LLNL-PROC-848841

LAWRENCE LIVERMORE NATIONAL LABORATORY

An Implicit Monte Carlo Acceleration Scheme

M. M. Pozulp, T. S. Haut, P. S. Brantley, J. L. Vujic

May 8, 2023

The International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C 2023)
Niagara Falls, Canada
August 13, 2023 through August 17, 2023

An Implicit Monte Carlo Acceleration Scheme

M. Pozulp[1,2], T. Haut[1], P. Brantley[1], and J. Vujic[2]

[1]Lawrence Livermore National Laboratory [2]University of California, Berkeley
{pozulp1, haut3, brantley1}@llnl.gov vujic@berkeley.edu

ABSTRACT

The implicit Monte Carlo (IMC) equations approximate the photon absorption and re-emission process with an effective scattering event. Calculation runtimes are dominated by photons traversing optically thick media, in which the IMC photons undergo numerous effective scattering events separated by relatively short paths. We rectify this problem by employing a transport-corrected diffusion approximation in a hybrid method that unconditionally eliminates scattering events everywhere in phase space. We demonstrate that the computational cost of our scheme is independent of the mean free path.

KEYWORDS: IMC, Hybrid, No Scattering

1. INTRODUCTION

Implicit Monte Carlo (IMC) [1] is a popular algorithm for solving the thermal radiative transfer (TRT) equations. IMC is a linearization of TRT that contains an "effective scattering" term which approximates the photon absorption and re-emission process: if an IMC photon undergoes an effective scattering event, that means that the physical photon it represents was absorbed and re-emitted. IMC is most expensive in optically thick highly-absorbing media, in which IMC photons undergo numerous effective scattering events separated by relatively short paths between collisions. The IMC photons have long life histories, which results in extended calculation runtimes.

The earliest solution for improving IMC calculation runtimes was the Random Walk (RW) method of Fleck and Canfield [2], which requires that certain conditions are satisfied. Spatial and temporal refinement, which increase cost, dampen the RW speedup because the conditions for its use are less likely satisfied in smaller zones and with shorter timesteps. More recent solutions, like implicit Monte Carlo Diffusion (IMD) [3], replace transport zones with diffusion zones at spatial locations at which the diffusion solution is appropriate. The calculation in the remaining transport zones proceeds as before, whereas the solution in the diffusion zones is determined by solving a linear system arising from a discretization of the diffusion equation using a Monte Carlo technique

A similar problem in deterministic transport is the slow convergence of source iteration in the aforementioned optically thick highly-absorbing media. Solutions to slow iterative convergence typically rely on using some diffusion-based method along with the transport solve [4]. One class of solution is the Variable Eddington Factor (VEF) method, also known as Quasidiffusion [5]. VEF is a moment method in which the diffusion equation includes a nonlinear transport-correction. A second moment method (SMM) which includes a linear transport correction was later proposed by Lewis and Miller [6].

Chacón *et al.* [7] surveyed the 60 year history of moment methods, demonstrated moment method acceleration of IMC, and acknowledged the presence of statistical noise caused by the combination of Monte Carlo (MC) methods with moment methods. Unlike Chacón *et al.*, we use SMM, which is a different moment method than the method that they used, but we believe that our solution resembles their solution in many important ways. For example, our choice to combine MC with SMM also introduces statistical noise. A naïve usage of MC with SMM involves numerical differentiation of a noisy quantity, which can amplify the noise. The authors in Hammer, Park, and Chacón [8] avoid this issue by using a deterministic particle method that is similar to the method of characteristics. In contrast, our strategy to address the noise amplification, which we hope to present in a follow up publication, will employ a sophisticated MC tally which avoids the differentiation. In this paper, we examine the naïve usage of MC with SMM in order to quantify performance of the scheme with respect to the mean free path, and assess the effect of noise amplification on solution quality.

