Joint International Topical Meeting on Mathematics & Computation and Supercomputing in Nuclear Applications (M&C + SNA 2007)
Monterey, California, April 15-19, 2007, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2007)

IMPROVED IMPLICIT MONTE CARLO SCHEMES BASED ON THE DIFFERENCE FORMULATION

Richard P. Smedley-Stevenson
AWE
Aldermaston, Reading, Berkshire, RG7 4PR, UK
richard.smedley-stevenson@awe.co.uk

ABSTRACT

This paper describes slab geometry studies of the difference formulation based on the original P_0 formulation and a new P_1 formulation. Hybrid schemes are presented, which combine the best features of the intensity and difference formulations. The aim is to design an optimal algorithm for general Monte Carlo radiation transport problems, suitable for use with both IMC and SIMC temporal discretizations. The pseudo-scattering term in the IMC temporal discretisation is identified as a major source of inefficiency in the difference formulation, making it less efficient than the original intensity formulation. By contrast, the SIMC hybrid P_1 difference formulation offers a performance improvement of 3−4 orders of magnitude over the original SIMC method for an example of a deep penetration problem.

Key Words: IMC, SIMC, difference formulation

1. INTRODUCTION

The Implicit Monte Carlo (IMC) method for thermal radiation transport proposed by Fleck and Cummings [1] has been highly successful for solving complex thermal radiation transport problems such as modeling the energetic behavior of laser Hohlraums. However, it is important to be able improve the efficiency of the tracking scheme in opaque materials due to the random walk behavior exhibited by the X-ray photons; a dominant fraction of the absorption and re-emission events are replaced by discrete scattering events in order to improve the temporal stability of the Monte Carlo transport scheme. A related problem occurs for the Symbolic Implicit Monte Carlo (SIMC) scheme [2], where a large number of small weight particles are required to accurately model the energy flow in opaque materials due to the short lifetimes of the individual particles.

The original mechanism suggested for overcoming this difficulty with the IMC method is to replace the explicit modeling of the particle tracks by statistical sampling of the random walk process associated with particle propagation [3]. This process is not without its difficulties. Multi-frequency effects are not necessarily accounted for correctly when sampling the random walk trajectories. Furthermore, the need to inscribe a sphere of 10–20 mean free paths radius within the cell in order to sample the particle position means that the algorithm cannot be used for particles approaching the cell boundaries; this limits the usefulness of the scheme for problems with large aspect ratio cells/highly resolved opaque regions.

Alternatively, the Monte Carlo solution can also be coupled to the solution of the corresponding diffusion equation in the opaque material, the key issue being how to choose the boundary between the diffusion and transport solution regions to obtain the correct physical behavior. This diffusion problem can be solved deterministically as in the scheme proposed by N'kaoua [4] and in an updated form suitable for multi-frequency application by Clouët and Samba [5]. Alternatively, particles can be tracked according to a discretized diffusion equation [6]. More accurate treatments for coupling transport and diffusion regions in Monte Carlo transport calculations have recently been reported [7] [8].

The difference formulation offers a promising alternative acceleration procedure for both IMC and SIMC methods, by reducing the number of particles require to model the radiation field accurately in opaque regions without compromising the integrity of the underlying transport algorithm. There are similarities with the reduced source method proposed by Fraley et al. [9]. In the reduced source method the offset function is obtained by solving the P_1 equations and the continuum P_1 equations are used to obtain a transport equation for the residual correction to the radiation field. The accuracy of this solution is therefore compromised by the discretisation errors arising from the approximate solution of the P_1 equations. By contrast, the difference formulation maintains the accuracy of the original transport scheme by subtracting an algebraic expression for the radiation field corresponding to the equilibrium diffusion approximation.

2. TEMPORAL DISCRETIZATION

We begin by describing the equations for grey thermal radiation transport equation in a purely absorbing medium. The radiation transport equation can be written as

$$\left[\frac{1}{c}\frac{\partial}{\partial t} + \underline{\Omega}\cdot\nabla + \sigma_a'\right] I\left(\underline{\Omega},\underline{r},t\right) = \sigma_a' B\left(T\left(\underline{r},t\right)\right) \tag{1}$$

and this must be solved concurrently with the material energy equation,

$$\rho C_V \frac{\partial}{\partial t} T\left(\underline{r},t\right) = \sigma_a' \left[\tfrac{1}{4\pi}\int_{4\pi} I\left(\underline{\Omega},\underline{r},t\right) d\Omega - B\left(T\left(\underline{r},t\right)\right)\right] \tag{2}$$

due to the stiff coupling associated with the radiation-material energy exchange. This set of equations is in general non-linear except for the special case of $C_V \propto T\left(\underline{r},t\right)^3$. For simplicity we assume that both the specific heat capacity C_V and the absorption opacity σ_a' are independent of temperature. In general the opacities are assumed to be piecewise constant within each element in order to simplify the Monte Carlo calculation.

