AN IMPROVED RANDOM WALK ALGORITHM FOR THE IMPLICIT MONTE CARLO METHOD

Kendra P. Keady^{a,*}, Mathew A. Cleveland^a

^a*CCS-2, Los Alamos National Laboratory
P.O. Box 1663
Los Alamos, NM 87545*

Abstract

In this work, we introduce a modified Implicit Monte Carlo (IMC) Random Walk (RW) algorithm, which increases simulation efficiency for multigroup radiative transfer problems with strongly frequency-dependent opacities.

To date, the RW method has only been implemented in “fully-gray” form; that is, the multigroup IMC opacities are group-collapsed over the full frequency domain of the problem to obtain a gray diffusion problem for RW. This formulation works well for problems with large spatial cells and/or opacities that are weakly dependent on frequency; however, the efficiency of the RW method degrades when the spatial cells are thin or the opacities are a strong function of frequency.

To address this inefficiency, we introduce a RW frequency group cutoff in each spatial cell, which divides the frequency domain into optically thick and optically thin components. In the modified algorithm, opacities for the RW diffusion problem are obtained by group-collapsing IMC opacities below the frequency group cutoff. Particles with frequencies above the cutoff are transported via standard IMC, while particles below the cutoff are eligible for RW. This greatly increases the total number of RW steps taken per IMC time-step, which in turn improves the efficiency of the simulation. We refer to this new method as Partially-Gray Random Walk (PGRW).

*Corresponding author: Email, keadyk@lanl.gov; Phone, +1 (505) 665-9261.
LA-UR-16-25113.

We present numerical results for several multigroup radiative transfer problems, which show that the PGRW method is significantly more efficient than standard RW for several problems of interest. In general, PGRW decreases runtimes by a factor of $\sim 2-4$ compared to standard RW, and a factor of $\sim 3-6$ compared to standard IMC. While PGRW is slower than frequency-dependent Discrete Diffusion Monte Carlo (DDMC), it is also easier to adapt to unstructured meshes and can be used in spatial cells where DDMC is not applicable. This suggests that it may be optimal to employ both DDMC and PGRW in a single simulation.

Keywords: Implicit Monte Carlo, Random Walk, Hybrid Methods

1. BACKGROUND AND MOTIVATION

The IMC method [1] has been used to model thermal radiative transfer in a large variety of physical systems, including problems in inertial confinement fusion and astrophysics. The IMC equations are numerically robust and capable of producing high-fidelity solutions (provided reasonable values are chosen for the time step and zone size). However, one of the major drawbacks of IMC is the large computational cost required to obtain a given precision for realistic problems. In particular, the IMC method is often prohibitively expensive when the problem involves very optically thick regions. IMC particles in these optically thick regions tend to have many closely-spaced collisions, which require an inordinate amount of computing time to track.

To address this issue, Fleck and Canfield introduced an acceleration scheme for IMC particles in optically thick spatial cells, which they named Random Walk (RW) [2]. This name is somewhat misleading: the purpose of the RW method is, in fact, to *replace* a portion of the IMC random walk with a single diffusion step. The diffusion step is carried out on a sphere centered at the particle location, with the radius of the sphere set to the orthogonal distance to the nearest cell boundary (this guarantees constant material properties for the diffusion problem). Before the RW step is carried out, particle attributes are queried and compared to applicability criteria to ensure the diffusion approximation is valid. In the original RW method, the diffusion equation is obtained by integrating the original IMC opacities over the entire frequency domain. In other words, the standard RW algorithm is “fully-gray.”

This fully-gray treatment degrades the efficiency of RW acceleration, especially in problems where the opacity is a strongly-varying function of particle frequency. Because the RW problem is fully-integrated in frequency, optically thin opacity groups are lumped in with optically thick, “diffusive” opacity groups. This, in turn, limits the use of the fully-gray RW method in cells that may appear optically thick to particles in certain frequency groups, but optically thin in others.

If we instead separate the cell-wise frequency domain into diffusive and transport regimes, the RW procedure can be applied to particles in diffusive frequency groups, while particles in transport frequency groups are treated with standard IMC. This should increase the total number of RW steps per time-step (which should consequently improve the efficiency of the simulation), and motivates the work presented here.

In addition, the new RW procedure may avoid some of the “teleportation” error observed in DDMC when particles up-scatter from the diffusion regime to the transport regime [3], because up-scattered PGRW particles will be spatially distributed in accordance with the (analytic) spherical diffusion solution. These DDMC teleportation errors are largely suppressed when particle upscatter locations are sampled from a linear “tilt” distribution, but persist in problems with very strong temperature gradients [3].

The remainder of this paper is organized as follows. Section 2 outlines the asymptotic analysis used to obtain the partially-gray PGRW diffusion equation, while Section 3 uses the resulting diffusion equation to derive expressions for PGRW particle outcomes. We then present numerical results in Section 4, which compare the performance of the new PGRW algorithm to other methods of interest. Finally, we summarize our findings and discuss future work in Section 5.

2. ASYMPTOTIC ANALYSIS

This analysis is based on the procedure in Densmore et. al. [4], which was carried out in support of the frequency-dependent DDMC method. Here, we also include coherent scattering.

We assume familiarity with the standard Fleck and Cummings IMC method for radiative transfer problems [1]. For a detailed derivation of the IMC equations (including an in-depth analysis of the strengths and drawbacks of the method and its variants), the reader is referred to [5].

We begin with the IMC intensity equation for a single energy packet in a

homogeneous medium,

$$\frac{1}{c} \frac{dI}{dt} + \Omega \cdot \nabla I + \sigma_t(\nu)I = \frac{1}{4\pi} \frac{\sigma_a(\nu)b(\nu, T)}{\sigma_p} \int_{4\pi} \int_0^\infty (1-f)\sigma_a(\nu)I d\nu d\Omega + \frac{1}{4\pi} \sigma_s(\nu) \int_{4\pi} I d\Omega . \quad (1)$$

We have omitted the usual time-step subscripts in Eqs. (1) because we consider transport within a single time-step ($t_n < t < t_{n+1}$). We assume all temperature-dependent quantities are evaluated using the material temperature at t_n and held fixed throughout the time-step (as is typically done in IMC calculations). The radiation intensity $I = I(r, \nu, \Omega, t)$; we omit the functional dependence for brevity.