Our solution also resembles that of Cooper and Larsen [9], where MC transport is interrupted by a VEF solve used to compute weight windows. Unlike Cooper and Larsen, we use SMM instead of VEF and we use it to eliminate effective scattering events rather than compute weight windows. Our solution is a “hybrid” method in that we use a stochastic method (MC) to solve a modified transport equation and a deterministic method to solve the moment equations. We demonstrate its value by showing that 1) the cost of the scheme is independent of mean free path and 2) the scheme is accurate in the diffusion limit. We verify our scheme by showing that 3) the numerical solution has an error like $\frac{C_0}{\sqrt{N}} + C_1 h$, where N is the number of IMC photons, h is the mesh width, and the constants C_0 and C_1 are independent of the mean free path.

2. METHOD

2.1. Derivation of the SMM equations

The following derivation of the SMM equations is a reproduction of the derivation by Olivier [10] using the same notation and algebra. Just like Olivier, we use a monoenergetic, steady-state, fixed-source, linear transport equation with isotropic scattering as a simplified model for a single time step of TRT:

$$\mathbf{\Omega} \cdot \nabla \psi + \sigma_t \psi = \frac{\sigma_s}{4\pi} \int \psi \, \mathrm{d}\Omega' + q \,, \quad \mathbf{x} \in \mathcal{D} \tag{1a}$$

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

Here, $\mathbf{x} \in \mathbb{R}^{\dim}$ and $\mathbf{\Omega} \in \mathbb{S}^2$ are the spatial and angular variables, $\psi(\mathbf{x}, \mathbf{\Omega})$ is the angular flux, $\mathcal{D} \subset \mathbb{R}^{\dim}$ the spatial domain and $\partial\mathcal{D}$ its boundary and $\mathbf{n}$ the outward unit normal to the boundary, $\sigma_t(\mathbf{x})$ and $\sigma_s(\mathbf{x})$ are the total and scattering macroscopic cross sections, $q(\mathbf{x}, \mathbf{\Omega})$ is the fixed-source, and $\bar{\psi}(\mathbf{x}, \mathbf{\Omega})$ the inflow boundary function. The zeroth and first angular moments of the transport equation are

$$\nabla \cdot \boldsymbol{J} + \sigma_a \varphi = Q_0 \,, \tag{2a}$$

$$\nabla \cdot \mathbf{P} + \sigma_t \boldsymbol{J} = \boldsymbol{Q}_1 \,, \tag{2b}$$

where $\sigma_a(\mathbf{x}) = \sigma_t(\mathbf{x}) - \sigma_s(\mathbf{x})$ is the absorption macroscopic cross section, $Q_0 = \int q \, \mathrm{d}\Omega$ and $\boldsymbol{Q}_1 = \int \mathbf{\Omega} q \, \mathrm{d}\Omega$ are the zeroth and first angular moments of the fixed-source and φ, $\boldsymbol{J}$, and $\mathbf{P}$ are the zeroth, first, and second angular moments of the angular flux. This is an unclosed system of $1 + 3 = 4$ equations (because Eq. (2a) is a scalar equation and Eq. (2b) is a vector equation) and $1 + 3 + 6 = 10$ unknowns (because φ is a scalar, $\boldsymbol{J}$ is a vector, and $\mathbf{P}$ is a symmetric tensor). Close the system using $\mathbf{P} = \mathbf{T} + \frac{1}{3}\mathbf{I}\varphi$ where $\mathbf{I}$ is the identity matrix to get

$$\nabla \cdot \boldsymbol{J} + \sigma_a \varphi = Q_0, \tag{3a}$$

$$\frac{1}{3}\nabla\varphi + \sigma_t \boldsymbol{J} = \boldsymbol{Q}_1 - \nabla \cdot \mathbf{T}, \tag{3b}$$

where

$$\mathbf{T} = \int \boldsymbol{\Omega} \otimes \boldsymbol{\Omega} \psi \, \mathrm{d}\Omega - \frac{1}{3}\mathbf{I} \int \psi \, \mathrm{d}\Omega \tag{3c}$$

