

Predicting Monte Carlo Variance and Runtime from Source-Term Scaling in the Thick Diffusion Limit

M Pozulp, P Brantley, T Haut

Sep 2026

International Conference on
Mathematics and Computational
Methods Applied to Nuclear Science
and ...

Columbus, OH, United States of America

2027-04-18 through 2027-04-22

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Predicting Monte Carlo Variance and Runtime from Source-Term Scaling in the Thick Diffusion Limit

M. Pozulp^{1,*}, T. Haut², P. Brantley³

¹Applications, Simulations, and Quality Division, Lawrence Livermore National Laboratory, Livermore, CA 94550, USA;

²Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, Livermore, CA 94550, USA;

³Design Physics Division, Lawrence Livermore National Laboratory, Livermore, CA 94550, USA

[leave space for DOI, which will be inserted by ANS]

ABSTRACT

Monte Carlo transport calculations can exhibit large estimator variance or long particle histories in optically thick, diffusive regimes. We use the thick diffusion limit (TDL) scaling to predict which behavior occurs from the asymptotic strength and Monte Carlo representation of source terms. When accumulating a path-length estimator in an infinite medium, a source whose leading coefficient grows as the TDL parameter tends to zero produces unbounded variance when sampled through particle weights and unbounded expected work per history when represented by repeated collision events. The analysis recovers the $O(\varepsilon^{-2})$ collision work of the standard Monte Carlo linear transport method and the $O(\varepsilon^{-1})$ variance of single-collision Monte Carlo source iteration. Numerical results exhibit the predicted trends. Applied to gray thermal radiative transfer, the same framework explains the variance/runtime tradeoff among explicit, implicit, source-iteration, difference-formulation, and transport-diffusion methods. It therefore provides a screening tool for identifying costly or noisy Monte Carlo formulations before implementation.

Keywords: Monte Carlo, thick diffusion limit, variance, runtime

1. INTRODUCTION

Monte Carlo methods for solving the neutral particle Boltzmann transport equation can suffer high variance or long runtimes in optically thick, diffusive regimes. Numerical experiments can determine the variance and runtime in these regimes: the method is implemented, then observations of the variance and runtime are recorded for the calculations of interest. It would be useful if variance and runtime could be determined theoretically, not just experimentally. Doing so would be helpful for Monte Carlo method development. If an analysis predicts high variance or long runtimes, a method developer could pursue a different method rather than taking the time to implement the noisy or slow method. The analysis could also guide variance reduction or alternative modeling choices to address the predicted issue.

Predictive analysis is helpful for deterministic transport method development. For example, Larsen, Morel, and Miller's analysis of the discrete ordinates equations predicts the accuracy of different spatial discretizations in optically thick, diffusive regimes: if the discretized diffusion equation derived from the thick

*pozulp1@llnl.gov

diffusion limit (TDL) scaling of the discretized transport equation is invalid, the transport solution is inaccurate [1]. This theoretical tool aids the development of discrete ordinates methods by identifying the spatial discretizations that lead to inaccurate solutions. Can Monte Carlo variance and runtime be predicted too?

The size of the time step required for explicit Monte Carlo thermal radiative transfer (TRT) methods to maintain stability is small enough that explicit treatment of the time dependence is impractical, and the results tend to be noisy in optically thick, diffusive regimes. More practical time step sizes are possible in Monte Carlo TRT methods that use Fleck and Cummings' linearized TRT equations combined with Fleck and Cummings' Monte Carlo technique for solving the linearized equations [2]. The linearized equations are unconditionally stable, the TRT equations are recovered as the size of the time step approaches zero, and the results tend to be less noisy than explicit Monte Carlo. However, a single time step can be orders of magnitude more expensive. Can an analysis of the TRT equations and Fleck and Cummings' linearized equations explain why Monte Carlo solutions of the latter tend to be less noisy and slower than the former?

Monte Carlo TRT methods that neither treat the time dependence explicitly nor use Fleck and Cummings' linearized equations can be noisy in optically thick, diffusive regimes. The methods of Brooks [3] and N'Kaoua [4] are two examples. The Monte Carlo component in both Brooks' method and N'Kaoua's method is faster than Fleck and Cummings' Monte Carlo method for solving Fleck and Cummings' linearized equations because effective scattering events are omitted from particle histories, but results from Brooks' method and N'Kaoua's method can be noisier. Szoke and Brooks' difference formulation reduces variance by transforming the emission source [5]. Can an analysis of the transport equation that arises in Szoke and Brooks' difference formulation explain why solutions of the transformed equation are less noisy?