We next define a frequency cutoff ν_c , and assume that the material is optically thick for frequencies below this cutoff and optically thin for frequencies above it (we discuss our procedure for determining ν_c in Section 3.2). With this in mind, we define the following asymptotic scaling:

$$\begin{aligned} c &\rightarrow \frac{c}{\epsilon} , \\ C_v &\rightarrow \epsilon C_v , \\ \sigma &\rightarrow \begin{cases} \frac{\sigma}{\epsilon} , & \nu < \nu_c , \\ \sigma \epsilon^2 , & \nu > \nu_c , \end{cases} \end{aligned} \quad (2)$$

where ϵ is a small parameter (this scaling is consistent with [4] and [6]), and the opacity scaling applies to both absorption (σ_a) and scattering (σ_s). With these definitions, the Fleck factor can be written [4]:

$$f = \epsilon^2 f^{(2)} + O(\epsilon^4) ,$$

while the Planck mean opacity scales as

$$\sigma_p = \frac{1}{\epsilon} \int_0^{\nu_c} \sigma_a(\nu)b(\nu, T)d\nu + \epsilon^2 \int_{\nu_c}^\infty \sigma_a(\nu)b(\nu, T)d\nu = \frac{1}{\epsilon} \sigma_{p,d} + \epsilon^2 \sigma_{p,t} . \quad (3)$$

On the right-hand side of Eq. (3), we use the additional subscripts d and t to differentiate between the diffusive and transport frequency regimes, respectively. These subscripts will also be used to denote the transport and

diffusion intensities, such that

$$I(r, \nu, \Omega, t) = \begin{cases} I_d(r, \nu, \Omega, t) , & 0 < \nu < \nu_c , \\ I_t(r, \nu, \Omega, t) , & \nu_c < \nu < \infty . \end{cases} \quad (4)$$

The asymptotic scalings in Eq. (2)-(3) are applied to the IMC radiation intensity equation for a single energy packet, written separately for a particle with $0 < \nu < \nu_c$ and one with $\nu_c < \nu < \infty$. This yields (to $\mathcal{O}(\epsilon^2)$):

$$\begin{aligned} \frac{\epsilon^2}{c} \frac{dI_d}{dt} + \epsilon \Omega \cdot \nabla I_d + \sigma_t(\nu) I_d = & \\ \frac{1}{4\pi} \frac{\sigma_a(\nu) b(\nu, T)}{\sigma_{p,d}} \int_0^{\nu_c} \int_{4\pi} (1 - \epsilon^2 f^{(2)}) \sigma_a(\nu') I_d d\Omega' d\nu' + & \\ \frac{1}{4\pi} \sigma_s(\nu) \int_{4\pi} I_d d\Omega , & \quad \nu < \nu_c , \end{aligned} \quad (5a)$$

$$\begin{aligned} \frac{\epsilon}{c} \frac{dI_t}{dt} + \Omega \cdot \nabla I_t + \epsilon^2 \sigma_t(\nu) I_t = & \\ \frac{1}{4\pi} \frac{\epsilon^2 \sigma_a(\nu) b(\nu, T)}{\sigma_{p,d}} \int_0^{\nu_c} \int_{4\pi} \sigma_a(\nu') I_d d\Omega' d\nu' + & \\ \frac{1}{4\pi} \epsilon^2 \sigma_s(\nu) \int_{4\pi} I_t d\Omega , & \quad \nu > \nu_c . \end{aligned} \quad (5b)$$

For a detailed description of how Eqs. (5) are obtained, see Appendix A.

After the asymptotic scaling is applied and the two frequency regimes are separated, we observe that Eqs. (5) are weakly coupled; to first order, Eq. (5a) involves only the diffusion intensity (I_d), while Eq. (5b) involves only the transport intensity (I_t). The transport intensity equation includes effective scattering from frequencies below the diffusion cutoff, but this effect is $\mathcal{O}(\epsilon^2)$. This weak coupling results from our choice of opacity scaling, which differs by a factor of ϵ^3 between the diffusion and transport frequency regimes.

We next develop boundary and initial conditions for Eqs. (5). To consider only the subset of random walks which lie entirely within the prescribed sphere, we impose a zero-intensity boundary condition on I_d and I_t at $r = R_o$ [2]. At $t = 0$, we initialize the system with a delta-function source below the diffusion frequency cutoff. These choices yield the following boundary and initial

conditions on I_t and I_d :

$$I_d(r = R_o, \nu, \Omega, t) = 0 , \quad (6a)$$

$$I_t(r = R_o, \nu, \Omega, t) = 0 , \quad (6b)$$

$$I_d(r, \nu, \Omega, t = 0) = \frac{1}{4\pi r^2} \delta(\nu - \nu') \delta(r) , \quad \nu' < \nu_c , \quad (6c)$$

$$I_t(r, \nu, \Omega, t = 0) = 0 . \quad (6d)$$

To proceed with the asymptotic analysis, the intensity is expanded in a power series in ϵ , then truncated such that terms of $\mathcal{O}(\epsilon^3)$ and higher are neglected. This yields:

$$I \approx I^{(0)} + \epsilon I^{(1)} + \epsilon^2 I^{(2)} + \mathcal{O}(\epsilon^3) .$$

If we define

$$\phi_d = \int_0^{\nu_c} \int_{4\pi} I d\nu d\Omega ,$$

and

$$\phi_t = \int_{\nu_c}^{\infty} \int_{4\pi} I d\nu d\Omega ,$$

then, similarly,

$$\begin{aligned} \phi_d^{(k)} &= \int_0^{\nu_c} \int_{4\pi} I^{(k)} d\nu d\Omega , \\ \phi_t^{(k)} &= \int_{\nu_c}^{\infty} \int_{4\pi} I^{(k)} d\nu d\Omega . \end{aligned}$$

