$ is comparable to σ. This occurs when σ is small or when the radiation energy is changing rapidly, which is when the diffusion approximation is not applicable. Flux limiters are essentially a way to employ diffusion in regions where a full transport solution would be more appropriate.

Flux limiters cause D to depend nonlinearly on ∇E. This nonlinearity would prevent formulating the discretized diffusion equation as a matrix equation, so t^n values for E are often used. Iteration is also sometimes employed.

When IMD is used to invert the matrix, the statistical noise in the values obtained for E will make the calculation of ∇E used in the flux limiter inaccurate. This makes the IMD algorithm a poor choice for solving flux-limited diffusion equations, although one might be able to use it if some kind of smoothing was imposed on the values used to calculate ∇E. IMD is useful when coupled to IMC in a hybrid method. Then IMC can be used where the transport equation is more appropriate (i.e., where a flux limiter would be needed), and IMD can be used without a flux limiter where the diffusion approximation is accurate.

In the calculations described in this paper, we have employed the diffusion approximation only in regions with σ so large that a flux limiter would not be expected to modify the value of D. Experimentation confirmed that the flux limiter given above did not modify the value of D obtained from $D = 1/(3\sigma)$. The hybrid results presented in Section 6 were not run with a flux limiter in IMD.

Advancing the IMD particles is very similar to advancing IMC particles. They both deposit energy in the matter, and reach census. Their tracking behavior is much simpler, however. IMD particles do not have a position in the zone, a time within the time step, or direction cosines. They are spatially associated with the whole zone and temporally associated with the whole time step. This is because the diffusion equation results from an integration over angles, and the discretization that produced the matrix employed zone and time average quantities.

In a zone in which the effective scattering is large, an IMC particle will perform many scatters. Each scatter is relatively expensive, requiring among other things, a calculation of the distance from the IMC particle's current location to the zone boundary, the computation of a natural logarithm to determine the distance to a scatter, and several 'if' tests to determine the particle's fate. Advancing an IMD particle is much cheaper, because this distance to boundary calculation is not needed, and because the particle always leaves the zone during a step. In effect, many expensive IMC steps can be replaced by one inexpensive IMD step, at the cost of replacing the solution of the transport equation with the solution of the diffusion equation. If some region of the problem has a large effective scattering, we may be willing to make this tradeoff, since the increase in speed may be quite large, while the decrease in accuracy may be tolerable. In the next section, we describe the IMC/IMD hybrid method. In the Section 5, we compare the IMC/IMD hybrid method to other hybrid methods. In Section 6, we will demonstrate the IMC/IMD hybrid method on three test problems and examine the results in terms of accuracy and performance.

4. AN IMC/IMD HYBRID METHOD

Hybrid algorithms, with IMC and diffusion running in the same problem, have been described in the literature [6]. In these hybrid algorithms, part of the problem is run with IMC, and the part with a large effective scattering, which would be prohibitively expensive to run with IMC, is run with diffusion. Dividing the problem up in this way means selecting the zones in which diffusion is an appropriate approximation, and forming a matrix equation on each contiguous set of these zones. IMC particles are advanced through the zones not selected as diffusion zones.

The similarity of IMD with IMC makes it easy to mate the two. The problem is divided into IMD and IMC regions by identifying zones where the diffusion approximation is acceptable. Since IMC particles are simulating the transport equation, which is accurate even in the diffusive region, we can allow IMC particles to penetrate freely into these regions. The IMC region does not provide a flux for the diffusion calculation, so no problems involving negative probabilities of reflection result. These negative probabilities arise in the hybrid method described in [6]. Since there is no flux into the diffusion region at its boundary, a vacuum boundary condition is imposed there.

The matter in the problem radiates because of its temperature. In the IMC regions, this radiation is modeled as a source of IMC particles. In the IMD regions, IMD particles are

radiated. Since IMC particles can travel into IMD regions, they may deposit energy into the matter in IMD regions. Thermal emission from these zones in subsequent time steps is done by emitting IMD particles.

In effect, we are solving the transport equation with a source from the matter temperature in zones which are not diffusive:

$$\frac{1}{c}\frac{\partial I}{\partial t} + \hat{\Omega}\cdot\nabla I = -\sigma_a I + \frac{c\sigma_a a T^4}{4\pi}. \tag{42}$$

In zones which are diffusive we are solving the transport equation without a source,

$$\frac{1}{c}\frac{\partial I}{\partial t} + \hat{\Omega}\cdot\nabla I = -\sigma_a I, \tag{43}$$

as well as the diffusion equation with a matter temperature source

$$\frac{1}{c}\frac{\partial E}{\partial t} - \frac{\partial}{\partial x}\left[D\frac{\partial E}{\partial x}\right] = \sigma_a a T^4 - \sigma_a E. \tag{44}$$

In diffusive regions, where both IMC and IMD particles are being advanced, the radiation energy density is the sum of E, solved for in the diffusion equation, and the angular integral of I, solved for in the transport equation. This representation is consistent since both equations are linear in the energy variable. Both IMC and IMD particles contribute to heating the matter in the diffusive regions. This is consistent since the equation for the rate of change of matter energy density, Eq. (2), is linear in the angular integral of I. Just as these linearities allow the IMC algorithm to treat IMC particles as independent of one another, they allow us to treat IMC and IMD particles as independent.

When an IMD particle jumps from a zone in the diffusive region into a zone in the IMC region, it is converted into an IMC particle at the boundary between the zones. The IMC particle has an energy equal to that of the IMD particle. A location on the interface between the zones is calculated stochastically, as it would be for a source particle created on that interface by, for example, a temperature source. Direction cosines are generated from the cosine distribution, and a time between t^n and t^{n+1} is randomly assigned to the particle.

Since there is no natural time associated with IMD particles, we have chosen to assign a random time to the IMC particles that result from the conversion of IMD particles. The selection of a random time for the newly created IMC particle is consistent with the usual treatment of surface fluxes in IMC codes. IMC particles generated from surface fluxes are evenly distributed over a time step, and the IMD particles that are converted are acting as a surface flux for the IMC. This is also the way times are assigned to IMC particles in the hybrid method in [6].

When running an IMC simulation, we usually attempt to keep the total number of IMC particles constant. In a hybrid simulation, we attempt to keep the sum of the total number of IMC and IMD particles constant by apportioning particles between IMC and IMD in the same ratio as the sum of source and census energy of the IMC and IMD particles. This apportionment does not prevent conservation of energy, because we do not demand that the energy of the IMC and IMD particles be exactly the same. However, we do perform Russian rouletting of IMC particles when they are brought out of census, in order to apportion the available IMC particles between census particles and source particles. This procedure does cause conservation of energy to be statistical rather than exact. We keep track of the energy

created or destroyed by this process. Including this energy allows us to calculate energy balance for the whole code, and we find that this is conserved to one part in 10^{10} in all the simulations presented in this paper.

Although the IMC particles will contribute to an accurate solution of the transport equation in regions of high effective scattering, they are still expensive to track in these regions. We might wish to convert IMC particles to IMD particles to reduce the cost of the calculation. We have experimented with two methods for converting IMC particles into IMD particles in diffusive regions.

The first is to convert IMC particles that reach census in diffusive regions into IMD particles in the next time step. That is, when these particles are brought out of census, they are converted to IMD particles rather than tracked as IMC particles.

The second is to convert IMC particles into IMD particles sometime after they cross into a diffusive region. We do not convert them immediately because this would represent a flux into the diffusive region, and we have imposed a vacuum boundary condition on the interface of the diffusive region. Also, many IMC particles entering the diffusive region will scatter back out. If we converted IMC particles into IMD particles immediately, many of the IMD particles would come out and have to be converted back into IMC particles. Multiple conversions from IMC to IMD and back are expensive. So we want to convert IMC particles to IMD particles only after we have given them a chance to scatter back out of the diffusive region. We do this by converting IMC particles only after they have scattered a certain number of times. We have found by experiment with several test problems that converting IMC particles in IMD regions to IMD particles after they take 50 IMC steps without leaving the diffusive region produces good results with a useful reduction in the run time.