is the correction tensor. To derive the boundary condition, let $J_n^{\pm} = \int_{\boldsymbol{\Omega}\cdot\mathbf{n}\gtrless 0} \boldsymbol{\Omega} \cdot \mathbf{n} \, \psi \, \mathrm{d}\Omega$ denote the partial currents and use them to express the net current, then manipulate the expression

$$\begin{aligned} \boldsymbol{J} \cdot \mathbf{n} &= J_n^- + J_n^+ \\ &= 2J_n^- + (J_n^+ - J_n^-) \\ &= 2J_n^- + \int |\boldsymbol{\Omega} \cdot \mathbf{n}| \, \psi \, \mathrm{d}\Omega \\ &= 2J_{\text{in}} + B \ , \end{aligned} \tag{4}$$

where $J_{\text{in}} = \int_{\boldsymbol{\Omega}\cdot\mathbf{n}<0} \boldsymbol{\Omega} \cdot \mathbf{n} \, \bar{\psi} \, \mathrm{d}\Omega$ is the incoming partial current computed from the inflow boundary function, $\bar{\psi}$, and $B(\psi) = \int |\boldsymbol{\Omega} \cdot \mathbf{n}| \psi \ \mathrm{d}\Omega$. Substitute $B = \beta + \frac{1}{2}\varphi$, where

$$\beta = \int |\boldsymbol{\Omega} \cdot \mathbf{n}| \psi \ \mathrm{d}\Omega - \frac{1}{2} \int \psi \ \mathrm{d}\Omega \tag{5}$$

is the boundary factor, to get the SMM boundary condition

$$\boldsymbol{J} \cdot \mathbf{n} = \frac{1}{2}\varphi + 2J_{\text{in}} + \beta \ . \tag{6}$$

The SMM equations are thus

$$\nabla \cdot \boldsymbol{J} + \sigma_a \varphi = Q_0 \, , \quad \mathbf{x} \in \mathcal{D} \, , \tag{7a}$$

$$\frac{1}{3}\nabla\varphi + \sigma_t \boldsymbol{J} = \boldsymbol{Q}_1 - \nabla \cdot \mathbf{T} \, , \quad \mathbf{x} \in \mathcal{D} \, , \tag{7b}$$

$$\boldsymbol{J} \cdot \mathbf{n} = \frac{1}{2}\varphi + 2J_{\text{in}} + \beta \, , \quad \mathbf{x} \in \partial\mathcal{D} \, . \tag{7c}$$

We solve the second-order form, found by eliminating the current:

$$-\nabla \cdot \frac{1}{3\sigma_t}\nabla\varphi + \sigma_a\varphi = Q_0 - \nabla \cdot \frac{\boldsymbol{Q}_1}{\sigma_t} + \nabla \cdot \frac{1}{\sigma_t}\nabla \cdot \mathbf{T} \, . \tag{8}$$

The iteration, depicted in Fig. 1, consists of a transport solve with a scattering source computed using the scalar flux from the moment solve φ, and a moment solve with a transport correction computed using the angular flux from the transport solve ψ. SMM replaces a single IMC step with multiple steps which are cheaper because the MC photons do not undergo scattering events.

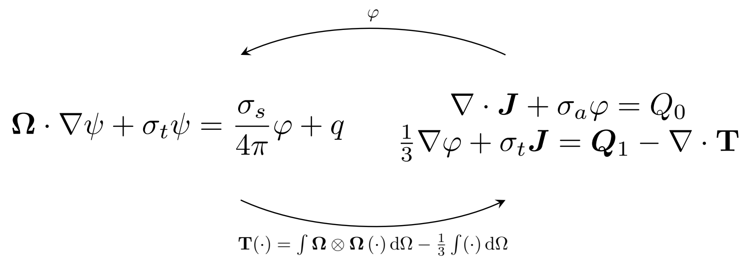

Figure 1: SMM iteration, reproduced from Olivier [10].