Monte Carlo methods do not necessarily need to model physical scattering, effective scattering, fission, and other types of collision sources as events in particle histories. The contributions of these source terms can be included in the particle weights. However, this approach requires iterating, because these sources depend on functions of the dependent variable. In every cycle of the iteration, the lagged sources are recomputed using iterates from the previous cycle, such as the angle-integrated solution iterate. The weights of the Monte Carlo particles are updated using the newly computed sources, and the particles are tracked through the geometry to compute new iterates. The iteration is finished when the lagged sources converge. Monte Carlo source iteration techniques can improve the runtime of Monte Carlo calculations when the iteration converges in a small number of inexpensive cycles. The effect of the additional weight from the lagged sources on the variance is minimal in transparent regimes, but it can be substantial in optically thick, diffusive regimes. Can an analysis of the transport equation predict which sources will cause high variance if they are lagged?

The answer to each question at the end of the previous four paragraphs is yes. Our analysis uses the same TDL scaling as Larsen, Morel, and Miller [1] to show that if the leading sampled source in the TDL-scaled equation grows as the TDL parameter $\varepsilon \rightarrow 0$, then representing that source in particle weights yields an unbounded path-length-estimator variance, whereas representing it through repeated collision events yields an unbounded expected work per history. The conclusion holds for both linear transport and TRT.

We begin our analysis in Section 2 with the linear transport equation. After making predictions, we calculate the variance and runtime of different methods and compare them to our predictions. The numerical results include an infinite medium and a problem with both optically thick and optically thin regions. We analyze the TRT equations in Section 3, where we explain the high variance of Monte Carlo TRT methods that do not solve Fleck and Cummings' linearized equations [2], and the long runtimes of those that do. We also use the analysis to explain why Szoke and Brooks' difference formulation [5] reduces variance by replacing unbounded sources with bounded ones, why Monte Carlo source iteration exchanges long runtimes for high variance, and why Monte Carlo methods that solve Fleck and Cummings' linearized equations and make a diffusion approximation [6, 7, 8] have bounded variance and runtime. We conclude in Section 4.

2. LINEAR TRANSPORT ANALYSIS AND NUMERICAL RESULTS

Our analysis follows that of Pozulp and Haut [9], who scaled a linear transport equation using the TDL scaling from Larsen, Morel, and Miller [1] to derive an expression for the variance of path-length estimators for single-collision Monte Carlo methods. Our analysis extends their results to consider runtimes by examining the scattering ratio, which we use to predict the runtime of standard Monte Carlo linear transport and Monte Carlo source iteration. The linear Boltzmann transport problem that we consider is,

$$\boldsymbol{\Omega} \cdot \nabla \psi + \sigma_t \psi = \frac{\sigma_s}{4\pi} \int_{\mathbb{S}^2} \psi \, d\boldsymbol{\Omega}' + q, \quad (1a)$$

where $\boldsymbol{\Omega} \in \mathbb{R}^3$ is the unit direction vector, ψ is the radiation intensity, σ_t is the total material opacity, σ_s is the scattering opacity, $\mathbb{S}^2$ is the unit sphere, and q is an arbitrary source function. Equation (1a) is subject to an inflow boundary condition, which defines the radiation intensity on the domain boundary in the hemisphere of directions that point into the domain at a given location on the domain boundary surface,

$$\psi(\mathbf{x}, \boldsymbol{\Omega}) = \psi_{\text{inc}}(\mathbf{x}, \boldsymbol{\Omega}), \quad \mathbf{x} \in \partial\mathcal{D} \text{ and } \boldsymbol{\Omega} \cdot \mathbf{n} < 0, \quad (1b)$$

where $\mathbf{x} \in \mathbb{R}^3$ is the spatial position coordinate, ψ_{inc} is the inflow source function, $\partial\mathcal{D}$ is the domain boundary, and $\mathbf{n}$ is the outward-facing unit normal vector on the domain boundary surface. The TDL scaling of the problem data in terms of the asymptotic analysis parameter $\varepsilon \in (0, 1]$ is,

$$\sigma_t^\varepsilon = \sigma_t / \varepsilon, \quad \sigma_a^\varepsilon = \sigma_a \varepsilon, \quad q^\varepsilon = q \varepsilon, \quad (2)$$

where $\sigma_a = \sigma_t - \sigma_s$ is the absorption opacity. The unscaled problem data are independent of ε and thus order unity,

$$\sigma_t \sim 1, \quad \sigma_a \sim 1, \quad q \sim 1. \quad (3)$$

Therefore, the order dependence of the scaled problem data on ε is,

$$\sigma_t^\varepsilon \sim \varepsilon^{-1}, \quad \sigma_a^\varepsilon \sim \varepsilon, \quad q^\varepsilon \sim \varepsilon. \quad (4)$$