Inserting the truncated power series expansion for I into the transport and diffusion IMC intensity equations, then equating $\mathcal{O}(1)$ terms, we obtain:

$$\begin{aligned} \sigma_t(\nu) I_d^{(0)} &= \frac{1}{4\pi} \frac{\sigma_a(\nu) b(\nu, T)}{\sigma_{p,d}} \int_0^{\nu_c} \int_{4\pi} \sigma_a(\nu) I_d^{(0)} d\Omega d\nu + \frac{1}{4\pi} \sigma_s(\nu) \int_{4\pi} I_d^{(0)} d\Omega , \\ \Omega \cdot \nabla I_t^{(0)} &= 0 . \end{aligned}$$

It is immediately clear that the leading-order diffusion intensity is isotropic in angle. Taking integrals of the $\mathcal{O}(1)$ IMC diffusion intensity equation over angle and frequency and performing some algebra, we arrive at the following expression for the leading-order component of I_d :

$$I_d^{(0)} = \frac{1}{4\pi} \frac{b(\nu, T)}{\int_0^{\nu_c} b(\nu, T) d\nu} \phi_d^{(0)}(r, t) .$$

The frequency dependence of the leading-order diffusion intensity is described by the scaled Planck function, which is renormalized by the fraction of the Planck spectrum below the diffusion cutoff, ν_c .

The gradient of the leading-order transport intensity is zero, which indicates that this quantity is constant along spatial characteristics. However, our boundary and initial conditions on I_t (Eqs. (6b) and (6d)) imply that

$$I_t^{(0)}(r = R_0, \nu, \Omega, t) = 0 , \quad (10)$$

which in turn means that

$$I_t^{(0)} = 0 . \quad (11)$$

Thus, the transport intensity is zero to leading order. Using Eq. (4), the $O(1)$ component of the full frequency-dependent intensity can hence be written:

$$I^{(0)} = \begin{cases} \frac{1}{4\pi} \frac{b(\nu, T)}{\int_0^{\nu_c} b(\nu, T) d\nu} \phi_d^{(0)}(r, t) , & \nu < \nu_c , \\ 0 , & \nu > \nu_c . \end{cases}$$

We next consider the $O(\epsilon)$ equations:

$$\begin{aligned} \Omega \cdot \nabla I_d^{(1)} + \sigma_t(\nu) I_d^{(1)} &= \frac{1}{4\pi} \frac{\sigma_a(\nu) b(\nu, T)}{\sigma_{p,d}} \int_0^{\nu_c} \int_{4\pi} \sigma_a(\nu') I_d^{(1)} d\Omega' d\nu' + \\ &\quad \frac{1}{4\pi} \int_{4\pi} I_d^{(1)} d\Omega , \\ \frac{d}{dt} I_t^{(0)} + \Omega \cdot \nabla I_t^{(1)} &= 0 . \end{aligned}$$

Beginning with the equation for $I_d^{(1)}$, we insert the $O(1)$ result, integrate over angle and frequency, and perform some algebra to obtain:

$$I_d^{(1)} = \frac{1}{4\pi} \frac{b(\nu, T)}{\int_0^{\nu_c} b(\nu, T) d\nu} \left(\phi^{(1)}(r, t) - \frac{1}{\sigma_t(\nu)} \Omega \cdot \nabla \phi^{(0)}(r, t) \right) .$$

We next consider the $O(\epsilon)$ equation for the transport intensity. Inserting the $I_t^{(0)}$ result, we observe that the first-order transport intensity is constant along spatial characteristics (as was the case for $I_t^{(0)}$). We again use the boundary and initial conditions for I_t to show that this constant must be zero, yielding:

$$I_t^{(1)} = 0 .$$

Therefore, by Eq. (4),

$$I^{(1)} = \begin{cases} \frac{1}{4\pi} \frac{b(\nu, T)}{\int_0^{\nu_c} b(\nu, T) d\nu} \left(\phi^{(1)}(r, t) - \frac{1}{\sigma_t(\nu)} \Omega \cdot \nabla \phi^{(0)}(r, t) \right), & \nu < \nu_c, \\ 0, & \nu > \nu_c. \end{cases}$$

Following the procedure of Densmore et. al. [4], we reconstruct the approximate intensities using the leading and first-order moments:

$$\begin{aligned} I_d &= \frac{1}{4\pi} \frac{b(\nu, T)}{\int_0^{\nu_c} b(\nu, T) d\nu} \left(\phi^{(0)}(r, t) + \epsilon \phi^{(1)}(r, t) - \frac{\epsilon}{\sigma_t(\nu)} \Omega \cdot \nabla \phi^{(0)}(r, t) \right) + O(\epsilon^2), \\ &= \frac{1}{4\pi} \frac{b(\nu, T)}{\int_0^{\nu_c} b(\nu, T) d\nu} \left(\phi(r, t) - \frac{\epsilon}{\sigma_t(\nu)} \Omega \cdot \nabla \phi(r, t) \right) + O(\epsilon^2). \end{aligned}$$

From Eq. (2), we recall that opacities below the diffusion frequency cutoff were originally scaled as $\sigma \rightarrow \frac{\sigma}{\epsilon}$. Thus, “unscaling” the total opacity yields:

$$I_d = \frac{1}{4\pi} \frac{b(\nu, T)}{\int_0^{\nu_c} b(\nu, T) d\nu} \left(\phi(r, t) - \frac{1}{\sigma_t(\nu)} \Omega \cdot \nabla \phi(r, t) \right) + O(\epsilon^2).$$

The transport intensity is simply:

$$I_t = 0 + O(\epsilon^2), \quad (14)$$

and thus the full frequency-dependent intensity can be written

$$I = \begin{cases} \frac{1}{4\pi} \frac{b(\nu, T)}{\int_0^{\nu_c} b(\nu, T) d\nu} \left(\phi(r, t) - \frac{1}{\sigma_t(\nu)} \Omega \cdot \nabla \phi(r, t) \right) + O(\epsilon^2), & \nu < \nu_c, \\ 0 + O(\epsilon^2), & \nu > \nu_c. \end{cases} \quad (15)$$