We wish to emphasize that the conversion of IMC particles to IMD particles is not a necessary part of the hybrid IMC/IMD algorithm. It is an attempt to make the algorithm faster by reducing the number of expensive IMC steps taken in the diffusive region of problems. In particular, the conversion of IMC particles in diffusive regions after 50 steps is not an attempt to use the Nth collided intensity approximation [11], although it appears to be similar. The value of 50 scatters has been determined by experiment and may be problem dependent.

We show in Section 6 that the hybrid method without conversions, and with the two methods of conversion discussed here, produces essentially the same answer on test problems.

IMC particles can travel through all zones of the problem within a radius of $c\Delta t$ from their point of origin. IMD particles can, by converting into IMC particles, also travel throughout the problem. Thus, all zones in a problem which are causally connected can exchange particles, even if they are in diffusive regions separated by IMC regions. These regions are coupled together implicitly.

5. COMPARISON TO OTHER HYBRID METHODS

There are several existing hybrid methods which use some form of diffusion to accelerate radiative transfer problems for which IMC alone would be prohibitively expensive. We now analyze the differences and similarities between the new IMC/IMD hybrid method and these other methods.

The first method we consider is the most straightforward: solving the diffusion equation by a standard matrix solver rather than by using IMD. The probabilities used to advance

the IMD particles are derived from a matrix representation of the differenced diffusion equation. Thus, we could consider solving this matrix by employing an appropriate solver rather than using IMD. This technique is the one described in [6].

As mentioned previously, Monte Carlo solution techniques for matrix equations are more expensive than standard techniques[10]. Thus we would expect this approach to be faster than the IMC/IMD hybrid in many cases. However, this speed advantage may be lessened in some problems by the fact that the Monte Carlo method does not need to obtain the entire solution. If IMD particles do not penetrate all the zones in a problem, then no computation effort is expended there. IMD, in effect, only solves part of the matrix equation. For a problem in which only a small amount of the diffusive region is actually heated, we could see a speed advantage for IMD. Since the IMD method is particle based, it should be easier to parallelize than a hybrid method that would require a parallel matrix solver.

A drawback avoided by the IMD method is the possibility of negative fluxes from the diffusion into the IMC. This problem is discussed in [6].

Next we consider two hybrid methods that use local (i.e., zonal or subzonal) solutions of the diffusion equation to advance IMC particles in time more quickly than the IMC algorithm. These are the so-called Random Walk method of Fleck and Canfield [12] and the Discrete Diffusion method of Morel, Urbatsch, Evans, and others [13, 14].

The Random Walk method [12] proceeds as follows: the radius of the largest sphere centered on the particle which is wholly contained inside the particle's current zone is calculated. This radius is equal to the minimum distance to any zone boundary. If this radius is less than some number of mean free paths (usually on the order of 5–10), the particle is advanced by the standard IMC method. If not, the particle's behavior is considered to be well approximated by diffusion. In this case, the diffusion equation is solved inside this sphere, with a delta function in the center (the particle's location) as the initial condition and with the solution being zero on the sphere's surface serving as the boundary condition. This solution is obtainable analytically as an infinite series. It represents the probability distribution for particles behaving diffusively which have not left the sphere. Particles are advanced in time by sampling this distribution, and placed either inside the sphere or on the surface. The particle's energy is decreased, and new direction cosines are assigned. (A discussion of the methods for generating direction cosines is provided in [15] and [16].) The particle is in effect restored as an IMC particle. Then the process is repeated.

It has been found that the accuracy of simulations using Random Walk is often improved if particles are prevented from taking two Random Walk steps consecutively. In other words, an IMC step must be taken between each Random Walk step. This is thought to be because the direction cosines assigned to the particle after the Random Walk step are not isotropically distributed. This means that the diffusion approximation is not strictly applicable to the particle. Allowing it to take an IMC step allows a chance for the direction cosines to become randomized through a scatter. This restriction was employed in the Random Walk simulations described in the next section. Work on alleviating this restriction by sampling particle position and angle from a distribution that more accurately represents the results of transport in opaque media is described in [16].

The Random Walk method, like the IMD method, moves the particle in one jump to a new location. This jump replaces the many particle trajectories that would be taken if the particle was advanced by an IMC calculation. A Random Walk step is more expensive to

calculate than an IMD step. This is because the IMD step uses precalculated probabilities, while the Random Walk step needs to do a one-and two-dimensional table lookup to determine the new position of the particle. It also takes more Random Walk steps to move a particle out of a zone. This is because the new position of the particle is inside or on the surface of the sphere constructed around the particle, and this sphere is entirely within the zone. When the particle is near a zone boundary, the sphere will no longer have a radius greater than a few mean free paths, and so at least one, and possibly several, IMC steps will be taken before the particle crosses the zone boundary. IMD, by contrast, always moves the particle to a neighboring zone in one step. The IMD method would be expected to be faster than the Random Walk method because of the number of IMC steps taken by the Random Walk method in diffusive regions. In the next section, we run a test problem comparing the speed and accuracy of the Random Walk and IMD methods.

Discrete Diffusion Monte Carlo [13, 14] also uses a representation of the diffusion equation in a zone to advance a particle in time. Rather than using an analytic solution in a convenient subzonal volume, as the Random Walk method does, it uses an approximate solution over the whole zone. This approximate solution is obtained by assigning values of the radiation energy density to the zone center and edges, imposing Marshak boundary conditions at the edges, and manipulating the resulting equations to get expressions for the outgoing fluxes at the boundary. Using these and a representation of the diffusion equation integrated over the zone, an energy balance equation is obtained, which relates the outgoing fluxes, the source energy, and the absorption rate in the zone. This energy balance equation is solved for probabilities similar to the ones in Eqs. (32)–(35). Unlike those expressions, however, the ones produced by Discrete Diffusion Monte Carlo involve only the properties of the given zone; IMD probabilities involve quantities from neighboring zones also.

The Discrete Diffusion Monte Carlo method has, as of this writing, been applied to the time-independent diffusion equation [13] and the time-dependent equilibrium diffusion equation [14], and only in one dimension. As it is similar to IMD, it will probably produce similar results when applied to the nonequilibrium diffusion equation, Eq.(10). When this occurs, a meaningful comparison of the two methods will possible .

Lastly, we consider the hybrid method of N'Kaoua [17], which is based on the Symbolic IMC (SIMC) method of Brooks [18].

The IMC procedure evolved from an attempt to estimate the future matter temperature in the source term of the radiative transfer equation. Symbolic IMC [18], rather than estimating this term, regards it as an unknown. Monte Carlo particles are advanced, as in IMC, but have unknown weights; the weight is a fraction of the unknown t^{n+1} energy of the zone from which the particle is emitted. The result of advancing these particles is a matrix equation. The elements in the unknown vector are the t^{n+1} zone energies. The elements of the matrix are the net gains and losses obtained from emission and absorption of the particles. The size of the matrix is equal to the number of zones in the problem squared. It has nonzero elements connecting every pair of zones which have exchanged energy via particle transfer. This matrix equation must then be solved to obtain the new matter energies.

N'kaoua [17] developed a hybrid method using SIMC. It essentially replaces some of the elements of the matrix obtained by SIMC with elements derived from a discretization of the diffusion equation. This replacement is done in regions where the diffusion equation is an accurate approximation. This accelerates the SIMC method by replacing a Monte Carlo calculation of the matrix elements in regions where it is expensive. As in

the pure SIMC method, the matrix equation calculated by the hybrid method must still be solved.