Applying the scaling in Equation (2) to the data in Equation (1a) produces a new Boltzmann transport problem for ψ^ε ,

$$\boldsymbol{\Omega} \cdot \nabla \psi^\varepsilon + \frac{\sigma_t}{\varepsilon} \psi^\varepsilon = \frac{\sigma_t - \sigma_a \varepsilon}{4\pi} \int_{\mathbb{S}^2} \psi^\varepsilon \, d\boldsymbol{\Omega}' + \varepsilon q, \quad (5a)$$

subject to a new inflow boundary condition with TDL-scaled inflow data,

$$\psi^\varepsilon(\mathbf{x}, \boldsymbol{\Omega}) = \psi_{\text{inc}}^\varepsilon(\mathbf{x}, \boldsymbol{\Omega}), \quad \mathbf{x} \in \partial\mathcal{D} \text{ and } \boldsymbol{\Omega} \cdot \mathbf{n} < 0. \quad (5b)$$

The scattering ratio is $\omega = \sigma_s / \sigma_t$, and so the TDL-scaled scattering ratio $\omega^\varepsilon = \sigma_s^\varepsilon / \sigma_t^\varepsilon$ is,

$$\omega^\varepsilon = 1 - \varepsilon^2 \frac{\sigma_a}{\sigma_t}. \quad (6)$$

The probability that a collision is absorbing is $p_a = 1 - \omega$ and the TDL-scaled probability $p_a^\varepsilon = 1 - \omega^\varepsilon$ is,

$$p_a^\varepsilon = \varepsilon^2 \frac{\sigma_a}{\sigma_t}. \quad (7)$$

Let N_c be the number of collisions before a particle is absorbed, including the absorbing collision. Since every collision independently absorbs with probability p_a^ε , N_c is a geometric random variable. The average

number of collisions in a particle history is the expectation of N_c , $E[N_c]$, which is the inverse of Equation (7). Assume that the runtime of a Monte Carlo calculation τ is proportional to the product of the number of particles N_p and the average history length, $\tau \propto N_p E[N_c]$. Using the inverse of Equation (7) for $E[N_c]$,

$$\tau \propto N_p \frac{1}{\varepsilon^2} \frac{\sigma_t}{\sigma_a} \sim \varepsilon^{-2}. \quad (8)$$

Inducing the asymptotic regime by taking the limit of Equation (8) as ε goes to zero, we conclude that the runtime is unbounded. Can we bound it? We could make a diffusion approximation [6, 7, 8], we could iterate [10, 11, 12], or we could do both [13]. We analyze iteration. Let i be the iteration index and consider a single cycle of an iteration in which the scattering source is lagged,

$$\mathbf{\Omega} \cdot \nabla \psi^{(i+1)} + \sigma_t \psi^{(i+1)} = \frac{\sigma_s}{4\pi} \phi^{(i)} + q, \quad (9a)$$

$$\psi^{(i+1)}(\mathbf{x}, \mathbf{\Omega}) = \psi_{\text{inc}}(\mathbf{x}, \mathbf{\Omega}), \quad \mathbf{x} \in \partial\mathcal{D} \text{ and } \mathbf{\Omega} \cdot \mathbf{n} < 0, \quad (9b)$$

where $\phi^{(i)}$ is an approximate angle-integrated intensity computed during cycle i . One choice for $\phi^{(i)}$ is,

$$\phi^{(i)} = \int_{\mathbb{S}^2} \psi^{(i)} d\mathbf{\Omega}', \quad (10)$$

where the integral in Equation (10) is approximated using Monte Carlo. However, the resulting iteration would require $O(\varepsilon^{-2})$ cycles to converge. For the analysis, we assume that $\phi^{(i)}$ is computed in a way that makes the iteration converge quickly. For example, if $\phi^{(i)}$ is computed by solving Lewis and Miller's Second Moment system [14] using the hybrid Monte Carlo-deterministic Second Moment method of Pozulp et al. [12], the iteration converges in ten cycles or less, with fewer cycles required as $\varepsilon \rightarrow 0$. Thus, the order dependence of the number of cycles required to converge on ε is $\varepsilon^0 = 1$ (it is constant with respect to ε).