The fundamental unknowns are $I\left(\underline{\Omega},\underline{r},t\right)$ the radiation intensity which will be modeled using Monte Carlo particles, $T\left(\underline{r},t\right)$ the material temperature distribution and $B\left(T\left(\underline{r},t\right)\right) = aT\left(\underline{r},t\right)^4$, the blackbody equilibrium radiation energy density which are modeled deterministically. We assume that the local values of T and B are consistent at the discretisation points, but allow each variable to be expanded separately in terms of the basis functions between the discretisation points; a continuous analogue of this scheme is described in [10]. This considerably simplifies our algorithm relative to a consistent treatment such as that developed by Brooks et al. [11] for linear elements.

2.1. Fleck and Cummings IMC

In this paper we consider two different strategies for discretizing this coupled set of equations in time. The Fleck and Cumming implicit Monte Carlo (IMC) method is based on linearizing the material energy equation and using this to eliminate the end of time-step material temperature from the transport equation. The thermal emission source is reduced in order to overcome the stability constraints associated with treating this term explicitly. The remaining fraction of the emission events are modeled by introducing an additional scattering term which accounts for the absorption and re-emission events without any net energy exchange.
For the IMC method the grey transport equation can be written as

$$\left[\frac{1}{c}\frac{\partial}{\partial t}+\underline{\Omega}\cdot\nabla+\sigma_a'\right]I(\underline{\Omega},\underline{r},t)=f\sigma_a'B(T(\underline{r},t))+(1-f)\sigma_a'\tfrac{1}{4\pi}\int_{4\pi}I(\underline{\Omega},\underline{r},t)d\Omega \tag{3}$$

where the Fleck parameter is given by

$$f=\left(1+\frac{\partial B/\partial T}{\rho C_V}\sigma_a'c\Delta t\right)^{-1} \tag{4}$$

Here we have adopted the fully implicit variant of the scheme which is unconditionally stable; the thermal emission source is bounded so that only a fraction of a cell's material energy component can be radiation away during the time-step. The material energy update determined based on energy conservation. In our implementation, f is treated as piecewise constant within each cell.

2.2. Symbolic Implicit Monte Carlo (SIMC)

The SIMC method was derived independently by [4] for thermal radiation transport problems and Brooks [2] for line transport problems. It is based on the idea of postponing the determination of the end of time-step cell temperatures (atomic level populations) until after the tracking step and then solving the set of algebraic equations associated with the energy equation in order to determine these cell temperature values. The thermal emission particles now carry an unknown contribution to their initial weight and the scoring procedure has to keep track of the contributions from each of the symbolic unknowns to the energy deposited in every cell in the mesh.

The original SIMC method is based on a piecewise constant discretisation [12], which produces erroneous results for optically thick cells. This can be generalized to allow either continuous or discontinuous finite element representations of the spatial variations, coupled to a Galerkin weighted energy equation [11]. In this paper we adopt a linear discontinuous spatial discretization and solve the fully non-linear form of the energy equation, but use the linearized equation in order to decide whether or not to invoke mass lumping at the head of the radiation wave in order to ensure that the nodal temperature values remain positive during the solution of the non-linear equations.

3. THE DIFFERENCE FORMULATION

The difference formulation was first proposed by Szöke and Brooks [13] and is based around the idea of subtracting the equilibrium blackbody radiation field from the intensity distribution and modeling this reduced "difference intensity" with Monte Carlo particles. The corresponding transport equation is given by

$$\left[\frac{1}{c}\frac{\partial}{\partial t}+\underline{\Omega}\cdot\nabla+\sigma_a'\right]\left(I(\underline{\Omega},\underline{r},t)-\hat{B}(\underline{\Omega},\underline{r},t)\right)=\overbrace{\sigma_a'\left(B(T(\underline{r},t))-\hat{B}(\underline{\Omega},\underline{r},t)\right)}^{\text{reduced thermal emission source}} \underbrace{-\left[\frac{1}{c}\frac{\partial}{\partial t}+\underline{\Omega}\cdot\nabla+\sigma_a'\right]\hat{B}(\underline{\Omega},\underline{r},t)}_{\text{new difference formulation source terms}} \tag{5}$$

where $\hat{B}(\underline{\Omega},\underline{r},t)$ is an offset function related to the blackbody energy density. In its simplest form $\hat{B}(\underline{\Omega},\underline{r},t)=B(\underline{r},t)$ and the thermal emission term is eliminated in favor of terms involving the spatial and temporal gradients of the blackbody function; for clarity we refer to this as the P_0 difference formulation, as it is based on the use of an isotropic offset function.

The time derivative terms can be treated in a manner analogous to the thermal emission source, while the spatial derivative term requires a new sampling strategy. When integrated over angle the spatial gradient term vanishes, consequently this term must be modeled by emitting pairs of positive and negative weight particles. Each particle is orientated relative to the direction of the spatial gradient, with a cosine angular distribution; note that we are free to choose whether to correlate their other properties such as starting positions, times, directions etc.

This combination of positive and negative weight particles model the net particle flux induced by spatial gradients in the offset function. Particles moving into a region with a larger $\hat{B}$ have negative weights in order to compensate for the increase in the background radiation field, while particles moving into a region where $\hat{B}$ is decreasing have positive weights to compensate the drop in the background radiation field. A similar interpretation can be made for the time evolution of the offset function.