SIMC, and the hybrid method of N'kaoua, are both truly implicit in the matter temperature. Thus both avoid the problems with stability and accuracy described in [4]. This comes at the cost of solving a matrix equation, which takes place after the Monte Carlo calculation. Since this matrix has nonzero elements connecting zones which exchange particles, it could conceivably be very dense where the mean free path of particles extends over many zones. Solution of a dense matrix of this size, especially in a three-dimensional problem, could be prohibitive. Although IMD employs the same estimate of future matter temperature as IMC, and is thus subject to the time step restrictions described in [4], it does not require solution of a matrix equation. This may make IMD more cost effective for many problems.

6. NUMERICAL RESULTS

We demonstrate that the IMC/IMD hybrid method produces a result essentially identical to the result of a standard matrix inversion method, differing only in that it has, like all Monte Carlo methods, some statistical noise. We then demonstrate the IMC/IMD hybrid method on three problems which possess both opaque and nonopaque regions, for which diffusion alone would be expected to produce an inaccurate answer, and for which IMC alone requires a prohibitively large amount of computational resources. We also compare the IMC/IMD hybrid method to the Random Walk method.

The first problem is the test problem from [7]. This test problem consists of a slab of material with mass density $\rho = 1$, specific heat $c_v = 1$, and opacity $\sigma_a = 1$ extending from $x = 0$ to $x = 10$. A temperature source with $T_{source} = \sqrt{2}$ at $x = 0$ generates a Marshak wave which penetrates into the initially cold slab, which has an initial temperature of $T_{initial} = 10^{-5/2}$. The units were chosen such that $c = a = k = 1$.

We compare the output of a pure IMD calculation to the results of a simulation using the difference equations given in [7] and solved with a tridiagonal solver as in [9]. This is done to demonstrate that the IMD method of solution produces the same result as the application of a tridiagonal solver to the discretized diffusion equation given by Eqs. (23) and (24). (As noted above, no flux limiter was used in either simulation.)

Figure 1 shows the radiation temperature at three different times during the calculation, as calculated with IMD and with the method given in [7]. Figure 2 shows the matter temperature. These figures plot the same variables as Figs. 1 and 2 in [7].

The two methods produce almost identical results. In this calculation, IMD was run with 10,000 particles. The IMD calculation took considerably longer than the calculation done with the traditional finite difference method. (Both calculations took one minute or less on a 533 MHz DEC alpha machine.) This is not surprising, as the tridiagonal matrix that the difference method in [7] uses may rapidly be inverted. Monte Carlo matrix inversions are usually much slower than traditional methods of inverting a matrix [10]. We are interested in the speed-up we can attain over IMC calculations on problems for which diffusion cannot be employed everywhere. We now examine two such problems, comparing pure IMC to the IMC/IMD hybrid method.

The first hybrid problem is a one-dimensional Cartesian problem in which a region with $\sigma_a = .1$ abuts a region with $\sigma_a = 100$. The low-opacity region extends from $x = 5$ to $x = 10$ and is heated by a temperature source at $x = 10$ with a constant temperature of $T_{source} = 1.0$. The high-opacity region extends from $x = 0$ to $x = 5$. A vacuum boundary condition is

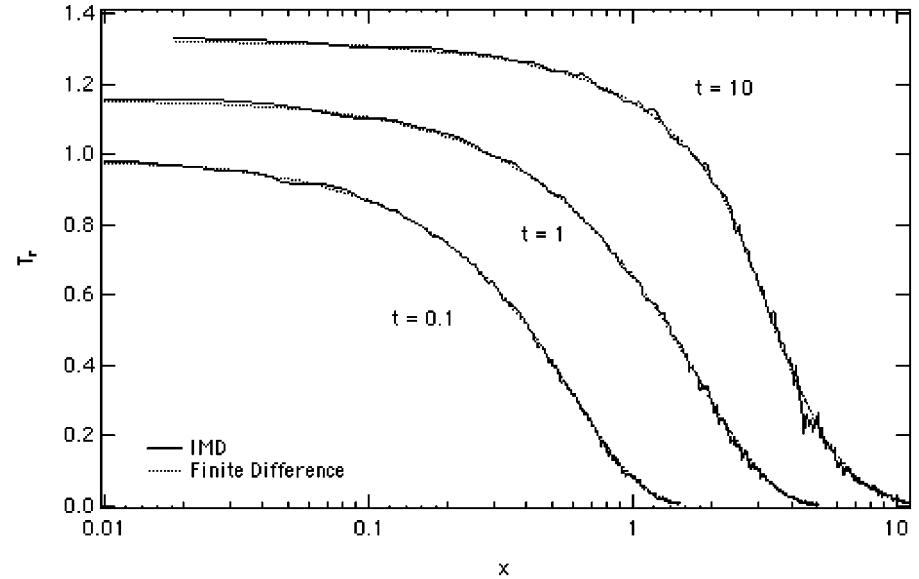

FIG. 1. Radiation temperature T_r vs. x for $t = 0.1$, 1.0, and 10 for the first test problem. A surface source with $T = \sqrt{2}$ at $x = 0$ drives a Marshak wave in the direction of increasing x. Results from IMD are compared to the results of a finite difference diffusion calculation. The two methods produce similar results.

imposed at $x = 0$. The initial temperature in both materials was $T_{initial} = 10^{-3}$, and the heat capacity was $48T^3$. The units were chosen such that $c = a = k = 1$.

In the high-opacity region, the scattering fraction $= 1 - f = .988$ (independent of temperature), so IMC particles in that region have a very small scattering mean free path. (This

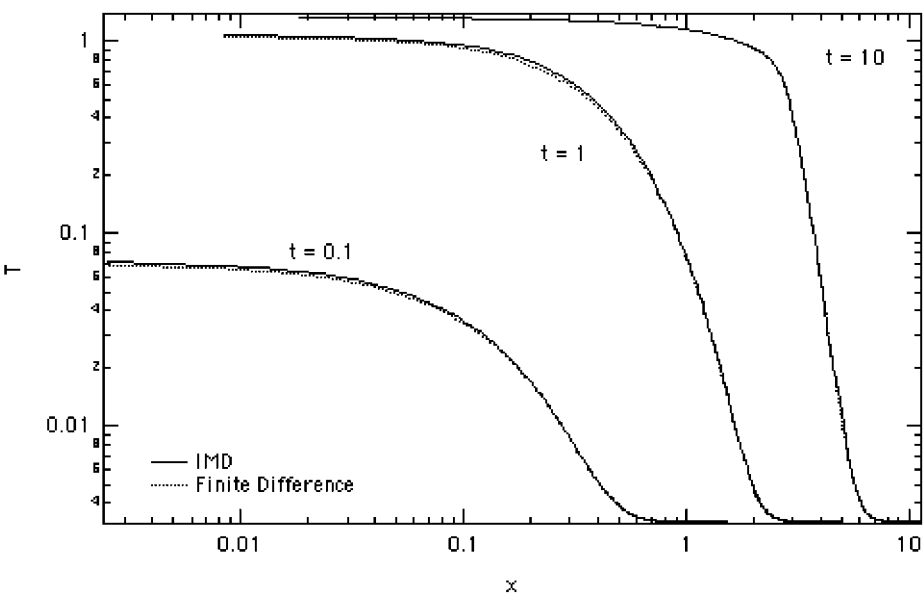

FIG. 2. Matter temperature T vs. x for $t = 0.1$, 1.0, and 10 for the first test problem. A surface source with $T = \sqrt{2}$ at $x = 0$ drives a Marshak wave in the direction of increasing x. Results from IMD are compared to the results of a finite difference diffusion calculation. The two methods produce similar results.

value of f results from using a time step of $\Delta t = 10$). We would expect that IMC would be very expensive in the high-opacity region and that diffusion would be inaccurate in the low-opacity region.