Monte Carlo methods for the source iteration process described by Equations (9a) and (9b) do not have scattering events. In an infinite medium, particles undergo a single collision (absorption), and the order dependence of the runtime of a single cycle on ε is $\varepsilon^0 = 1$. Assuming that the total runtime is proportional to the product of the number of particles, the average history length, and the number of cycles required to converge, we find that the order dependence of the total runtime on ε is $\varepsilon^0 = 1$.

Figure 1 shows a plot demonstrating numerical evidence in support of the order ε^{-2} and ε^0 runtimes for an infinite medium modeled in 2D-XY geometry with an 8×8 mesh of squares of equal size on $\mathcal{D} = [0, 1]^2$ using 1 million particles for different values of the TDL scaling parameter ε . The ε values are $2^0, 2^{-1}, 2^{-2}, \dots, 2^{-13}$. Whereas the runtime of standard Monte Carlo (MC) grows quadratically with the inverse of ε , the runtime of Monte Carlo source iteration (MC SI) is approximately constant. In this problem, the MC SI runtime decreases slightly as $\varepsilon \rightarrow 0$. This is because the facet crossing events disappear from the MC SI particle histories as the scaled mean free path $\lambda^\varepsilon = 1/\sigma_t^\varepsilon = \varepsilon/\sigma_t$ goes to 0 when $\varepsilon \rightarrow 0$.

Unfortunately, the bounded runtime achieved by sampling the scattering source through particle weights is accompanied by unbounded variance, as we will now show. The TDL scaling of Equation (9a) is,

$$\mathbf{\Omega} \cdot \nabla \psi^\varepsilon + \frac{\sigma_t}{\varepsilon} \psi^\varepsilon = \frac{\sigma_t - \sigma_a \varepsilon}{4\pi} \phi^\varepsilon + \varepsilon q, \quad (11)$$

where we suppressed the iteration index. Let $\hat{F}$ be a path-length estimator computed during a Monte Carlo solve of Equation (11). Pozulp and Haut [9] showed that the order dependence of the variance of $\hat{F}$ on ε is,

$$\text{Var}[\hat{F}] \sim \varepsilon^{2\gamma+1}, \quad (12)$$

where γ is the smallest exponent on the ε coefficient of the right-hand-side terms in Equation (11), so $\gamma = -1$. Substituting $\gamma = -1$ into Equation (12) gives,

$$\text{Var}[\hat{F}] \sim \varepsilon^{-1}, \quad (13)$$

which is unbounded as $\varepsilon \rightarrow 0$. Iteration exchanges unbounded runtime for unbounded variance. Fortunately, the variance can be bounded by cancelling the scattering source. Write $\psi^{(i+1)}$ as the sum of an isotropic term and an anisotropic term,

$$\psi^{(i+1)} = \frac{\phi^{(i)}}{4\pi} + \tilde{\psi}^{(i+1)}. \quad (14)$$

Substituting Equation (14) into Equations (9a) and (9b) preserves the left-hand-side operator ($\mathbf{\Omega} \cdot \nabla + \sigma_t$), but changes the right-hand side,

$$\mathbf{\Omega} \cdot \nabla \tilde{\psi}^{(i+1)} + \sigma_t \tilde{\psi}^{(i+1)} = -\frac{1}{4\pi}(\sigma_a \phi^{(i)} + \mathbf{\Omega} \cdot \nabla \phi^{(i)}) + q, \quad (15a)$$

$$\tilde{\psi}^{(i+1)}(\mathbf{x}, \mathbf{\Omega}) = \psi_{\text{inc}}(\mathbf{x}, \mathbf{\Omega}) - \frac{1}{4\pi} \phi^{(i)}(\mathbf{x}), \quad \mathbf{x} \in \partial\mathcal{D} \text{ and } \mathbf{\Omega} \cdot \mathbf{n} < 0. \quad (15b)$$

The unbounded scattering source is gone. The TDL scaling of Equation (15a) is,

$$\mathbf{\Omega} \cdot \nabla \tilde{\psi}^\varepsilon + \frac{\sigma_t}{\varepsilon} \tilde{\psi}^\varepsilon = -\frac{1}{4\pi}(\varepsilon \sigma_a \phi^\varepsilon + \mathbf{\Omega} \cdot \nabla \phi^\varepsilon) + \varepsilon q, \quad (16)$$

where we suppressed the iteration index. Let $\hat{G}$ be a path-length estimator computed during a Monte Carlo solve of Equation (16). The smallest exponent on the ε coefficient of the right-hand-side terms in Equation (16) is zero due to the $\mathbf{\Omega} \cdot \nabla \phi^\varepsilon$ term, which has the coefficient ε^0 . The variance of $\hat{G}$ is thus $\varepsilon^{2(0)+1}$,

$$\text{Var}[\hat{G}] \sim \varepsilon, \quad (17)$$

which is bounded because $\text{Var}[\hat{G}] \rightarrow 0$ as $\varepsilon \rightarrow 0$.