This is the original formulation of the difference scheme, which was undertaken based on an SIMC temporal discretisation, where the end of time-step blackbody function can be expressed in terms of the symbolic unknowns. This same value is used in the offset function in order to completely eliminate the thermal emission term i.e.

$$\hat{B}(\underline{\Omega},\underline{r},t)\rightarrow B(T(\underline{r},t))=B\left(T\left(\underline{r},t^{n+1}\right)\right)\quad t^n<t\le t^{n+1} \tag{6}$$

However, there is no reason why we cannot apply the same strategy to the IMC method, except that we must reformulate the algorithm in terms of quantities available at the start of the time-step.

3.1. Explicit variant

In this section we define an explicit variant of the SIMC difference formulation, where the start of time-step blackbody function is subtracted from the intensity, rather than the fully implicit end of time-step value. Note that this does not affect the temporal accuracy of the scheme, as we are only modifying the function being subtracted and the transport equation itself is unchanged. More generally we express the offset function as

$$\hat{B}(\underline{r},t) = B(\underline{r},t^{n+\alpha}) = (1-\alpha)B(T(\underline{r},t^{n})) + \alpha B(T(\underline{r},t^{n+1})), \quad t^{n} < t \le t^{n+1} \tag{7}$$

so that $\alpha = 1$ corresponds to the fully implicit scheme and $\alpha = 0$ is the explicit variant described above.

For the purposes of the SIMC implementation this explicit variant has some particularly advantageous features. Specifically, in order to control the noise we must estimate the number of symbolic particles required to model each of the source terms, based on the likely changes in the blackbody function during the step. This is quite straightforward for the intensity term, however for the gradient terms this involves modeling the differences between values in adjacent locations and the change in these values during the time-step.

This can lead to severe cancellation errors, and in order to ensure that the SIMC scheme is robust we must be overly cautious in terms of the minimum number of particle required to model each term. Furthermore, treating these terms implicitly means that there will be more entries in the matrix associated with the energy deposition making it harder to solve this non-linear system of equations.

Consequently, adopting the explicit variant of the difference formulation allows us to ensure that the matrix equation remains identical to that associated with the intensity formulation. Furthermore, it ensures that we do not need to worry about the interplay between the values of the blackbody function in neighboring locations when sourcing the particles.

The only downside of this approach is that the thermal emission source is no longer completely eliminated from the SIMC equations. The additional cost of modeling the residual term is given by

$$(1-\alpha)\sigma_a' c\Delta t \frac{1}{c}\frac{\partial B}{\partial t} \tag{8}$$

which is a factor $(1-\alpha)\sigma_a' c\Delta t$ larger than time derivative term in the fully implicit scheme. The presence of this extra term means that the fully implicit scheme will eventually outperform the explicit method for large enough time-steps, despite the significant cost savings for the other terms due to modeling them explicitly. In order to ensure that our scheme remains optimal, we can compute the total number of particles that would be emitted for both the explicit and implicit difference schemes and use this information to determine which variant to adopt for a particular time-step. We refer to this as the “variable alpha” scheme.

3.2. P_1 difference formulation

During this research one aim has been to seek to improve on the original difference formulation in optically thick media. Specifically, the difference formulation is effective at reducing the cost in opaque media so that it is comparable to that in transparent media, but can we do any better? The original difference formulation replaces isotropic thermal emission "energy" particles with linearly anisotropic "flux" particles. Can we eliminate these "flux" particles from opaque regions, and replace them by "net heat-flow" particles?

This can be accomplished by the addition of a linearly anisotropic contribution to the offset function which exactly cancels the spatial gradient term i.e.

$$\begin{aligned}\hat{B}(\underline{\Omega},\underline{r},t) &= B(\underline{r},t) - \sigma_a'^{-1}\underline{\Omega}\cdot\nabla B(\underline{r},t) \\ &= \hat{B}_0(\underline{r},t) + \underline{\Omega}\cdot\underline{\hat{B}}_1(\underline{r},t)\end{aligned} \tag{9}$$

In the more general case the second term is scaled by the total opacity and this is simply the radiation intensity in the equilibrium diffusion limit. For linear elements the spatial gradients are piecewise constant within each cell, consequently we no longer need to emit any particles in the interior of opaque cells (aside from those associated with the time evolution of the offset function). Instead, particles are created on the boundaries between cells according to the flux jumps in order to model the net energy exchange. In SIMC calculations these particles penetrate only a few mean free paths into the interior of optically thick cells, before being terminated by the tracking scheme.

This is reminiscent of the treatment of the gradient particles in a piecewise constant discretisation, however by using a piecewise linear representation of the blackbody function we have not compromised the diffusion limit behavior. Unfortunately, this formulation is only likely to be efficient in regions where the isotropic term dominates i.e. where the radiation field is reasonably diffuse and the spatial gradients are well defined (rather than being dominated by statistical noise). We expect the P_1 difference scheme to significantly outperform the P_0 formulation for modeling behavior such as the penetration of the equilibrium diffusion (Marshak) wave into the Hohlraum wall.

3.3. Hybrid formulations

As discussed above, the difference formulation does not always lead to improved efficiency and/or reduced numerical noise. The magnitude of the spatial gradients of the blackbody function (which may also be contaminated by statistical noise) can be large enough that the resulting difference source terms exceed those associated with the thermal emission process. Furthermore, for systems where the radiation field is significantly weaker than it would be in equilibrium, the associated difference intensity is no longer smaller than the radiation intensity.