The problem has an approximate steady state solution. In the steady state, $E_m = E = aT^4$. The temperature in the low-opacity region will approach $T_{source} = 1$. In the high-opacity region, the matter and radiation energy densities will be equal and will satisfy the time-independent diffusion equation

$$-\frac{\partial}{\partial x}\left[D\frac{\partial E}{\partial x}\right] = 0 \tag{45}$$

on $0 < x < 5$, with $E = 1$ at $x = 5$ and $D = 1/300$. A vacuum boundary condition is imposed at $x = 0$:

$$E_{x=0} + 2D\frac{\partial E}{\partial x}\bigg|_{x=0} = 0. \tag{46}$$

This yields the solution

$$E(x) = \frac{x + 2D}{5 + 2D} \tag{47}$$

for $0 < x < 5$. We can get the equilibrium value of temperature from $T(x) = E(x)^{.25}/a$.

The problem was run with 110 zones, with zoning of $\Delta x = 1$ for $0 < x < 4$ and $5 < x < 10$. In $4 < x < 5$ there were 20 zones with geometrically decreasing size, with the largest zone spanning the region $x = 4$ to $x = 4.0928033856$. The time step $\Delta t = 10$, and the problem was run to $t = 50{,}000$. This was considered to be steady state because the solution at $t = 100{,}000$ was essentially the same as at $t = 50{,}000$. The IMC calculation was run with 25,000 particles per time step. In the hybrid simulation, the sum of the number of IMC and IMD particles was 25,000 in each time step.

Figure 3 shows matter temperature for a simulation run with pure IMC, and with a hybrid method with IMC particles run everywhere and IMD particles run in the high-opacity region $0 < x < 5$. Both methods show good agreement with the approximate analytic answer, which is also shown.

The pure IMC calculation took approximately 1.057×10^5 s, about 29.4 h. The hybrid method took approximately 3800 s. (1.06 h), giving a speed-up of approximately a factor of 27.8. Both calculations were run on a 533 MHz DEC alpha chip.

Figure 4 shows a plot of the number of particle steps taken by both methods. The pure IMC method takes most of its steps in the high-opacity region, which spans zones 1 to 60. In the hybrid calculation, many fewer IMC steps were taken in the high-opacity region. Since each IMC step takes about the same amount of computer time, and the IMC steps are more expensive than the IMD steps, the time of each simulation is approximately proportional to the area under the curve representing the number of IMC steps taken. Thus the speed-up of almost 28 times is understandable as a result of the fact that so many fewer IMC steps were taken in the high-opacity material.

The statistical behavior of the hybrid method is examined in Fig. 5. Fig. 5 shows the error in the temperature, and the variance in the temperature at three different spatial locations in the problem, at the final time $t = 50{,}000$. The three locations are zone 1, at $x = 0$, at the vacuum boundary in the diffusive region; zone 60, at $x = 5$, the diffusive zone which abuts

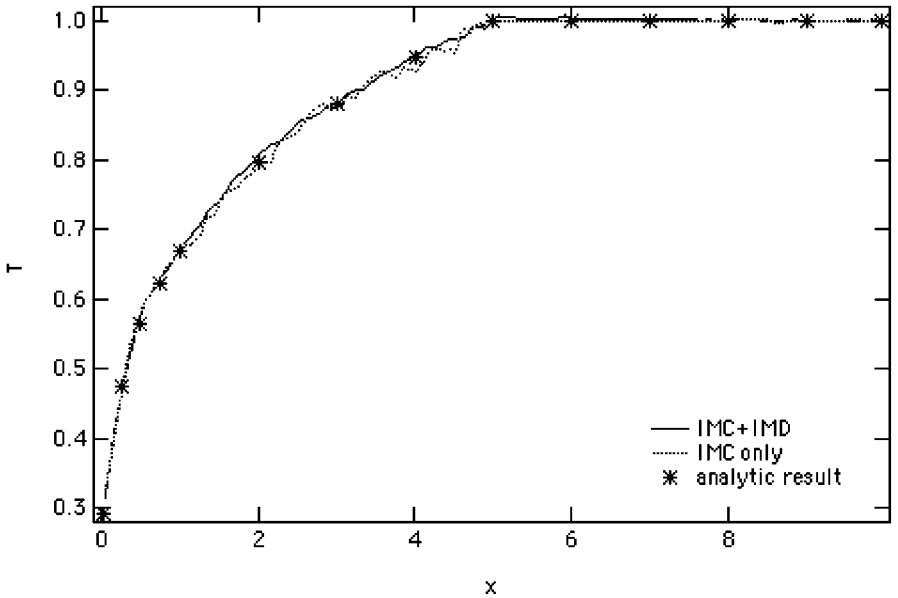

FIG. 3. Matter temperature T vs. x at $t = 50{,}000$ for the first hybrid test problem. A surface source with $T = 1$ at $x = 10$ drives a Marshak wave in the direction of decreasing x. Results from a calculation using only IMC are compared to results obtained from the hybrid method. Both calculations were done with 25,000 total particles. The approximate analytic answer at several points is also shown. The two methods produce similar results and agree well with the approximate analytic results.

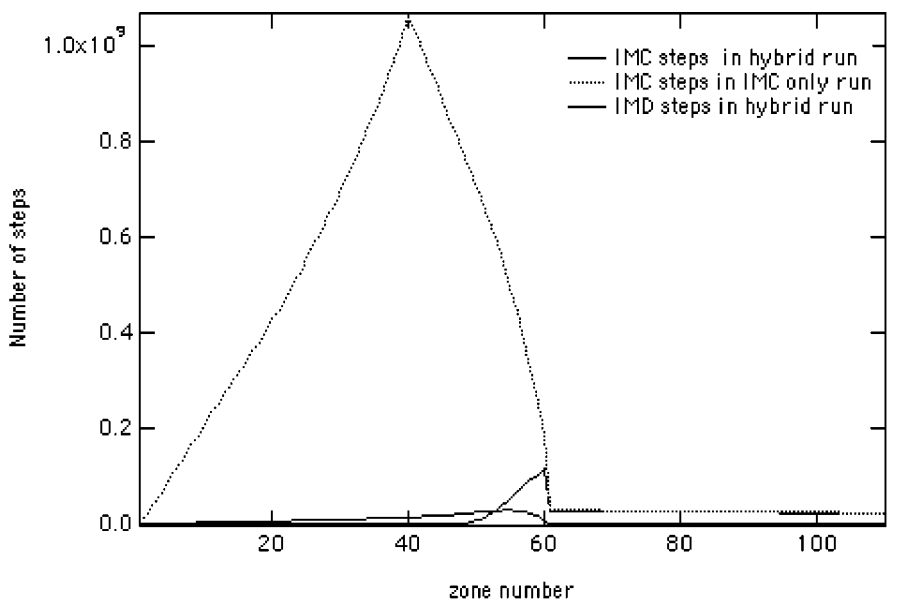

FIG. 4. Number of particle steps vs. zone number for the first hybrid test problem. The number of steps taken by IMC particles, which are the most expensive part of the calculation, is much smaller when the hybrid method is used. Both calculations were done with 25,000 total particles.