Figure 2 shows plots demonstrating numerical evidence in support of the order ε^{-1} variance for $\hat{F}$ and the order ε^1 variance for $\hat{G}$. In these plots, we estimated the variance of a path-length estimator for the angle-integrated intensity in a one-zone infinite medium with different values of the TDL scaling parameter ε . The angle-integrated intensity was calculated 600 times for each value of ε using a different seed for each calculation. The variance was estimated as the sum of the squared differences between the estimator realizations and the mean estimate, divided by 600:

$$\text{Var}[\hat{F}] \approx \frac{1}{600} \sum_{i=1}^{600} (\hat{F}_i - \bar{F})^2, \quad \text{Var}[\hat{G}] \approx \frac{1}{600} \sum_{i=1}^{600} (\hat{G}_i - \bar{G})^2, \quad (18)$$

where $\bar{F}$ and $\bar{G}$ are the mean estimates: $\bar{F} = (1/600) \sum_{i=1}^{600} \hat{F}_i$ and $\bar{G} = (1/600) \sum_{i=1}^{600} \hat{G}_i$.

Figure 3 presents solutions for a problem with both optically thick and optically thin regions, demonstrating that Monte Carlo methods with high variance in thick problems also have issues in problems with both thick and thin media. The quantity plotted in Figure 3 is a path-length estimator for the angle-integrated intensity computed before and after variance reduction for the Linearized Crooked Pipe from Olivier et al. [15], which is based on Graziani and LeBlanc's Crooked Pipe [16].




Figure 1. Runtime of standard Monte Carlo (MC) and the hybrid Monte Carlo-deterministic Second Moment (HSM) method. Reproduced with permission from the authors of [12].




Figure 2. Estimates of the variance of the standard path-length estimator for the angle-integrated intensity from a Monte Carlo source iteration method before variance reduction (a) and after variance reduction (b). Adapted with permission from the authors of [12].




Figure 3. Numerical solution before (a) and after (b) applying variance reduction similar to Equation (14). Reproduced with permission from the authors of [9].

3. THERMAL RADIATIVE TRANSFER ANALYSIS

We now consider the frequency-integrated (gray) equations of thermal radiative transfer (TRT). These equations are nonlinear, with the strongest nonlinearity appearing in the emission source, which is proportional to the fourth power of the temperature. The gray equations of thermal radiative transfer are,

$$\frac{1}{c} \frac{\partial I}{\partial t} + \mathbf{\Omega} \cdot \nabla I + \sigma I = \frac{c\sigma}{4\pi} U(T), \quad (19a)$$

$$C_v \frac{\partial T}{\partial t} = \sigma \int_{\mathbb{S}^2} I \, d\mathbf{\Omega} - c\sigma U(T), \quad (19b)$$

where c is the speed of light, $I = I(\mathbf{x}, \mathbf{\Omega}, t)$ is the radiation intensity, t is the time coordinate, σ is the absorption opacity, C_v is the heat capacity, $T = T(\mathbf{x}, t)$ is the temperature, and $U(T)$ is the function,

$$U(T) = aT^4, \quad a = \frac{8\pi^5 k^4}{15h^3 c^3},$$

where h and k are Planck's constant and Boltzmann's constant, respectively. Equations (19a) and (19b) are subject to an inflow boundary condition on I and initial conditions on I and T . The TDL scaling of the TRT problem data by $\varepsilon \in (0, 1]$ from Larsen, Pomraning, and Badham [17] and Larsen and Densmore [18] is,

$$c^\varepsilon = c/\varepsilon, \quad C_v^\varepsilon = \varepsilon C_v, \quad \sigma^\varepsilon = \sigma/\varepsilon, \quad (20)$$

where $c^\varepsilon = c/\varepsilon$ is the scaling of the time-derivative coefficient, $1/c \mapsto \varepsilon/c$, but is not used anywhere else (c in the emission is not scaled). The unscaled problem data are independent of ε and thus order unity,

$$c \sim 1, \quad C_v \sim 1, \quad \sigma \sim 1. \quad (21)$$

Therefore, the order dependence of the scaled problem data on ε is,

$$c^\varepsilon \sim \varepsilon^{-1}, \quad C_v^\varepsilon \sim \varepsilon, \quad \sigma^\varepsilon \sim \varepsilon^{-1}. \quad (22)$$