We now return to the original IMC radiation intensity equation for a single energy packet in a homogeneous medium, given by

$$\begin{aligned} \frac{1}{c} \frac{dI}{dt} + \Omega \cdot \nabla I + \sigma_t(\nu) I &= \frac{1}{4\pi} \frac{\sigma_a(\nu) b(\nu, T)}{\sigma_p} \int_{4\pi} \int_0^\infty (1-f) \sigma_a(\nu) I d\nu d\Omega + \\ &\quad \frac{1}{4\pi} \sigma_s(\nu) \int_{4\pi} I d\Omega. \end{aligned} \quad (16)$$

To obtain a gray diffusion equation for all $\nu < \nu_c$, we first insert the approximate expression for I (Eq. (15)) into the IMC equation (Eq. (16)). The emission integral on the right-hand side of the equation is split into two parts, one from $0 < \nu < \nu_c$ and one from $\nu_c < \nu < \infty$. The bounds of the latter integral are above the diffusion frequency cutoff, where the intensity is zero to

second order; as a result, this term drops out. We operate on the resulting IMC equation by $\int_0^{\nu_c} \int_{4\pi} (\cdot) d\Omega d\nu$. After algebraic simplification, this yields:

$$\frac{1}{c} \frac{d}{dt} \phi(r, t) - D \nabla^2 \phi(r, t) + \bar{\sigma}_t \phi(r, t) = (1 - f) \frac{\int_0^{\nu_c} \sigma_a(\nu) b(\nu, T) d\nu}{\sigma_p} \bar{\sigma}_a \phi(r, t) + \bar{\sigma}_s \phi(r, t) ,$$

where

$$\bar{\sigma}_x = \frac{\int_0^{\nu_c} \sigma_x(\nu) b(\nu, T) d\nu}{\int_0^{\nu_c} b(\nu, T) d\nu} , \quad (x = t, a, s) ,$$

and

$$D = \frac{1}{3} \frac{\int_0^{\nu_c} \frac{b(\nu, T)}{\sigma_t(\nu)} d\nu}{\int_0^{\nu_c} b(\nu, T) d\nu} .$$

If we further define

$$\begin{aligned} \gamma &= \frac{\int_0^{\nu_c} \sigma_a(\nu) b(\nu, T) d\nu}{\int_0^{\infty} \sigma_a(\nu) b(\nu, T) d\nu} , \\ &= \frac{\int_0^{\nu_c} \sigma_a(\nu) b(\nu, T) d\nu}{\sigma_p} , \end{aligned}$$

and combine like terms, the diffusion equation can be simplified to:

$$\frac{1}{c} \frac{d}{dt} \phi(r, t) - D \nabla^2 \phi(r, t) + [\bar{\sigma}_a - (1 - f) \gamma \bar{\sigma}_a] \phi(r, t) = 0 .$$

Equivalently,

$$\frac{1}{c} \frac{d}{dt} \phi(r, t) - D \nabla^2 \phi(r, t) + [f \bar{\sigma}_a + (1 - f)(1 - \gamma) \bar{\sigma}_a] \phi(r, t) = 0 . \quad (17)$$

The first bracketed term on the left-hand side,

$$f \bar{\sigma}_a \phi(r, t) ,$$

describes particles that are absorbed and not re-emitted within the time step.

The second bracketed term,

$$(1 - f)(1 - \gamma) \bar{\sigma}_a \phi(r, t) ,$$

is non-standard, but can be interpreted by considering each component individually. The quantity $(1 - f)$ represents the total fraction of absorbed particles

that are re-emitted, while $(1 - \gamma)$ is the fraction of the emission spectrum re-emitted above the cutoff frequency ν_c . Finally, $\phi(r, t)$ represents the time and space-dependent particle distribution in the RW sphere. Thus, their product represents those particles that are absorbed in the diffusion frequency regime and re-emitted (“upscattered”) into the transport frequency regime during a time-step.

Assuming a solution of the form

$$\phi(r, t) = \exp(-At)\Psi(r, t) ,$$

and inserting this into the diffusion equation, we obtain:

$$\left([f\bar{\sigma}_a + (1 - f)(1 - \gamma)\bar{\sigma}_a] - \frac{A}{c} \right) \Psi(r, t) + \frac{1}{c} \frac{d\Psi}{dt}(r, t) - D\nabla^2 \Psi(r, t) = 0 .$$

In this form, it is obvious that if we choose

$$A = c [f\bar{\sigma}_a + (1 - f)(1 - \gamma)\bar{\sigma}_a] ,$$

we obtain a source-free diffusion equation for $\Psi(r, t)$:

$$\frac{1}{c} \frac{d\Psi}{dt}(r, t) - D\nabla^2 \Psi(r, t) = 0 .$$

Boundary and initial conditions on $\Psi(r, t)$ can be determined from the original problem for I_d . For the single energy-packet case, the initial condition on I_d is:

$$I_d(r, \nu, \Omega, t = 0) = \frac{1}{4\pi r^2} \delta(\nu - \nu') \delta(r) , \quad \nu' < \nu_c ,$$

where ν' is the initial frequency of the particle. If we integrate over angle and frequency, we obtain an initial condition for $\Psi(r, t)$:

$$\Psi(r, t = 0) = \frac{1}{r^2} \delta(r) .$$

From our zero-intensity boundary condition on I_d , we also obtain a zero boundary condition on $\Psi(r = R_o, t)$:

$$\Psi(r = R_o, t) = 0 .$$

This diffusion equation has the solution:

$$\Psi(r, t) = \frac{1}{2R_o^2} \sum_{n=1}^{\infty} \left(\frac{n}{r}\right) \sin\left(\frac{n\pi}{R_o}r\right) \exp\left(-\left(\frac{n\pi}{R_o}\right)^2 cDt\right). \quad (18)$$

The full solution for $\phi(r, t)$ can thus be written:

$$\begin{aligned} \phi(r, t) &= \exp(-At)\Psi(r, t) \\ &= \exp(-c(1-f)(1-\gamma)\bar{\sigma}_a t) \exp(-cf\bar{\sigma}_a t) \\ &\quad \frac{1}{2R_o^2} \sum_{n=1}^{\infty} \left(\frac{n}{r}\right) \sin\left(\frac{n\pi}{R_o}r\right) \exp\left(-\left(\frac{n\pi}{R_o}\right)^2 cDt\right) \end{aligned} \quad (19)$$

In the next section, results from the asymptotic analysis will be used to derive probability distributions for various PGRW particle events. We will also discuss the practical implementation of the PGRW method, including our straightforward algorithm for determining the cell-wise frequency group cutoff.