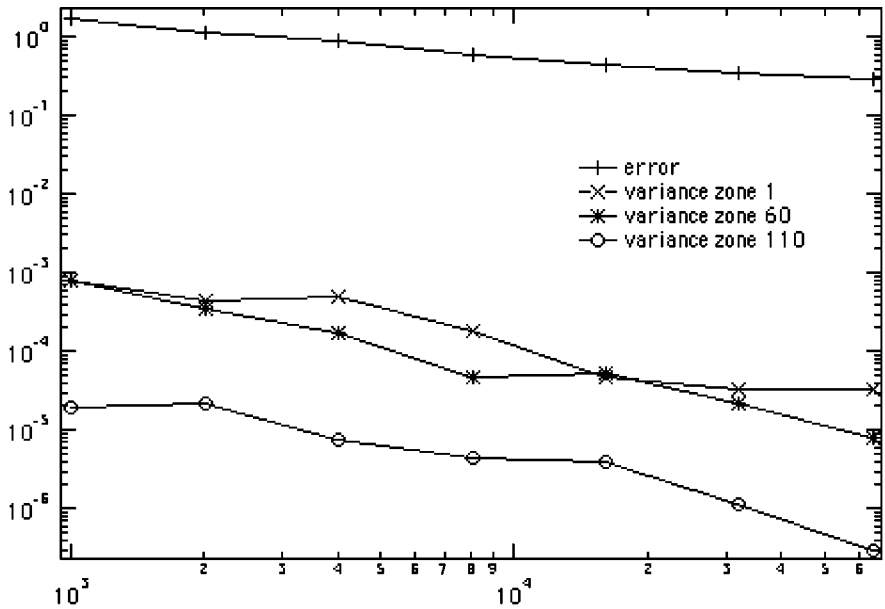

FIG. 5. Total error, and variance at three different x locations, plotted vs. number of particles in the first hybrid test problem. The error should be proportional to $1/\sqrt{N}$, and the variance proportional to $1/N$, where N is the number of particles. Curve fits show approximate agreement with this expected behavior. The error declines as $N^{-.43}$; the variance of zone 1 like $N^{-0.88}$; for zone 60, like $N^{-1.06}$; for zone 110, like $N^{-0.99}$. The overall value of the variance for zone 110 is lower than that for the other zones because there are more particles in that zone since it is near the temperature source.

the thin material; and zone 110, the zone in the nondiffusive region which abuts the temperature source.

These quantities are plotted against the number of particles used in the calculation. The particle number used in these calculations began at 1000 and was increased by a factor of two, with the final value being 64,000. For each different number of particles, 10 runs were done. The error was averaged over these 10 runs, and the variance in T at the three zones was calculated. The error was approximated as the absolute value of the difference between the value of T and the approximate analytic value in every zone in the problem. This error was averaged over the 10 runs.

The error should decline as $N^{-1/2}$, where N is the number of particles. A fit of the error curve in Fig. 5 shows that the error declines as $N^{-.43}$. This value is slightly larger than expected, probably due to the fact that there is a residual error in the calculation due to zoning. That is, we cannot expect the error to decline like $N^{-1/2}$ for very large N since there is a minimum discretization error. The variance should decline like N^{-1}. For zone 1, it declines like $N^{-0.88}$; for zone 60, like $N^{-1.06}$; for zone 110, like $N^{-0.99}$. These results are consistent with the statistical behavior we expect from a Monte Carlo simulation. Thus, the IMC/IMD hybrid method exhibits the behavior we expect from a Monte Carlo simulation.

In Fig. 6, we examine the effects of converting IMC particles to IMD particles in the diffusive region of the problem. Figure 6 compares the three different results of the one-dimensional hybrid test problem. In one case, IMC particles were never converted to IMD particles. This is the result discussed above and used in generating the data pictured in Figs. 3 and 4. In the second case, IMC particles reaching census in the diffusive region at

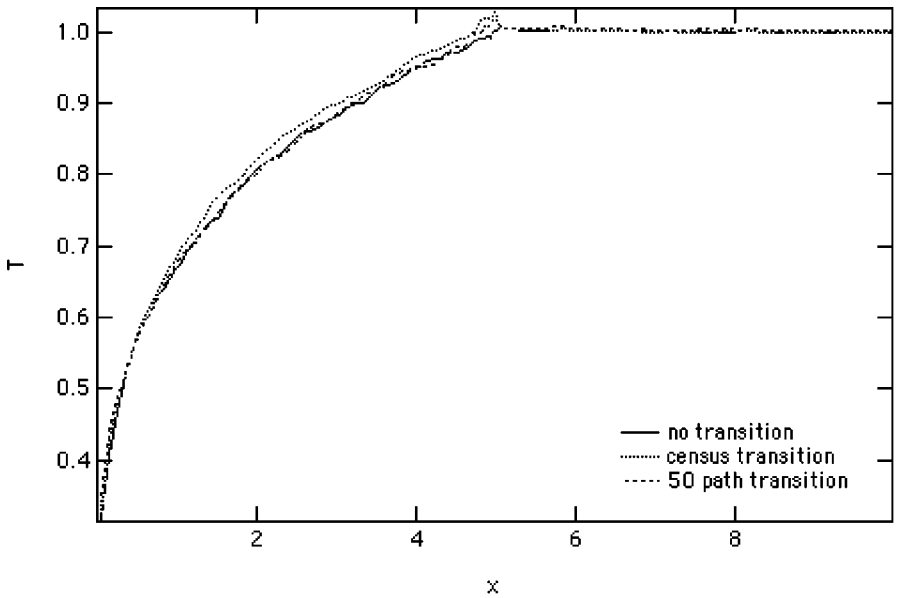

FIG. 6. Matter temperature T vs. x at $t = 50{,}000$ for the first hybrid test problem. A surface source with $T = 1$ at $x = 10$ drives a Marshak wave in the direction of decreasing x. Results from three calculations using the hybrid method are shown. In one, IMC particles in the diffusive region are tracked with IMC until they are killed off through Russian rouletting or leave the problem. In another, IMC particles reaching census in the diffusive region were converted to IMD particles in the next time step. In the third, IMC particles which scattered 50 times in the diffusive region were converted to IMD particles in the current time step. All calculations were done with 25,000 total particles. All methods produce similar results.

the end of a time step were converted to IMD particles at the beginning of the next time step. In the third case, IMC particles which scattered 50 times in the diffusive region without leaving it were converted to IMD particles.

Figure 6 shows that all three cases produce similar results. The methods that convert IMC particles to IMD particles complete the simulation faster. When IMC particles are not converted, the simulation takes about 3800 s, when census particles are converted, it takes about 1340 s, when IMC particles are converted after 50 scatters, it takes about 1100 s. This is easily understood, since most of the time in simulations is taken up by advancing IMC particles, and converting them to IMD particles means that fewer IMC steps are simulated. Converting IMC particles into IMD particles gives approximately another factor of 3 speed-up in the one-dimensional hybrid problem.

The second hybrid problem is used to compare the IMC/IMD hybrid acceleration method to the Fleck and Canfield Random Walk acceleration method [12]. The problem has every zone beginning at $T = 1$, with temperature sources with $T = 1$ at either end. This problem has a simple analytic answer: the temperature stays at $T = 1$ everywhere at all times. This problem is run on the same mesh as the first hybrid problem described above, and with the same units. The material properties are the same, except that the opaque material has an opacity of 500 rather than 100. This was done so that the opaque zones, most of which have a length of 0.1, are more than 5 mean free paths across. Because of the restriction that the size of the sphere the Random Walk algorithm uses be at least 5 mean free paths, Random Walk would never be invoked in the problem if the opacity remained 100. No sphere of this size could fit into the zones in the problem.