Applying the scaling in Equation (20) to the data in Equation (19a) produces a new transport equation in terms of I^ε and T^ε ,

$$\frac{\varepsilon}{c} \frac{\partial I^\varepsilon}{\partial t} + \mathbf{\Omega} \cdot \nabla I^\varepsilon + \frac{\sigma}{\varepsilon} I^\varepsilon = \frac{c}{4\pi} \frac{\sigma}{\varepsilon} U(T^\varepsilon), \quad (23)$$

Using the same logic as in Section 2, the runtime of Monte Carlo methods for Equation (23) is bounded because it is order unity in the TDL, but the variance is unbounded because it is order ε^{-1} . Thus, the variance of Monte Carlo methods that solve Equation (19a) is unbounded in the TDL. This includes explicit Monte Carlo, Brooks' method [3], and N'Kaoua's method [4]. Can we bound it? Szoke and Brooks' difference formulation [5] bounds the variance by replacing the unbounded emission source with bounded ones. The difference formulation for the gray TRT equations is derived by writing $I_\nu = B_\nu + \tilde{I}_\nu$ and integrating over all photon frequencies to get,

$$I = \frac{c}{4\pi} U(T) + \tilde{I}. \quad (24)$$

Substituting Equation (24) into Equation (19a) preserves the left-hand-side operator ($\frac{1}{c} \frac{\partial}{\partial t} + \mathbf{\Omega} \cdot \nabla + \sigma$), but changes the right-hand side,

$$\left(\frac{1}{c} \frac{\partial}{\partial t} + \mathbf{\Omega} \cdot \nabla + \sigma \right) \tilde{I} = -\frac{aT^3}{\pi} \left(\frac{\partial T}{\partial t} + c\mathbf{\Omega} \cdot \nabla T \right), \quad (25)$$

where we used $\frac{\partial U}{\partial t} = 4aT^3 \frac{\partial T}{\partial t}$ and $\nabla U = 4aT^3 \nabla T$. Scaling the problem data according to Equation (20) gives,

$$\left(\frac{\varepsilon}{c} \frac{\partial}{\partial t} + \mathbf{\Omega} \cdot \nabla + \frac{\sigma}{\varepsilon} \right) \tilde{I}^\varepsilon = - \frac{a(T^\varepsilon)^3}{\pi} \left(\frac{\partial T^\varepsilon}{\partial t} + c \mathbf{\Omega} \cdot \nabla T^\varepsilon \right). \quad (26)$$

The smallest exponent on the ε coefficient of the right-hand-side terms in Equation (26) is zero due to the $\mathbf{\Omega} \cdot \nabla T^\varepsilon$ term, so the variance of Monte Carlo TRT methods using the difference formulation is bounded.

Now consider Fleck and Cummings' linearized equations. Applying Fleck and Cummings' linearization to the gray TRT system of Equations (19a) and (19b) gives,

$$\frac{1}{c} \frac{\partial I}{\partial t} + \mathbf{\Omega} \cdot \nabla I + \sigma I = \frac{1-f}{4\pi} \sigma \int_{\mathbb{S}^2} I \, d\mathbf{\Omega} + \frac{f}{4\pi} c \sigma U(T), \quad (27a)$$

$$C_v \frac{\partial T}{\partial t} = f \sigma \int_{\mathbb{S}^2} I \, d\mathbf{\Omega} - f c \sigma U(T), \quad (27b)$$

where f is the Fleck factor,

$$f = \frac{1}{1 + \alpha \beta \sigma c \Delta t}. \quad (27c)$$

In the Fleck factor, $\alpha \in [0, 1]$ is the implicitness factor and β is the strength of the coupling between the material and the radiation,

$$\beta = \beta(\mathbf{x}, t) = \frac{\partial U(T)}{\partial U_m} = \frac{dU(T)/dT}{dU_m/dT} = \frac{4aT^3}{C_v}. \quad (27d)$$

Applying the scaling in Equation (20) to Equation (27a) gives,

$$\frac{\varepsilon}{c} \frac{\partial I^\varepsilon}{\partial t} + \mathbf{\Omega} \cdot \nabla I^\varepsilon + \frac{\sigma}{\varepsilon} I^\varepsilon = \frac{1-f^\varepsilon}{4\pi} \frac{\sigma}{\varepsilon} \int_{\mathbb{S}^2} I^\varepsilon \, d\mathbf{\Omega} + \frac{f^\varepsilon}{4\pi} c \frac{\sigma}{\varepsilon} U(T^\varepsilon), \quad (28)$$

where f^ε is the TDL scaling of the Fleck factor defined in Equation (27c),

$$f^\varepsilon = \frac{1}{1 + \alpha \beta \sigma c \Delta t / \varepsilon^2} = \frac{\varepsilon^2}{\varepsilon^2 + \alpha \beta \sigma c \Delta t}, \quad (29)$$

so the order dependence of the scaled Fleck factor on ε is $f^\varepsilon \sim \varepsilon^2$, provided that $\alpha \neq 0$.