To overcome these difficulties we propose a hybrid scheme which seamlessly transitions between the intensity and difference formulations on a cell by cell basis according to the characteristics of the local radiation field. In order to develop a hybrid formulation we express the offset function as follows,

$$\hat{B}(\underline{\Omega},\underline{r},t) = g(\underline{r},t)B(\underline{r},t) - h(\underline{r},t)\sigma_a'^{-1}\underline{\Omega}\cdot\nabla\left(g(\underline{r},t)B(\underline{r},t)\right) \tag{10}$$

where for simplicity both $g(\underline{r},t)$ and $h(\underline{r},t)$ are treated as piecewise constant within each element. These functions are defined in terms of estimates of the amount of computational work associated with the different schemes and are updated at the start of each time-step.

Our initial studies focused on the "hybrid P_0 formulation" with the magnitude of the P_0 term determined locally based on the computational cost within each cell[1]. Unfortunately, the additional overheads associated with the gradients of the transition function can make this scheme significantly more expensive than using the full P_0 term. As an alternative, we propose determining the P_0 term based on the departure from equilibrium (a much smoother criterion than the aforementioned efficiency measure).

Combining the P_0 term determined above with a variable P_1 correction gives rise to the "eq P_0+var P_1 difference" scheme, which provides optimal performance in thick regions and a cost that decreases with increasing penetration depth. We contrast this with the P_0 difference formulation, where the cost asymptotes to a constant value independent of the optical depth. The superiority of our new variable P_1 scheme will be demonstrated in the results section.

4. RELATIVE EFFICIENCY OF THE INTENSITY, DIFFERENCE AND HYBRID FORMULATIONS

4.1. Cost function

For the slab geometry test problems studied in this paper we define a cost function for the Monte Carlo calculation associated with the amount of work required to perform the simulation. This is based on the accumulated magnitudes of the different source terms together with an estimate of the cost of modeling the pseudo-scattering process in IMC calculations.

$$\begin{aligned} C = \int_{x=0}^{L} \Bigg[& \overbrace{\sigma_a'\left|B^{n+1}-\hat{B}_0\right|}^{\text{SIMC thermal emission source term (fully implicit)}} + \sigma_s\left|J-\hat{B}_0\right| + \overbrace{f\sigma_a'\left|B^{n}-B^{n+1}\right| + (1-f)\sigma_a'\left|J-B^{n+1}\right|}^{\text{additional overhead of IMC relative to SIMC}} \Bigg] dx \\ & + \int_{x=0}^{L} \left[\frac{1}{c}\left|\frac{\partial B}{\partial t}\right| + \overbrace{\sigma_a'\left|\frac{\partial}{\partial t}\left(B-\alpha\hat{B}_0\right)\right|\Delta t}^{\text{SIMC only}} + \tfrac{1}{2}\frac{1}{c}\left|\frac{\partial \hat{B}_1}{\partial t}\right| \right] dx \\ & + \int_{x=0}^{L} \left[\tfrac{1}{2}\left|\frac{d\hat{B}_0}{dx}\right| + \tfrac{1}{3}\left|\frac{d\hat{B}_1}{dx}\right| \right] dx + \left[\tfrac{1}{4}\left|J_{\text{ext}}-\hat{B}_0\right| + \tfrac{1}{6}\left|\frac{d\hat{B}_1}{dx}\right| \right]_{x=0}^{L} \end{aligned} \tag{11}$$

where f is the Fleck parameter in the IMC method. The constants account for the angular dependence of the different terms. Here each gradient term is expanded into separate boundary

[1] The results for related scheme, the "hybrid P_1 formulation", which combines this term with a similar treatment for the P_1 term are also presented in this paper.

and cell contributions, and the modulus operator is applied independently to each contribution. The boundary terms have been included separately from the terms in the interior of the problem domain in order to illustrate their non-zero contribution in the intensity formulation. This cost function approximates the computational cost of the transport solution reasonably accurately while removing the sensitivity of the results to the precise details of the implementation.

4.2. Pseudo-scattering and the difference formulation

The cost functions associated with equivalent IMC and SIMC intensity formulation simulations are very similar, suggesting that the numerical noise characteristics of the two schemes are unaffected by the differences in the corresponding transport operators. However, including a non-zero offset function significantly alters the noise characteristics of the schemes, by changing the function being modeled by the Monte Carlo particles.

The introduction of negative weight particles affects our estimate of the expense of pseudo-scattering term. Specifically, we cannot simply use the magnitude of this term as this does not account for cancellations between positive and negative weight particles. Instead, we scale this term by the ratio of the track length estimate derived from the absolute magnitude of the particle weights and the corresponding signed estimate for $J - \hat{B}_0$ i.e.

$$J - \hat{B}_0 \to \left(J - \hat{B}_0\right)\frac{\left|w^+\right| + \left|w^-\right|}{\left|w^+ + w^-\right|} \tag{12}$$

In the IMC method, the reduced absorption coefficient leads to a dramatic increase in the number of particles used to model the difference intensity, while the pseudo-scattering process randomizes their propagation directions. A significant fraction of the scattered particles (that originated from the spatial gradient term) now make zero net contribution to the absorption tallies. The growth in this overhead makes the combination of an IMC temporal discretisation and the difference formulation less efficient than even the original IMC scheme, for the problems considered in this paper. Some form of "Russian roulette" procedure is required in order to effectively address this problem.