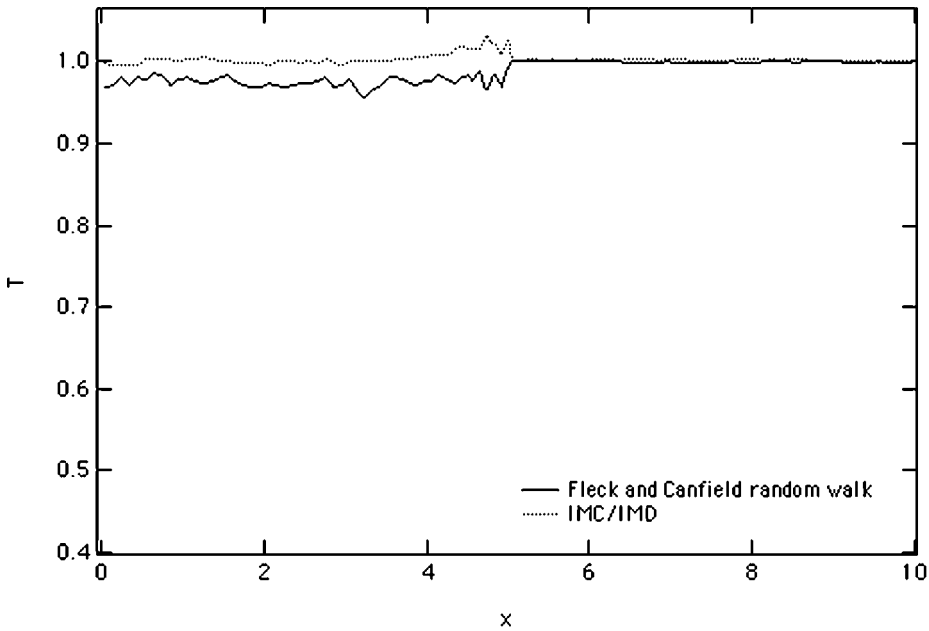

FIG. 7. Matter temperature T vs. x at $t = 10{,}000$ for the second hybrid test problem. Initial conditions are $T = 1$ everywhere, and surface sources with $T = 1$ at $x = 0$ and $x = 10$ maintain this steady state. Results from two calculations using the Random Walk method and the IMC/IMD hybrid method are shown. Both acceleration methods reproduce the steady state solution with reasonable accuracy. The Random Walk method took approximately 2.45×10^4 s, while the IMC/IMD method took 450 s.

The problem was run with both the IMC/IMD hybrid algorithm and with IMC using the Random Walk algorithm. The size of the spheres used in the Random Walk algorithm was required to be 5 mean free paths or larger. Both simulations used 50,000 particles and were run to $t = 10{,}000$ with $\Delta t = 10$. Running this problem with IMC alone would have taken a prohibitively large amount of computer time, so this was not done.

Figure 7 shows the matter temperature at the end of the simulation. Both methods produce a solution that agrees reasonably well with the analytic solution. The calculation using Random Walk took about 2.54×10^5 s (about 6.8 h). The IMC/IMD hybrid took about 450 s, providing a speed-up of approximately 54.

The IMC/IMD hybrid algorithm is much faster than the Random Walk method on this problem for two reasons. First, many more IMC steps are taken in the opaque material when Random Walk is used than when IMD is used. This is because IMC steps are taken when the particle is near the boundary of a zone, and because we have required that an IMC path be taken between consecutive IMD steps. IMD particles quickly replace IMD particles in the opaque regions and very few IMC steps are taken there. Second, more Random Walk steps than IMD steps are taken in the opaque material. This is because a Random Walk path will terminate in or on the sphere centered on the particle and will not get the particle out of the zone. Many Random Walk and IMC steps are required to get a particle to leave a zone. Each IMD path moves the particle to a neighboring zone. This effect on the relative run times of the methods is enhanced by the fact that IMD steps are less expensive than Random Walk steps, which require sampling from a table.

The reason the IMC/IMD algorithm is so much faster than Random Walk on this problem is illustrated in Fig. 8. This figure shows the number of steps taken by particles in the two

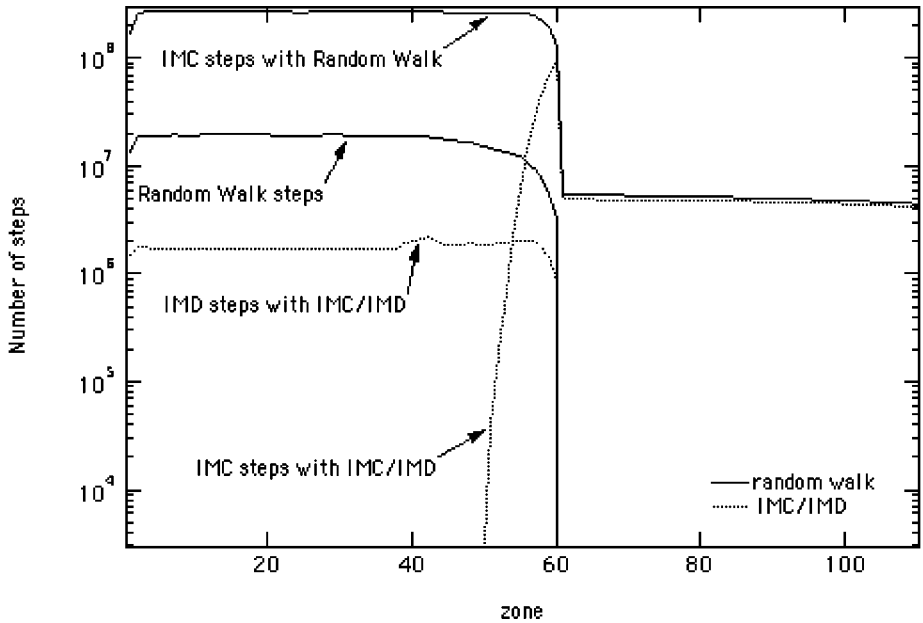

FIG. 8. Number of steps taken by particles using the different acceleration methods in the two simulations of the problem pictured in Fig. 7. For the Random Walk simulation, the number of IMC steps and the number of Random Walk steps are shown as solid lines. For the IMC/IMD hybrid simulation, the number of IMC steps and IMD steps taken are shown as dotted lines. Many more IMC steps are taken in the opaque material when Random Walk is used. The number of Random Walk steps taken is much higher than the number of IMD steps.

simulations. For the Random Walk simulation, it shows the number of IMC steps and the number of Random Walk steps. For the IMC/IMD hybrid simulation, it shows the number of IMC steps and IMD steps taken. Many more IMC steps are taken in the opaque material when Random Walk is used, and the number of Random Walk steps taken is much higher than the number of IMD steps.

The third hybrid problem is a two-dimensional problem developed by Frank Graziani and Jim Le Blanc often referred to as the tophat problem [19]. In this problem a cylindrical section of dense, opaque material is embedded in less dense, less opaque material, which is itself embedded in a cylinder of the dense, opaque material. The outer region extends from $z = 0$ to $z = 7$, and from $r = 0$ to $r = 2$. Inside of this, the inner region of dense material extends from $z = 3$ to $z = 4$ and from $r = 0$ to $r = 1$, and the region of less dense material extends from $z = 0$ to $z = 2.5$, and from $z = 4.5$ to $z = 7$ with a radius of 0.5, and from $z = 2.5$ to $z = 4.5$ with a radius of 1.5.

The dense material has a mass density $\rho = 10$ g/cm^3 and an opacity $\sigma_a = 2000\,\text{cm}^{-1}$. The less dense material has $\rho = 0.01$ g/cm^3 and $\sigma_a = 0.2\,\text{cm}^{-1}$. Both materials have a heat capacity $c_v = 10^{15}$ erg/g $-$ keV.

A temperature source with $T_{source} = 0.5$ keV is applied at one end of the less dense material. Vacuum boundary conditions are applied at all other boundaries. The initial temperature $T_{initial} = 0.05$ keV everywhere. Time units are such that $c = 300$. The time step begins at $\Delta t = 10^{-3}$ and is increased by a factor of 1.1 each time step until it reaches 1.0, and is held constant at 1.0 thereafter.