Following the logic in Section 2, the runtime of Fleck and Cummings' Monte Carlo method applied to Fleck and Cummings' linearized equations Equation (28) is unbounded because it is order ε^{-2} . Monte Carlo source iteration methods that solve Fleck and Cummings' linearized equations without effective scattering events by incorporating the contribution of the effective scattering source into the particle weights, such as Lam et al. [10], have unbounded variance in the TDL. However, Monte Carlo methods that solve Fleck and Cummings' linearized equations and make a diffusion approximation like Fleck and Canfield [6], Gentile [7], and Densmore et al. [8] have bounded variance and runtime. This is because when ε is small, a diffusion equation is solved (using a Monte Carlo interpretation) instead of a transport equation. The diffusion equation is derived by first integrating Equation (27a) over all directions,

$$\frac{1}{c} \frac{\partial \phi}{\partial t} + \nabla \cdot \mathbf{F} + f \sigma \phi = f c \sigma U(T), \quad (30)$$

where we used,

$$\phi = \int_{\mathbb{S}^2} I \, d\mathbf{\Omega}, \quad \text{and} \quad \mathbf{F} = \int_{\mathbb{S}^2} \mathbf{\Omega} I \, d\mathbf{\Omega}. \quad (31)$$

Using Fick’s law to write $\mathbf{F}$ in terms of ϕ as,

$$\mathbf{F} = -\frac{1}{3\sigma} \nabla \phi, \quad (32)$$

and substituting Equation (32) into Equation (30) gives a diffusion equation for ϕ ,

$$\frac{1}{c} \frac{\partial \phi}{\partial t} - \nabla \cdot \left(\frac{1}{3\sigma} \nabla \phi \right) + f\sigma\phi = fc\sigma U(T). \quad (33)$$

Applying the scaling in Equation (20) to Equation (33) gives a new diffusion equation for ϕ^ε ,

$$\frac{\varepsilon}{c} \frac{\partial \phi^\varepsilon}{\partial t} - \nabla \cdot \left(\frac{\varepsilon}{3\sigma} \nabla \phi^\varepsilon \right) + \frac{f^\varepsilon \sigma}{\varepsilon} \phi^\varepsilon = \frac{f^\varepsilon c \sigma}{\varepsilon} U(T^\varepsilon). \quad (34)$$

The order of each term in Equation (34) is ε because $f^\varepsilon \sim \varepsilon^2$ (with the assumption that $\alpha \neq 0$). Division by ε yields an equation where each term is order $\varepsilon^0 = 1$. Thus, the exponent on the ε coefficient in the emission term is zero, and so the variance of Monte Carlo solutions to Equation (34) is bounded (it is order ε). The runtime is bounded because the scattering ratio is zero. Consequently, Monte Carlo methods that solve Fleck and Cummings’ linearized equations in optically thin media and make a diffusion approximation in optically thick media have bounded variance and runtime in the thick diffusion limit¹.

4. CONCLUSION

By applying the thick diffusion limit (TDL) scaling used by Larsen, Morel, and Miller [1], we found that if the leading sampled source in the TDL-scaled equation grows as the TDL parameter $\varepsilon \rightarrow 0$, then representing that source in particle weights yields an unbounded path-length-estimator variance, whereas representing it through repeated collision events yields an unbounded expected work per history. We presented numerical runtimes in Figure 1 that are consistent with the runtimes predicted by our analysis, and numerical variances in Figure 2 that are also consistent with our predictions.

Our analysis predicts the unbounded runtime of Monte Carlo methods that solve Fleck and Cummings’ linearized thermal radiative transfer (TRT) equations using effective scattering events. The analysis predicts unbounded variance for Brooks’ method [3] and N’Kaoua’s method [4]. It predicts bounded variance for Szoke and Brooks’ difference formulation [5]. It predicts unbounded variance for Monte Carlo methods that solve Fleck and Cummings’ linearized equations without effective scattering events by incorporating the contribution of the effective scattering source into the particle weights, such as Lam et al. [10]. It also predicts that Monte Carlo methods that solve Fleck and Cummings’ linearized equations and make a diffusion approximation, such as Fleck and Canfield [6], Gentile [7], and Densmore et al. [8], have bounded variance and runtime.