This also has a bearing on the efficiency of the difference formulation for problems where the physical scattering term is significant. The cancellations in the estimate of the scattering term are present in both IMC and SIMC calculations, so for scattering dominated regions it may be appropriate to revert to using the intensity formulation in order to avoid this additional overhead.

4.3. Temporal gradients

Numerical experiments indicate that the spatial gradient terms dominate the cost of the difference formulation, consequently we omit the time derivative terms from the cost function (aside from the dominant term associated with using an explicit offset function) to enable the Monte Carlo algorithms to be studied without the need to refer to the time-step used in the simulations. Instead, we focus on the issue of statistical noise in the spatial gradient terms.

Any gains derived from the cost function estimates are likely to be supplemented by the positive effect of reduced noise on the magnitude of the temporal gradient terms, associated with reducing the magnitude of the function being simulated. For the Szilard and Pomraning [14] problem described in the results section, the initial value of the cost function is unity due to the source on the left hand boundary and it asymptotes to a value of two for the difference formulation in the absence of statistical noise[2].

5. RESULTS

The results presented in this paper were generated from a slab geometry test code which includes both IMC and fully implicit SIMC temporal discretizations, with a linear discontinuous spatial representation of the material temperature with mass lumping invoked in order to ensure the positivity of the solution [11]. All sources were modeled based on generating fixed weight particles[3].

The implementation details of this test code merit an entire paper in their own right, especially the strategies used to solve the non-linear SIMC energy equation, but in this paper we will concentrate on the computational results and what they mean in terms of using the difference formulation for general radiation transport problems. Cost functions will be compared rather the computational performance, in order to draw conclusions that do not depend on the algorithm specifics.

In order to demonstrate the hybrid scheme on a realistic problem, we consider the following constant opacity Marshak wave problem defined by Szilard and Pomraning [14]. We use dimensionless units ($c = a = 1$) and the slab is treated as an initially cold ($T(x,t=0) = T_r(x,t=0) = T_0 = 10^{-5}$) pure absorber ($\sigma'_a = \rho C_V = 1$). An isotropic inward source of radiation ($I(\mu > 0, x = 0, t) = 4$) corresponding to a drive temperature of $\sqrt{2}$ is incident on the left hand end of the slab. For our purposes

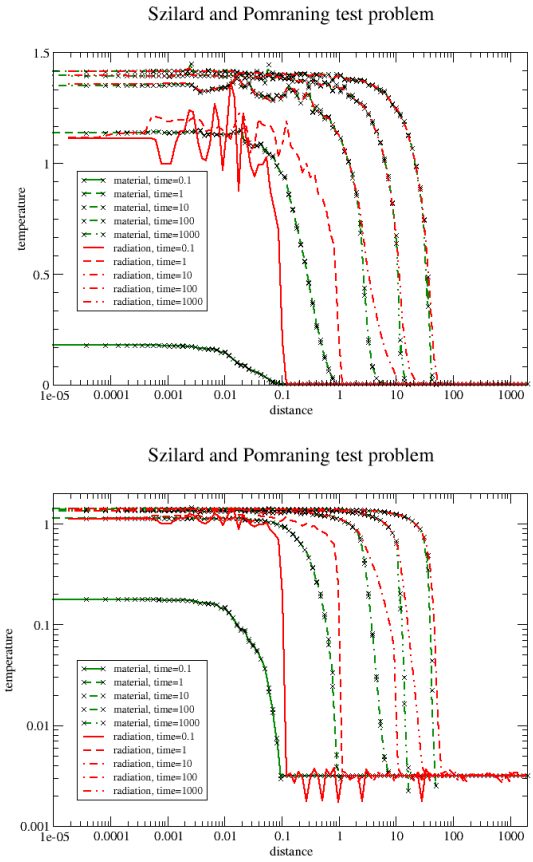

Fig. 1a-b. Sample Monte Carlo results for the Szilard and Pomraning problem illustrating the spatial mesh resolution

[2] Here we have assumed that the blackbody function is monotonically decreasing away from the drive once the drive region reaches equilibrium and the slab is semi-infinite.

[3] The source strength is compared to a maximum weight value (fixed at the start of the calculation) in order to determine the number of uniform weight particles required to model this source term.

the slab is treated as semi-infinite, rather than of fixed length, in order to access the equilibrium diffusion regime.

This is a challenging test of transport algorithms as the early time behavior is dominated by transport effects and equilibration. However, once the material equilibrates the transport effects diminish and the energy flow is then driven by the diffusion of the photons through the medium. There is a corresponding reduction in the speed of the radiation wave as the flow character changes from linear propagation of the drive signal, to absorption and subsequent re-emission by the medium. This latter phase poses a challenge for Monte Carlo algorithms due to the contrast between the short life-time of the photons and the long time-scale for the propagation of the radiation wave.