3. IMPLEMENTATION

In this section, we describe the version of PGRW recently implemented in the Jayenne Project software at Los Alamos National Laboratory (LANL) [7]. This discussion is divided into two distinct parts; first, in Section 3.1 we discuss the PGRW process itself, including determination of particle eligibility, event selection, and post-RW particle weight and frequency adjustment. Then, in Section 3.2 we describe the frequency group cutoff calculation, which takes place at the beginning of each transport time-step.

3.1. PGRW Eligibility and Event Outcomes

The PGRW method is applied during the standard IMC particle transport process. During each particle transport step, a RW sphere radius is determined using the orthogonal distance to the nearest cell boundary. Initial eligibility for PGRW is then evaluated using current particle attributes (i.e. spatial location and frequency); for a PGRW step to take place, (i) the radius of the RW sphere must be larger than the distance to collision, (ii) the distance to census must

not be the limiting event, and (iii) the particle frequency group must be below the PGRW cutoff for the spatial cell.

Provided these conditions are met, the RW sphere radius is then compared to a set minimum optical thickness requirement (τ_{rw} , which will be discussed further in Section 3.2). If the radius of the RW sphere is smaller than the minimum value, the particle is transported using standard IMC. If, however, the radius is greater than or equal to the minimum value, the particle is deemed eligible for PGRW.

Once PGRW eligibility is established, the solution of the spherical diffusion problem presented in Section 2 is used to sample a particle outcome. Three possible outcomes must be considered: (i) the particle remains within the sphere and is censused (without upscattering), (ii) the particle upscatters to a transport frequency, or (iii) the particle escapes through the surface of the sphere within the time-step (without upscattering). This differs from the standard RW algorithm, in which all particles either go to census or escape through the outer radius of the sphere (because upscattering is not possible).

In the Jayenne implementation of PGRW, the outcome of the RW step is determined by sampling a time-to-occurrence for each possible event. First, we sample a time to upscatter (t_{up}) using

$$t_{up} = \frac{-\ln(\xi)}{(1-f)(1-\gamma)\bar{\sigma}_a c} , \quad (20)$$

where ξ is a uniformly distributed pseudo-random number on $[0, 1]$. The expression for t_{up} is derived from the first exponential term in Eq. (19). Next, the value of t_{up} is compared to the (known) time to census, t_{cen} . If $t_{cen} < t_{up}$, the time to upscatter is discarded; otherwise, t_{cen} is discarded. Next, a time-to-leak (t_{leak}) is sampled from

$$\xi = \int_0^{R_o} \Psi(r, t_{leak}) r^2 dr , \quad (21)$$

where R_o is the outer radius of the RW sphere and $\Psi(r, t)$ is the solution of the source-free diffusion equation (see Eq. (18)). The quantity t_{leak} is the time required for the particle to escape through the surface of the random walk

sphere. Because Eq. (21) cannot be solved explicitly for t_{leak} , values of the distribution are stored in tabular form and used for interpolation. Once obtained, the estimated value of t_{leak} is compared to the minimum of t_{up} and t_{cen} , and the smallest value retained. The event associated with the smallest t -value is selected as the PGRW outcome.

Once an event is chosen, we must sample a new position (and frequency group, in the case of upscattering) for the particle. If the particle leaks from the sphere, it is moved to a new, randomly sampled location on the surface.

If the particle goes to census while still in the sphere, we sample an effective radius (R_{cen}) using a random number and the solution of the spherical RW diffusion problem:

$$\xi = \frac{\int_0^{R_{cen}} \Psi(r, t_{cen}) r^2 dr}{\int_0^{R_o} \Psi(r, t_{cen}) r^2 dr} . \quad (22)$$

Like the t_{leak} calculation, the effective radius at t_{cen} cannot be determined explicitly, and is instead estimated using interpolation between tabulated solution values.

Similarly, if the particle upscatters, we sample a random number and linearly interpolate between tabulated solution values to estimate the radius R_{up} at which the upscatter occurs (with $R_{up} < R_o$). In the upscatter case, we must also sample a new (transport) frequency group for the particle. The transport group is sampled from the frequency groups above the cutoff, according to the emission spectrum for the cell.

Regardless of the PGRW outcome (census, escape, or upscatter), the particle energy-weight must be adjusted to account for absorption. This is done using the second exponential term in Eq. (19),

$$w_{new} = w_{old} \exp(-cf\bar{\sigma}_a \Delta t) , \quad (23)$$

where w_{old} and w_{new} are the old and new particle energy-weights, respectively, and Δt is the elapsed time for the PGRW step. Once the particle energy-weight and frequency have been adjusted, the particle continues transporting according to standard IMC.

3.2. Group Cutoff Calculation

While the outcomes of the PGRW event (census, escape, or upscatter) are well-defined and thus straightforward to implement, the group cutoff determination is largely subjective. In general (excluding resonance regions), opacity curves are a relatively smooth function of frequency. Thus, no clear delineation between “optically thin” and “optically thick” frequency regimes can be made, and the cutoff must be determined according to some user- or developer-defined cell optical thickness criteria.

In this work, a straightforward algorithm is used to determine the frequency group cutoff. First, we set τ_{rw} , the minimum optical thickness of a PGRW sphere (a value of 5 mean-free-paths is used in this work, which is consistent with previous implementations of the RW method [8]). This parameter is used throughout the simulation to determine which particle transport steps are eligible for RW.