ACKNOWLEDGEMENTS

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-CONF-2024263.

REFERENCES

- [1] E. W. Larsen, J. E. Morel, and W. F. Miller Jr. “Asymptotic solutions of numerical transport problems in optically thick, diffusive regimes.” *J Comput Phys*, **volume 69**(2), pp. 283–324 (1987).

¹A more thorough analysis would examine intermediate optical depths, where it is unclear whether it is better to solve a transport equation or a diffusion equation, due to the relatively higher cost of the former and the relatively lower accuracy of the latter.

- [2] J. A. Fleck and J. D. Cummings. “An implicit Monte Carlo scheme for calculating time and frequency dependent nonlinear radiation transport.” *J Comput Phys*, **volume 8**, pp. 313–342 (1971).
- [3] E. D. Brooks. “Symbolic implicit Monte Carlo.” *J Comput Phys*, **volume 83**(2), pp. 433–446 (1989).
- [4] T. N’Kaoua. “Solution of the Nonlinear Radiative Transfer Equations by a Fully Implicit Matrix Monte Carlo Method Coupled with the Rosseland Diffusion Equation via Domain Decomposition.” *SIAM Journal on Scientific and Statistical Computing*, **volume 12**(3), pp. 505–520 (1991).
- [5] A. Szoke and E. D. Brooks III. “The transport equation in optically thick media.” *Journal of Quantitative Spectroscopy and Radiative Transfer*, **volume 91**(1), pp. 95–110 (2005).
- [6] 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 Comput Phys*, **volume 54**, pp. 508–523 (1984).
- [7] N. Gentile. “Implicit Monte Carlo diffusion - An acceleration method for Monte Carlo time-dependent radiative transfer simulations.” *J Comput Phys*, **volume 172**, pp. 543–571 (2001).
- [8] J. D. Densmore, T. J. Urbatsch, T. M. Evans, and M. W. Buksas. “A hybrid transport-diffusion method for Monte Carlo radiative-transfer simulations.” *J Comput Phys*, **volume 222**(2), pp. 485–503 (2007).
- [9] M. Pozulp and T. Haut. “A Proof of the Asymptotic Variance of Path Length Estimators for Single-Collision Monte Carlo Source Iteration in the Thick Diffusion Limit.” *J Comput Theor Transport*, **volume 55**(3), pp. 271–294 (2026).
- [10] A. Q. Lam, T. S. Palmer, T. A. Brunner, and R. M. Vega. “A Monte Carlo Thermal Radiative Transfer Solver with Nonlinear Elimination.” *J Comput Theor Transport*, **volume 52**(3), pp. 221–245 (2023).
- [11] S. Pasmann, I. Variansyah, C. T. Kelley, and R. McClarren. “A Quasi-Monte Carlo Method With Krylov Linear Solvers for Multigroup Neutron Transport Simulations.” *Nuclear Science and Engineering*, **volume 197**(6), pp. 1159–1173 (2023).
- [12] M. Pozulp, T. Haut, P. Brantley, and S. Olivier. “A hybrid Monte Carlo-deterministic second moment method with efficient variance reduction.” *J Comput Phys*, **volume 565**, p. 115121 (2026).
- [13] H. Park, J. Densmore, A. Wollaber, D. Knoll, and R. Rauenzahn. “Monte Carlo Solution Methods in a Moment-Based Scale-Bridging Algorithm for Thermal Radiative Transfer Problems: Comparison with Fleck and Cummings.” In *Proceedings of M&C 2013*. Sun Valley, Idaho (2013).
- [14] E. Lewis and W. Miller. “A comparison of P1 synthetic acceleration techniques.” *Transactions of the American Nuclear Society*, **volume 23** (1976).
- [15] S. Olivier, W. Pazner, T. S. Haut, and B. C. Yee. “A family of independent Variable Eddington Factor methods with efficient preconditioned iterative solvers.” *J Comput Phys*, **volume 473**, p. 111747 (2023).
- [16] F. Graziani and J. LeBlanc. “The Crooked Pipe Test Problem.” Technical report, UCRL-MI-143393 (2000).
- [17] E. Larsen, G. Pomraning, and V. Badham. “Asymptotic analysis of radiative transfer problems.” *Journal of Quantitative Spectroscopy and Radiative Transfer*, **volume 29**(4), pp. 285–310 (1983).
- [18] J. D. Densmore and E. W. Larsen. “Asymptotic equilibrium diffusion analysis of time-dependent Monte Carlo methods for grey radiative transfer.” *J Comput Phys*, **volume 199**(1), pp. 175–204 (2004).