In order to test out the behavior of the scheme for cells of varying optical thickness we use a fine mesh close to the drive region, with geometrically expanding cells to model the interior of the slab (see Fig. a-b). We compare the different schemes by evaluating the cost function from the current values of the blackbody function and the track length tallies (see Fig. 2a-d). In addition to the scheme used to generate the results we can also predict cost functions for various other combinations, based on the same particle distribution. This allows us to compare the relative merit of these different approaches, noting that these estimates are only hypothetical (changing

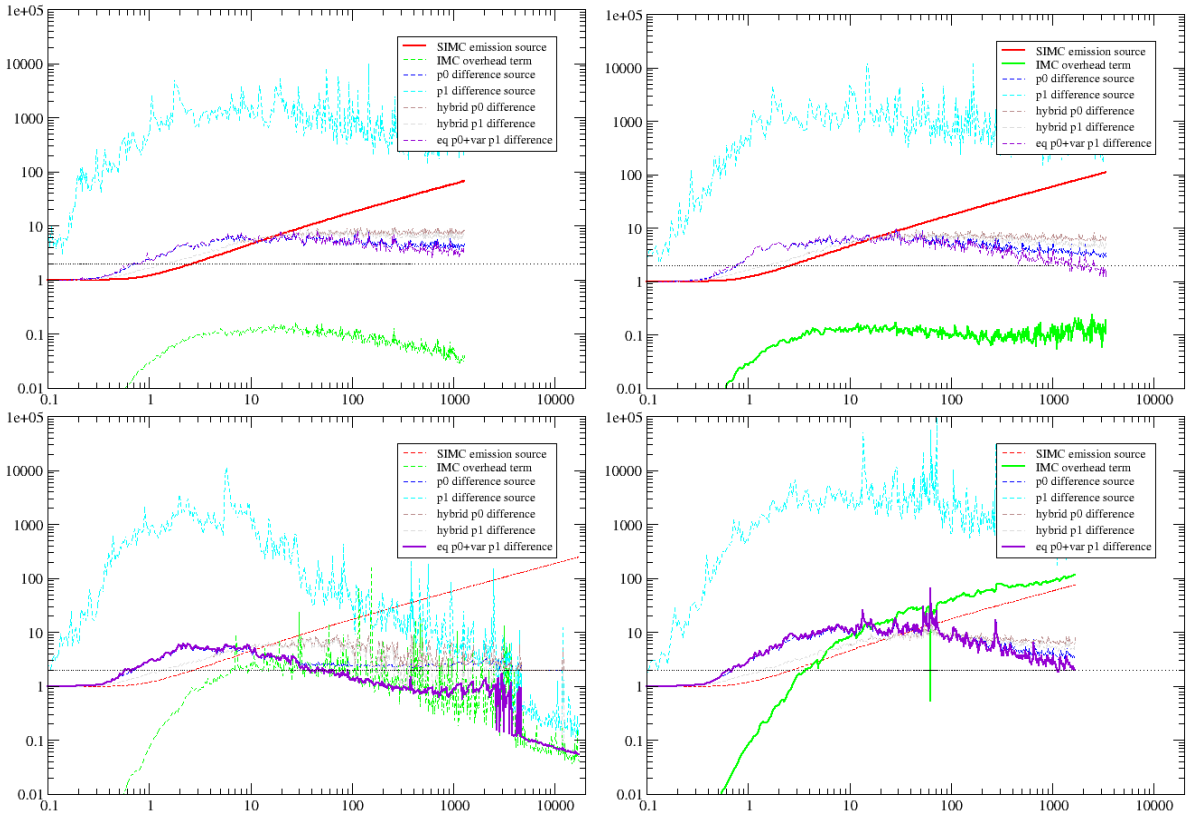

Fig. 2a-d. Cost function verses time (left to right, top to bottom) SIMC, IMC, SDIMC (g=1, variable h and α) and DIMC (g=1, variable h, α=1).

the scheme fundamentally alters the particle statistics).

For the SIMC and IMC calculations using the intensity formulations, the cost of the calculation (bold red line in Fig. 2a verses the sum of the bold red and green lines in Fig. 2b) are almost identical, with the cost of the calculation growing as $t^{1/2}$ once the transport effects diminish. The SIMC calculation spends most of its time tracking thermal emission particles through the mesh, while the IMC calculation is dominated by explicit modeling of the discrete pseudo-scattering events. Despite the differences between the algorithms, there is a great deal of similarity between the cost functions, illustrating the similarities between the noise characteristics of the corresponding radiation fields.

We contrast this with the hybrid P_1 SDIMC (difference SIMC) results in Fig. 2c, where there is an initial overhead relative to the SIMC scheme. This reduces as the radiation starts to penetrate more than a few mean free paths into the slab and this scheme begins to outperforms the original P_0 scheme. In contrast to the P_0 scheme which reaches an asymptotic value of 2, the cost of the hybrid P_1 scheme is proportional to the reciprocal of the penetration depth i.e. the scheme gets cheaper per unit time interval as the radiation wave penetrates deeper into the slab; the jumps around $t = 1{,}000 - 4{,}000$ are caused by changes to the implicitness parameter for the offset function associated with the "variable alpha" scheme.