We next choose a constant (cell- and group-independent) scaling factor, C . The product of the minimum PGRW radius and the scaling factor,

$$\tau_{min} = C\tau_{rw} , \quad (24)$$

defines the minimum cell optical thickness required for a frequency group to qualify as “optically thick.” Then, to determine the maximum qualifying group in a particular spatial cell, we simply (i) query the mean chord length of the cell (denoted $\bar{s}$, and calculated using the standard expression $\bar{s} = \frac{4V_{cell}}{A_{cell}}$), (ii) compute the optical thickness of the cell (in mean-free-paths) for each frequency group g using

$$\tau_g = \sigma_{t,g}\bar{s} , \quad (25)$$

and (iii) compare this effective group optical thickness to the minimum optical thickness τ_{min} from Eq. (24). If $\tau_g \geq \tau_{min}$, we accept the group as sufficiently optically thick, update the group cutoff for the current cell, and repeat the calculation for group $g + 1$. Otherwise, we break from the loop over energy groups and move to the next spatial cell.

Using Eqs. (24) and (25), we can describe how a particular choice of the scaling factor (C) affects both the group cutoff calculation and the subsequent transport calculation.

When C is small ($\sim \mathcal{O}(1)$), the group optical thickness required to qualify for PGRW is also small. This maximizes the number of frequency groups which are eligible for PGRW; however, it also restricts the permissible size range of RW spheres in a given spatial cell. As the scaling factor limits to zero, all frequency groups meet the optical thickness criterion, and the PGRW method becomes indistinguishable from the standard RW method.

Conversely, when C is very large ($\sim \mathcal{O}(100)$), fewer frequency groups will meet the optical thickness requirement in a particular cell. However, in this case the range of permissible RW sphere radii is broadened; because the collapsed PGRW opacity is very large, RW spheres with small radii (in cm) will still appear optically thick (in mean-free-paths). As the scaling factor becomes infinite, no frequency groups will qualify for PGRW, and the simulation will revert to standard IMC.

The preceding paragraphs describe the central tradeoff of PGRW; by treating some portion of the frequency domain as ineligible for RW, the PGRW method increases the range of valid RW sphere sizes. This tradeoff proves especially beneficial in problems where the multigroup opacity is a strong function of frequency group.

4. NUMERICAL RESULTS

In this section, we present numerical results demonstrating the performance of the new PGRW method. In addition, we compare the performance of PGRW to IMC, standard RW, and DDMC. First, we consider a pair of 1-D problems taken from [4], one of which is homogeneous in space and the other heterogeneous. We then present results for a 2-D hohlraum problem, with geometry taken from [9]. To better test the performance of our PGRW method, we substitute multigroup opacities for the original gray treatment of the hohlraum problem.

4.1. 1-D Homogeneous Marshak Wave

The first test problem considered in this work is a homogeneous 1-D Marshak wave. The material begins at a uniform temperature of 1 eV, and is heated by a 1 keV boundary source at $x = 0$. The opacity is temperature and frequency-dependent, and is given by

$$\sigma_a(\nu, T) = \frac{1000}{(h\nu)^3 \sqrt{kT}}, \quad (26)$$

where the quantities $(h\nu)$ and (kT) are reported in keV.

We approximate the frequency dependence of the opacity with a multi-group structure, which consists of 30 groups logarithmically spaced from $E_{min} = 0.1$ eV to $E_{max} = 100$ keV. Group-wise opacities are determined using Eq. (26) evaluated at the group midpoint frequency, and span over 17 orders of magnitude. Thus, while particles in the lower frequency groups will have very short mean-free-paths, particles in the higher groups will stream long distances between collisions.

The spatial grid size is uniform and fixed at $\Delta x = 0.02$ cm, and the total extent of the slab is 5 cm with a reflecting boundary at the right-hand edge. We use a time-step size of $\Delta t = 0.01$ ns, and run the simulation to a total time of $t = 1$ ns (100 time-steps). 10^5 particles are simulated per time-step.

In Figure 1, we plot the material temperature profile at $t = 1.0$ ns for the left-most centimeter of the problem domain. Excellent agreement is observed between the PGRW, RW, DDMC, and IMC solutions. Since the DDMC implementation in Jayenne uses advanced techniques to mitigate upscatter teleportation, we do not observe the error described in [3] in our DDMC output.

Next, in Figure 2a we show the PGRW material temperature profile for $0 < x < 1$ cm at $t = 0.01$ ns, 0.2 ns, and 1.0 ns. The results shown here use the optimal scaling factor for this problem ($C = 20$). The second graph (Figure 2b) plots the cell-wise PGRW frequency group cutoff at $t = 0.01$ ns, 0.2 ns, and 1.0 ns (note that the frequency cutoff is given as the index of the maximum diffusive frequency group).

Figure 1: Material temperature at $t = 1.0$ ns, homogeneous problem

At $t = 0$ ns, the material is at a constant temperature. As a result, multigroup opacities are uniform across the slab, and the PGRW group cutoffs are uniform as well. As the boundary source at $x = 0$ cm begins to heat the material, the multigroup opacities decrease ($\propto (kT)^{-0.5}$, from Eq. (26)) and the material becomes less opaque than it was in the previous time-step. This, in turn, causes fewer frequency groups to qualify as “optically thick” according to the PGRW criteria.

This behavior is clearly observed in Figure 2b. As the elapsed problem time (t) increases, the PGRW frequency group cutoff decreases, first at the leftmost (hottest) edge of the slab. As the material heats from left to right, the frequency group cutoff decreases accordingly.

Finally, in Figure 3, we plot the simulation wall-time versus the PGRW cutoff criterion (purple curve). All simulations for the homogeneous problem were run with 4 MPI ranks in domain replication mode (where each MPI rank owns a full copy of the mesh). Simulation times for IMC (yellow line), standard RW (green line), and DDMC (blue line) are included for reference. Timing data is also summarized in Table 1, with two variants of the PGRW scaling factor (C , from Eq. (24)) included to provide reasonable upper and lower bounds. This

(a) Material temperature

(b) Frequency group cutoffs

Figure 2: Material temperature and PGRW group cutoff, homogeneous problem

data shows that PGRW is roughly 2.7 times faster than standard IMC when the *least-optimal* scaling factor is used, and ~ 6.5 times faster than standard IMC when the scaling factor is nearly optimal. In addition, the PGRW simulation is faster than the standard RW simulation for all values of the scaling factor considered here.

Figure 3 shows that the PGRW method is reasonably insensitive to the choice of scaling factor, over the range of values considered in this work (2 - 125).