This problem cannot be run accurately with diffusion alone because diffusion cannot accurately model the flow of radiation around the corners. There is no analytic answer for the

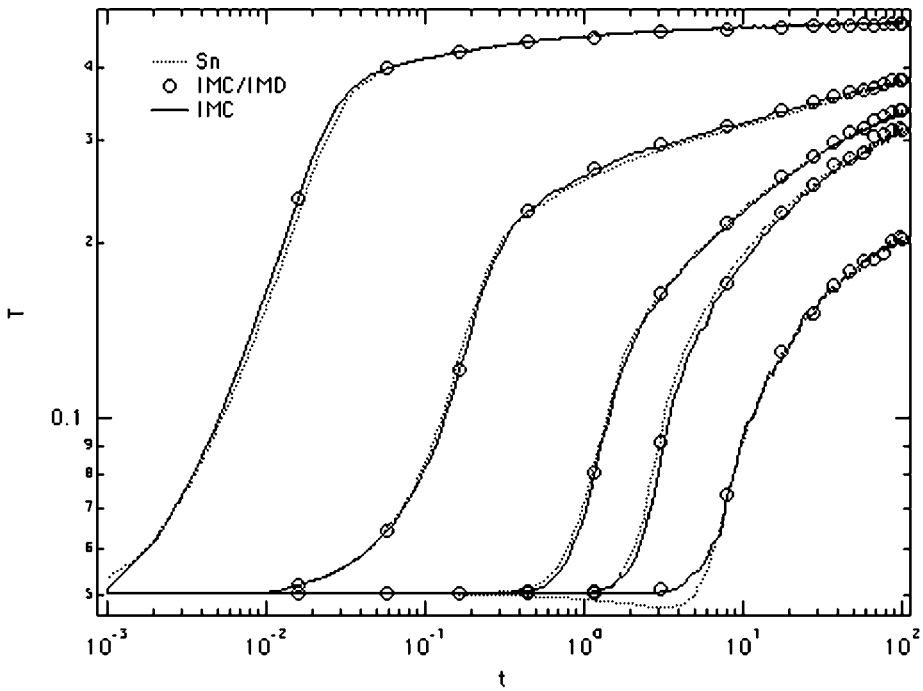

FIG. 9. T vs. t for the five fiducial points located in the thin material in the tophat problem. This plot compares results from pure IMC, the IMC/IMD hybrid method, and an Sn calculation. The time step for the IMC and the hybrid calculation was the same, but only certain times for the hybrid calculation have been plotted to make the plot clearer. All three codes produce similar results. (The difference between Sn and IMC for the last point to heat up is due to a lack of particles in the cold region of the problem. This lack of particles prevented this zone from cooling off slightly before the thermal wave reached it.)

problem. The usual figures of merit are the matter temperature at five points in the optically thin material. These points are: $(r = 0, z = 0.25)$, $(r = 0, z = 2.75)$, $(r = 1.25, z = 3.5)$, $(r = 0, z = 4.25)$, and $(r = 0, z = 6.75)$.

As time proceeds in this problem, radiation travels down the thin material and around the section of thick material in the middle of the problem, and the thick material slowly heats up. It is in the heated parts of the thick regions that IMC will be most computationally intensive.

In Fig. 9 we have plotted matter temperature at the five fiducial points from the hybrid simulation and the pure IMC calculation, as well as an Sn solution [20]. In this calculation, we have converted IMC particles which take 50 steps in the opaque region to IMC particles, as described above. Both Monte Carlo calculations had 500,000 total particles per time step. We see good agreement between all three simulations for all the points at most times. Not all points for the hybrid simulation are shown, in order to make the plot less cluttered.

The only notable area of disagreement in this plot is in the behavior of the point farthest from the heat source. In the Sn simulation there is a small decrease in the temperature at this point before the radiation wave begins to heat it. This does not occur in either Monte Carlo simulation. This disagreement occurred because the IMC and hybrid calculations were run with a variance reduction technique; particles were not allocated to zones that were still at the initial temperature. The zone farthest from the heat source, which is near

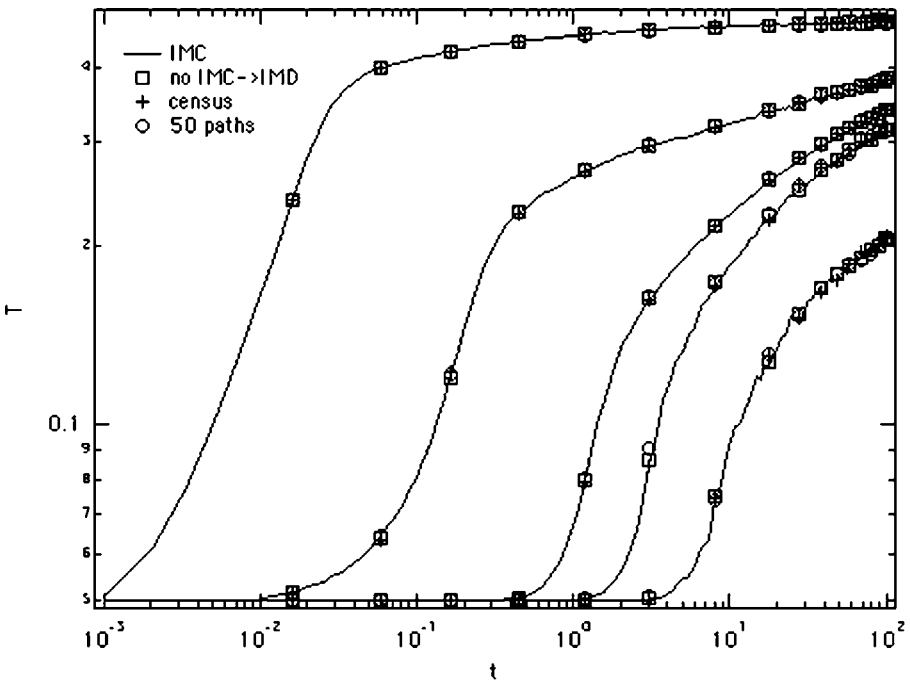

FIG. 10. T vs. t for the five fiducial points located in the thin material in the tophat problem. This plot compares results from pure IMC to the IMC/IMD hybrid method. Results from three calculations using the hybrid method are shown. In one, IMC particles in the diffusive region are tracked with IMC until they are killed off through Russian rouletting or leave the problem. In another, IMC particles reaching census in the diffusive region were converted to IMD particles in the next time step. In the third, IMC particles which scattered 50 times in the diffusive region were converted to IMD particles in the current time step. All calculations were done with 500,000 total particles. All methods produce similar results.

the end of the problem, could actually cool off slightly before the radiation wave reached it. The Sn calculation captured this transient, while the Monte Carlo calculations did not, although they could have if that zone had been allowed to radiate. As soon as the radiation reaches that zone, the effect of the cooling is quickly rendered unnoticeable, and all three simulations show good agreement for this point thereafter.

In Fig. 10, we examine the effects of converting IMC particles to IMD particles in the diffusive region of the problem. In one case, IMC particles were never converted to IMD particles. In the second case, IMC particles reaching census in the diffusive region at the end of a time step were converted to IMD particles at the beginning of the next time step. In the third case, IMC particles which scattered 50 times in the diffusive region without leaving it were converted to IMD particles. The three cases are compared to a pure IMC calculation. Figure 10 shows that all three cases produce similar results.

At $t = 100$, the hybrid calculation where IMC particles were converted to IMD particles after 50 steps took about 2.42^4 s (about 6.7 h). The hybrid calculation where IMC particles reaching census in the opaque region were converted took approximately 1.83×10^5 s (about 51 h). With no conversion, the calculation took approximately 1.86×10^5 s (about 52 h). The pure IMC took about 5.73×10^5 s (about 159 h). Depending on which conversion method is used, the hybrid method is faster by a factor of from 3 to almost 24. All simulations were run on a 533 MHz DEC alpha.