For the DIMC (difference IMC) results in Fig. 2d we see that the noise in the radiation field persists with time and is similar (slightly larger) than that for the IMC and SIMC calculations. The overheads associated with the pseudo-scattering term dominate the calculation (and are even greater than the cost of modeling the thermal emission source term). This term is responsible for sustaining the numerical noise in the difference intensity field at a level greater than that seen in any of the other calculations and we conclude that, unless we can find an effective "Russian roulette" strategy to remove this source of inefficiency, the difference formulation is incompatible with an IMC temporal discretisation.

We have also compared the effects of mesh resolution on the SDIMC results (see Fig. 3a-c). For optically thick cells, the gradients in the blackbody function remain small close to the drive and the difference formulation outperforms the intensity scheme. However, as we refine the left hand edge of the slab we see that the efficiency of the P_1 difference formulation degrades dramatically, due to the strong flux gradients in this region. A similar but less significant effect is observed for P_0 scheme, but this can be attributed to residual statistical noise in the calculation.

These plots illustrate the important of understanding the behavior in optically thin cells when designing a general purpose transport scheme based on the difference formulation. This is especially important for the P_1 scheme, which can only be employed in regions where the radiation flux is sufficiently well defined and small enough that it inclusion improves the efficiency of the calculation and the accuracy of the results. In transparent cells the P_1 scheme can degrade the effectiveness of the difference formulation by several orders of magnitude and even the P_0 formulation may be a factor of few times less efficient that the intensity formulation.

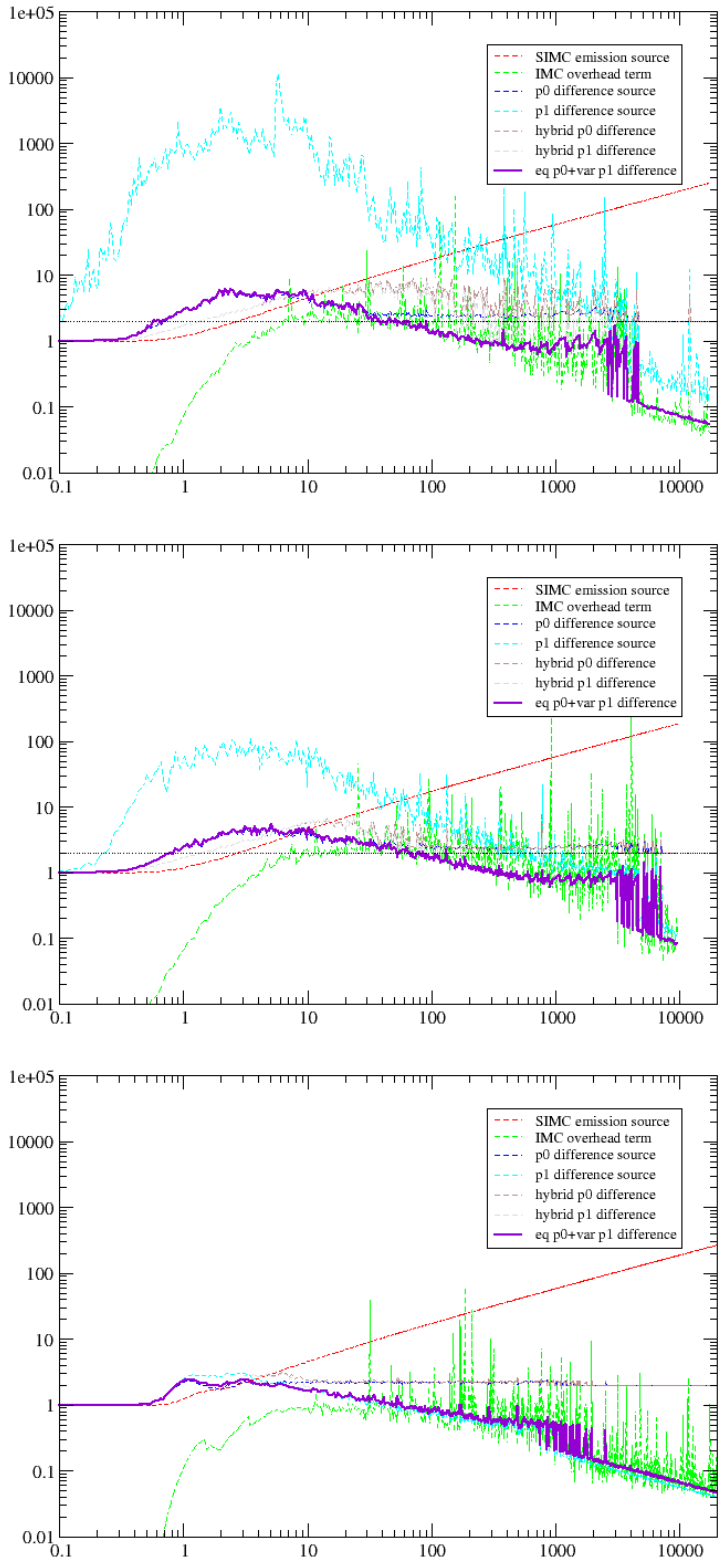

Fig. 3a-c. Cost function vs mesh resolution: min(dx)=3.8e−5, 9.5e−3, 0.69 (top to bottom) for SDIMC (g=1, variable h and α)