Figure 3: Homogeneous problem timing

Table 1: Homogeneous problem speedup, relative to IMC

Method	Wall time (s)	Speedup
PGRW, $C = 25$ (optimal)	619	6.53
PGRW, $C = 2$ (worst)	1496	2.70
RW	2884	1.40
DDMC	102	39.6
IMC	4042	—

We observe a noticeable increase in wall time for scaling factors $C < \sim 10$, but the trend is relatively flat for $10 < C < \sim 100$. At worst, the PGRW method is still considerably faster than standard RW for this problem. However, it is slower than DDMC for all values of the scaling factor.

4.2. 1-D Heterogeneous Marshak Wave

Next, we present results for a 1-D Marshak wave problem with two material regions. The total slab size is 3 cm: material 1 extends from $x = 0$ to $x = 2$ cm, with opacities given by

$$\sigma_a(\nu, T) = \frac{10}{(h\nu)^3 \sqrt{kT}} , \quad (27)$$

while material 2 extends from $x = 2$ to $x = 3$ cm, with opacities given by

$$\sigma_a(\nu, T) = \frac{1000}{(h\nu)^3 \sqrt{kT}} . \quad (28)$$

We note that material 1 and material 2 opacities differ only by a multiplicative constant of 100. Thus, at a given temperature, material 1 will always appear optically thinner than material 2 (to a particle with a given frequency).

As in the homogeneous problem, we approximate the continuous frequency dependence of the opacity using a 30-group structure, with logarithmically-spaced bins from 0.1 eV to 100 keV. Multigroup opacities are evaluated using Eq. (27) and Eq. (28) at the midpoint of each frequency bin for materials 1 and 2, respectively. The spatial discretization is $\Delta x = 0.02$ cm thick in the material 1 region, and $\Delta x = 0.005$ cm thick in the material 2 region. The left-most cell of material 2 (from $2.0 < x < 2.005$ cm) is further divided into 10 subcells of increasing size, with each subcell a factor of 1.58 thicker than the last. 10^5 total particles were simulated per 0.01 ns time-step, and the simulation was run for a total of 100 time-steps.

Figure 4: Material temperature at $t = 1$ ns, heterogeneous problem

First, in Figure 4 we compare the IMC, RW, and PGRW material temperature profiles at $t = 1$ ns. We were unable to simulate the heterogeneous

Marshak wave with DDMC, because the mesh type used is not yet supported by the DDMC implementation in Jayenne.

Figure 4 shows excellent agreement between the IMC, RW, and PGRW temperature solutions. This suggests that use of the new PGRW algorithm does not noticeably bias the solution.

(a) Material temperature

(b) Group cutoffs

Figure 5: Material temperature and PGRW group cutoff, heterogeneous problem

In Figure 5, we plot the material temperature and group cutoff distributions at $t = 0.01, 0.2$, and 1.0 ns. The frequency group cutoff is flat near the

beginning of the simulation, when the material is uniformly cold. We observe a step change in the group cutoff at the material interface, due to the variation in opacities between the two materials.

As the simulation proceeds, the material heats from left to right, and the group cutoff decreases from left to right accordingly. At $t = 1.0$ ns, the steep temperature gradient established near the material interface ($x = 2.0$ cm) causes a local dip in the frequency group cutoff for the optically thick material.

Timing data for the heterogeneous problem is shown in Figure 6, and summarized in Table 2. In Figure 6, the PGRW time versus scaling factor curve is shown in purple, with IMC and RW wall times included for comparison. All heterogeneous problem simulations were run on 4 MPI ranks with domain replication.

Figure 6: Heterogeneous problem timing

Table 2: Heterogeneous problem speedup, relative to IMC

Method	Wall time (s)	Speedup
PGRW, $C = 20$ (optimal)	220	3.81
PGRW, $C = 2$ (worst)	361	2.32
RW	582	1.44
IMC	838	—

The PGRW simulation is faster than both standard RW and IMC for all values of the scaling factor, C . The speedup factor with respect to IMC varies from 2.32 (for PGRW with the least-optimal scaling factor, $C = 2$) to 3.81 (for PGRW with the near-optimal scaling factor, $C = 20$).

4.3. 2-D Hohlraum

Finally, we apply the new PGRW method to a modified version of the Inertial Confinement Fusion (ICF) hohlraum problem originally presented in [9]. Geometry for the hohlraum is taken directly from [9], while material opacities are taken from the homogeneous Marshak wave presented in Section 4.1. This combination allows us to test the performance of the PGRW method for a 2-D problem with frequency-dependent opacities (the opacities specified in [9] depend only on temperature).

Figure 7: 2-D hohlraum problem geometry

Problem geometry is shown in Figure 7. Regions colored blue indicate hohlraum material, while regions colored white are void. Because this problem is modeled in cylindrical geometry, the $r = 0$ boundary is treated as reflecting. All other boundaries are vacuum. The problem is initially cold ($T = 0$ everywhere), with a 1 keV source at $z = 0$ cm. The spatial discretization is fixed at $\Delta r =$

$\Delta z = 0.01$ cm, for a total of 9100 spatial cells. We use a temporal discretization of $\Delta t = 0.1$ ns, and run the simulation for 100 time-steps. 10^5 total particles are simulated per time-step.

In Figure 8, we compare material temperature distributions at $t = 10$ ns for the IMC, RW, PGRW, and DDMC methods. The PGRW simulation uses the optimal scaling factor for this problem ($C = 15$).

Figure 8: Material temperature (keV) at $t = 10$ ns, hohlraum problem

Though the material temperature distributions are somewhat noisy, overall trends agree well between the four simulations. Specifically, the material temperature is highest in the wall near the source at $z = 0$, and is also high in the portion of the hohlraum nearest this wall. We note that some of these spatial cells violate the maximum principle, which is likely a result of our coarse temporal discretization. The coolest material regions are located at the backside of the hohlraum (near $z = 1$ cm) and the rear/upper walls. Slight temperature dips are observed in spatial cells adjacent to the void.

Solutions for all methods appear to exhibit increased noise near $r = 0$ cm,

which is probably due to correlations between particles entering and exiting the reflecting boundary.