The reason that the conversion method affects the run time of this problem more than the one-dimensional one is due to the energetics of the problem. The energy input from the external temperature source and the radiating opaque matter during each time step is large compared to the amount of energy in radiation. Many census particles are Russian rouletted when they reach census, so that particles may be apportioned to represent the energy from the sources at the beginning of the time step. This means that much of the run time is spent advancing IMC particles in the opaque material to the end of the time step in which they entered the thick material. So converting the relatively few particles which survive the Russian rouletting to IMD particles in the next time step does not provide a significant reduction in run time. Converting these IMC particles to IMD particles after 50 steps provides a significant decrease in run time because it converts the IMC particles which take the most time to simulate into the less expensive IMD particles.

Figure 11 shows a contour plot of the matter temperature in the problem at $t = 10$. The top half shows the hybrid calculation (with conversion after 50 steps), while the bottom shows the reflected results of the pure IMC calculation. (Note that these are two separate calculations, whose results have been placed on the same plot by mapping the pure IMC results to "negative" radius.) The contour lines are ten equally spaced lines between $T = 0.05$ and $T = 0.5$ keV. The hybrid calculation shows less penetration of the thermal wave into the thick material, but the agreement in the thin material is very good, and the overall accuracy seems adequate.

We believe the smaller amount of heating in the thick material results from the fact that the diffusion approximation is not accurate in the cold, thick material, because the scattering opacity is small while the material is cold. The factor f given by Eq. (4) is approximately

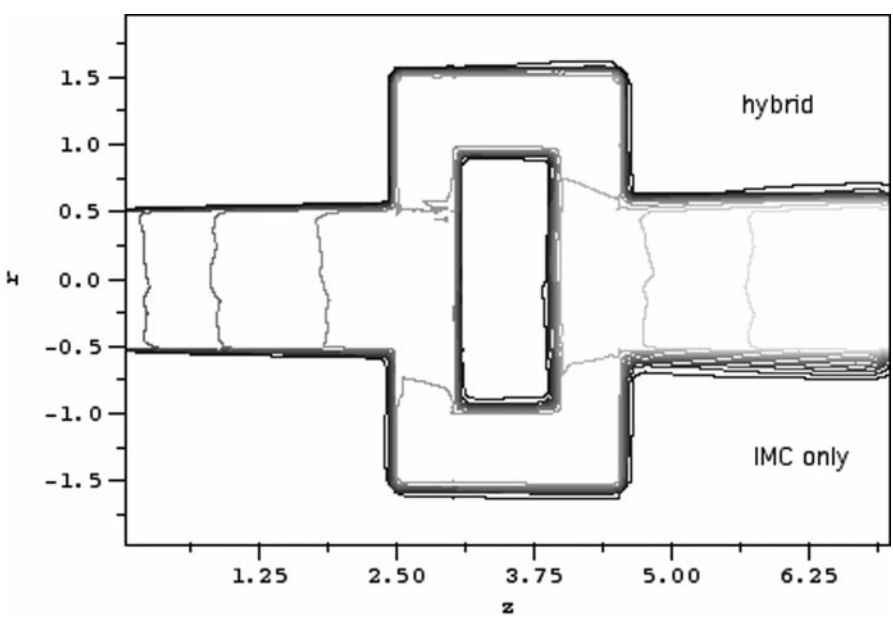

FIG. 11. T at $t = 100$ for tophat problem. This plot shows matter temperature for every zone in the problem. It compares results from pure IMC to the IMC/IMD hybrid. IMC/IMD hybrid results are shown on the top part of the plot. Pure IMC results were mapped to "negative r" and are shown on the bottom part of the plot. Contour lines are equally spaced between 0.05 and 0.5. The hybrid method underpredicts heating where the wave enters the thick material but shows good agreement in the thin material.

unity until the material heats up. A more accurate way of using the hybrid algorithm would be to use IMD only in zones with small f, rather than in all zones with a large opacity.

7. SUMMARY

We have developed a method for increasing the speed of IMC radiation transport calculations which contain regions where the opacity is large. This method employs a Monte Carlo method of matrix inversion to solve the diffusion equation. We have dubbed the method Implicit Monte Carlo Diffusion, which we abbreviate IMD. The IMD method employs particles which resemble the particles used in the IMC radiation transport method. Because of this resemblance, the two methods can easily be combined into a hybrid method, in which IMD is used in optically thick parts of the problem, where the diffusion approximation provides an accurate approximation to the results of the transport equation. Since the IMD method is considerably faster than the IMC method in these regions, the hybrid method can run much faster than IMC alone on some problems. We have demonstrated this method on one-dimensional Cartesian and two-dimensional cylindrical orthogonal mesh calculations. Test problems on these meshes show that the hybrid method is accurate and can provide speed-ups greater than an order of magnitude. Future work will include extending the method to 3D unstructured grids and to problems with frequency-dependent opacities.

ACKNOWLEDGMENTS

This work was performed under the auspices of the U.S. Department of Energy by the University of California Lawrence Livermore National Laboratory under Contract W-7405-Eng-48. The author thanks George Zimmerman, Todd Palmer, Jim Rathkopf, Mal Kalos, Abe Szoke, and Laura Burgardt for interesting and useful discussions. Special thanks go to Marc Cooper, who wrote the code implementing the Fleck and Canfield algorithm used in this paper.

REFERENCES

1. G. C. Pomraning, Equations of radiation hydrodynamics, in *International Series of Monographs in Natural Philosophy*, edited by D. ter Harr (Pergamon, Elmsfords, N. Y. 1973), Vol. 54.
2. J. A. Fleck, Jr., The calculation of nonlinear radiation transport by a Monte Carlo method, in *Computational Methods in the Physical Sciences*, edited by B. Alder and S. Fernbach, (McGraw-Hill, New York, 1963) Vol. 1, p. 43.
3. J. A. Fleck, Jr. and J. D. Cummings, *J. Comput. Phys.* **8**, 313 (1971).
4. E. W. Larsen and B. Mercier, *J. Comput. Phys.* **83**, 433 (1989).
5. E. W. Larsen, *Ann. Nucl. Ener.* **7**, 249 (1980).
6. G. C. Pomraning and G. M. Foglesong, *J. Comput. Phys.* **32**, 420 (1979).
7. R. H. Szilard and G. C. Pomraning, *Nucl. Sci. Eng.* **112**, 256 (1992).
8. G. L. Olson, L. H. Auer, and M. L. Hall, *J. Quant. Spectros. Radiat. Transfer* **64**, 619 (2000).
9. W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes*, 2nd ed. (Cambridge Univ. Press, New York, 1992).
10. B. R. Eccleston and T. S. Palmer, *Trans. Am. Nucl. Soc.* **81**, 148 (1999).
11. E. E. Lewis and W. F. Miller, Jr., *Computational Methods of Neutron Transport* (Am. Nucl. Soc. Illinois, La Grange Park, IL, 1993).
12. J. A. Fleck, Jr. and E. H. Canfield, *J. Comput. Phys.* **54**, 508 (1984).

13. T. J. Urbatsch, J. E. Morel, and J. C. Gulick, *Proceedings of the ANS Conference on Mathematics and Computation, Reactor Physics, and Environmental Analysis in Nuclear Applications, Madrid*, 1999, pp. 262–271.
14. T. M. Evans, T. J. Urbatsch, and H. Lichtenstein, *Proceedings of the Monte Carlo 2000 International Conference, Lisbon*, 2000.
15. M. A. Cooper, *UCRL-ID-139692* (2000).
16. W. B. Bateson, *UCRL-JC-137551* (2000).
17. T. N'kaoua, *J. Sci. Stat. Comput.* **12**, 505 (1991).
18. E. D. Brooks III, *J. Comput. Phys.* **83**, 433 (1989).
19. F. Graziani and J. LeBlanc, *UCRL-MI-143393* (2000).
20. P. Nowak and M. Nemanic, *Proceedings of the ANS Conference on Mathematics and Computation, Reactor Physics, and Environmental Analysis in Nuclear Applications, Madrid*, 1999, pp. 379–390.