6. CONCLUSIONS

In this paper we have demonstrated the effectiveness of new variants of the difference formulation (combined with an SIMC temporal discretisation) for thermal radiation transport problems. The original P_0 difference formulation is always more efficient than the corresponding intensity scheme, in regions which have equilibrated. However, the new P_1 scheme is only efficient in optically thick cells, where the gradients are sufficiently small that the flux correction term does not dominate the blackbody energy density. Consequently, a hybrid scheme is required for general applications which transitions between the intensity formulation and the optimum combination of P_0 and P_1 correction terms. This hybrid scheme outperforms all previous schemes for an example of a deep penetration problem, over a wide variety of different mesh resolutions.

In addition to investigating the SIMC difference formulation, we have also attempted to incorporate the difference formulation into the Fleck and Cummings IMC method. Unfortunately, the resulting scheme is inefficient due to the cancellations between the positive and negative weight particles arising from the difference formulation. The net effect of these overheads is that the difference formulation simulations are more expensive than modeling the original IMC formulation. It may be possible to introduce some variance reduction mechanism to mitigate these effects, but this is unlikely to change our general conclusion that the difference formulation and the IMC method are incompatible.

The success of the hybrid SIMC difference formulation in multi-dimensions depends on the ability to efficiently solve the set of non-linear equations associated with the SIMC energy equation. Investigations are currently underway to examine the viability of this approach for realistic multi-dimensional problems, prior to commencing the multi-dimensional implementation of the ideas presented in this paper.

ACKNOWLEDGEMENTS

I would like to thank Eugene Brooks and his co-workers at LLNL for their hospitality during a recent visit and Mike Clover (SAIC Inc.) for introducing the author to the reduced source Monte Carlo method.

REFERENCES

1. J. A. Fleck, Jr. and J. D. Cummings, “An Implicit Monte Carlo Scheme for Calculating Time and Frequency Dependent Nonlinear Radiation Transport”, *Journal of Computational Physics*, **8**, pp. 313–342 (1971).
2. E. D. Brooks III, “Symbolic implicit Monte Carlo,” *Journal of Computational Physics*, **83**, 433–446 (1989).
3. J. A. Fleck, Jr. and E. H. Canfield, “A random walk procedure for improving the computational efficiency of the implicit Monte Carlo method for nonlinear radiation transport”, *Journal of Computational Physics*, **54**, 508–523, (1984).
4. T. N’kaoua, “Solution of the nonlinear radiative transfer equations with a fully implicit matrix Monte Carlo method coupled with the Rosseland diffusion equation via domain decomposition”, *SIAM J. Sci. Stat. Comput.*, **12**, 505–520 (1991).
5. J–F. Clouët and G. Samba, “A Hybrid Symbolic Monte–Carlo method for radiative transfer equations”, *Journal of Computational Physics*, **188**, 139–156 (2003).
6. N. A. Gentile, “Implicit Monte Carlo Diffusion – An Acceleration Method for Monte Carlo Time–Dependent Radiative Transfer Simulations”, *Journal of Computational Physics*, **172**, 543–571 (2001).
7. J–F. Clouët and G. Samba, “Improvement of the symbolic Monte–Carlo method for the transport equation: P1 extension and coupling with diffusion”, Proceedings of the M&C 2005 meeting, Avignon, France, September 12–15, 2005 (2005).
8. J. D. Densmore, T. J. Urbatsch, T. M. Evans and M. W. Buksas, “Discrete diffusion Monte Carlo for grey implicit Monte Carlo simulations”, Proceedings of the M&C 2005 meeting, Avignon, France, September 12–15, 2005 (2005).
9. G. S. Fraley, E. J. Linnebur, R. J. Mason and R. L. Morse, “Thermonuclear Burn characteristics of compressed deuterium-tritium microspheres”, *Physics of Fluids*, **17**, pp. 474-489 (1974).
10. J. E. Morel, B. T. Adams, T. Noh, J. M. McGhee, T. M. Evans and T. J. Urbatsch, “Spatial discretisations for self adjoint forms of the radiative transfer equations”, *Journal of Computational Physics*, **214**, 12–40 (2006).
11. E.D. Brooks III, A. Szöke, and J.D.L. Peterson, "Piecewise Linear Discretization of Symbolic Implicit Monte Carlo Radiation Transport in the Difference Formulation," ”, *Journal of Computational Physics*, **220**, 471–497 (2006).
12. E. D. Brooks III, M. S. McKinley, F. Daffin and A. Szöke, “Symbolic implicit Monte Carlo radiation transport in the difference formulation: a piecewise constant discretisation,” *Journal of Computational Physics*, **205**, pp. 737–754 (2005).
13. A. Szöke and E. D. Brooks III, “The transport equation in optically thick media”, *Journal of Quantitative Spectroscopy and Radiative Transfer*, **91**, 95–110 (2005).
14. R. H. Szilard and G. C. Pomraning, “Numerical transport and diffusion methods in radiative transfer”, *Nuclear Science and Engineering*, **112**, pp. 256-269 (1992).