We next present timing results for the 2-D hohlraum problem, with all cases run using domain replication on 8 MPI ranks. This data is shown visually in Figure 9, and also summarized in Table 3.

Figure 9: Hohlraum problem timing

Table 3: Hohlraum problem speedup, relative to IMC

Method	Wall time (s)	Speedup
PGRW, $C = 2$ (worst)	4068	1.46
PGRW, $C = 15$ (optimal)	1564	3.79
RW	5277	1.12
DDMC	321	18.5
IMC	5923	—

The PGRW time versus scaling factor curve for the hohlraum problem closely resembles the curves presented for other problems in this section, and is minimized when the scaling factor $C \approx 15$. The trend is largely flat for $10 < C < 125$, indicating that the efficiency of the PGRW method is relatively insensitive to the scaling factor. The PGRW algorithm is more efficient than standard RW for all values of C tested here; however, PGRW is less efficient

than DDMC for all values of C .

5. DISCUSSION

Here, we have described a modification of the standard RW method, which was recently implemented in the Jayenne Project software at LANL. While the original RW method was fully-gray in frequency, our modified PGRW algorithm introduces a frequency group cutoff that divides the frequency domain into transport and diffusion regimes. Particles above the frequency group cutoff are transported using standard IMC, while particles below the cutoff are eligible for PGRW, which replaces a number of closely-spaced collisions with a single diffusion step. When compared to standard RW, the PGRW method significantly increases the total number of RW steps performed per transport time-step; this, in turn, reduces simulation time considerably for multigroup problems.

The performance of the new PGRW method appears reasonably insensitive to the optical thickness scaling factor (C) across a wide range of values. In addition, the optimal value of C changed very little across the three problems considered in this paper, which suggests that the PGRW method can be used as a “black box” without additional tuning by the user. Based on the numerical results presented in Section 4, near-optimal performance is achieved when C is between 15 and 25.

While preliminary results for the PGRW method are promising, a number of outstanding questions remain. First, the problems presented here use simplified opacity distributions, which are monotonically decreasing functions of frequency. This simplifies the analysis by allowing us to define a single cutoff that separates optically thick and optically thin opacity groups. However, real opacity data often contains high-amplitude “spikes” at resonant frequencies, which complicate the curves significantly. These opacity curves may require the implementation of a modified group cutoff calculation. If this is the case, it is likely that the arbitrary group collapse methodology already developed for

DDMC [10] and Implicit Monte Carlo Diffusion (IMD) [3] can be adapted to the PGRW method.

The opacities used in this paper all vary as $(h\nu)^{-3}$ and $(kT)^{-0.5}$. For completeness, the powers of $(h\nu)$ and (kT) should be varied to include other reasonably realistic opacity distributions.

In addition, we have not systematically studied the effect of the spatial discretization on the PGRW method. We intuit that the PGRW method will become less effective as the spatial grid is refined (similar to standard RW and DDMC), but we did not quantify this effect. In any case, we still expect the PGRW method to outperform the standard RW method regardless of spatial discretization.

In its present form, the PGRW method is still significantly slower than DDMC; thus, we suggest that DDMC be used in problems where it is applicable. However, the DDMC method has not yet been extended to problems with unstructured spatial grids. In addition, differing applicability requirements between DDMC and PGRW may cause some portion of the spatial cells in a problem to be eligible for PGRW, but not DDMC. Thus, in some cases it may be optimal to combine the DDMC and PGRW methods in a single simulation.

Currently, we use a simple linear scaling of the mean chord length to determine the diffusive frequency group cutoff in each spatial cell. However, parametric studies presented in Section 4 show that this straightforward frequency cutoff calculation yields good results over a wide range of scaling factors. At worst, the new PGRW method is comparable to standard RW; at best, it reduces runtimes by a factor of four or more, compared to standard RW.

6. Acknowledgments

This work was performed under U.S. government contract DE-AC52-06NA25396 for Los Alamos National Laboratory, which is operated by Los Alamos National Security, LLC, for the U.S. Department of Energy.

7. REFERENCES

- [1] J. A. Fleck and J. D. Cummings, “An Implicit Monte Carlo Scheme for Time and Frequency Dependent Nonlinear Radiation Transport,” *J. Comp. Phys.* **8**, 313 (1971).
- [2] J. A. Fleck and E. H. Canfield, “A Random Walk Procedure for Improving the Computational Efficiency of the Implicit Monte Carlo Method for Nonlinear Radiation Transport,” *J. Comp. Phys.* **54**, 508 (1984).
- [3] M. Cleveland and N. Gentile, ”Mitigating Teleportation Error in Frequency-dependent Hybrid Implicit Monte Carlo Diffusion Methods,” *J. Comp. and Theor. Trans.* **43**, 6 (2014).
- [4] J. Densmore, K. G. Thompson, and T. J. Urbatsch, ”A Hybrid Transport-diffusion Monte Carlo Method for Frequency-dependent Radiative-transfer Simulations,” *J. Comp. Phys.* **231**, 6924 (2012).
- [5] A. B. Wollaber, “Four Decades of Implicit Monte Carlo,” *J. Comp. and Theor. Trans.*, DOI: 10.1080/23324309.2016.1138132 (2016).
- [6] J. Morel and J. Densmore, ”A Two-component Equilibrium-diffusion Limit,” *Ann. Nucl. Energy* **32**, 233, (2004).
- [7] T. Urbatsch and T. Evans, “Milagro Version 2 An Implicit Monte Carlo Code for Thermal Radiative Transfer: Capabilities, Development, and Usage,” LA-14195-MS (2006).
- [8] J. Giorla and R. Sentis, “A Random Walk Method for Solving Radiative Transfer Equations,” *J. Comp. Phys.* **70**, 145 (1987).
- [9] R. McClarren and T. Urbatsch, “A Modified Implicit Monte Carlo Method for Time-dependent Radiative Transfer with Adaptive Material Coupling,” *J. Comp. Phys.* **228**, 5669 (2009).

- [10] R. T. Wollaeger et. al., “Radiation Transport for Explosive Outflows: A Multigroup Hybrid Monte Carlo Method,” *The Astro. J. Supp. Series* **209**, 36 (2013).