2.2. Implementation details

We use MC for the transport solve, so the correction tensor $\mathbf{T}$ has MC noise, and in Eq. (8) we compute two derivatives of $\mathbf{T}$. How to avoid the amplification of MC noise in $\mathbf{T}$ caused by differentiation is a research question which we are investigating. The boundary factor β also has MC noise but we do not differentiate β. We use three tallies to compute $\mathbf{T}$ and β. Let $\hat{\mathbf{T}} = \hat{\mathbf{P}} - \frac{1}{3}\mathbf{I}\hat{\phi}$ and $\hat{\beta} = \hat{K} - \frac{1}{2}\hat{\phi}$ be MC estimators for $\mathbf{T}$ and β, where

$$\hat{\phi} = \frac{1}{V}\sum_i d_i w_i \ , \quad \hat{\mathbf{P}} = \frac{1}{V}\sum_i \mathbf{\Omega}_i \otimes \mathbf{\Omega}_i d_i w_i \ , \quad \text{and} \quad \hat{K} = \frac{1}{V}\sum_i |\mathbf{\Omega}_i \cdot \mathbf{n}| d_i w_i \tag{9}$$

are path-length estimators, so the sum is over paths of length d_i in the volume V, and path i is traversed by an MC photon with weight w_i moving in the direction $\mathbf{\Omega}_i$. Thus, $\hat{\mathbf{T}}$ and $\hat{\beta}$ are piecewise-constant tallies computed in each zone, and each boundary zone, respectively, of a tally mesh. The tally mesh is a uniform tessellation of the spatial domain $\mathcal{D}$. The MC photons do not undergo scattering events. Instead, we include the effects of scattering during sourcing, in which the fixed-source $q(\mathbf{x}, \mathbf{\Omega})$ is augmented by the scattering source $\varphi(\mathbf{x})\sigma_s/4\pi$. We enforce the global requirement that the sum of the MC photon weights $\sum_i^N w_i$ must equal the integral of the scattering-and-fixed source $\int\int(\varphi\sigma_s/4\pi + q)\ \mathrm{d}\Omega\,\mathrm{d}\mathbf{x}$ by enforcing it locally within each zone of the source mesh. The source mesh is a uniform tessellation of $\mathcal{D} \times \mathbb{S}^2$ phase space. The sourcing algorithm is:

i. uniformly sample $\mathbf{x}$ and $\mathbf{\Omega}$ for each MC photon,

ii. compute $N_{\mathbf{x}} \times N_{\mathbf{\Omega}}$ volume integrals Φ_{ij} of the scattering-and-fixed source on the source mesh, where $N_{\mathbf{x}}$ and $N_{\mathbf{\Omega}}$ are the number of spatial and angular zones in the source mesh, and i and j are the spatial and angular zone indices,

iii. set the weight of each MC photon in zone ij to Φ_{ij}/N^{ij}, where N^{ij} is the number of MC photons sourced into zone ij.

We use bilinear interpolation in step ii to interpolate the moment flux φ to the source mesh zone centers, which is where we evaluate the scattering-and-fixed source function in a 1-point integration rule to compute the volume integrals. Our sourcing algorithm produces equal weight intra-zone MC photons, but inter-zone MC photon weights may differ. A more sophisticated sourcing algorithm, or a post-sourcing global variance reduction step like splitting high-weight particles and

Russian rouletting low-weight particles, would provide a more uniform weight distribution. We track MC photons until they get absorbed by the medium or escape through the spatial domain boundary $\partial\mathcal{D}$.

We use finite difference (FD) for the moment solve, so we need to interpolate the zonal $\hat{\mathbf{T}}$ and $\hat{\beta}$ tallies to the FD nodes in order to form the right-hand side of Eq. (8). We simply assign tally values to FD nodes without averaging, but we could use spatial averaging, as depicted in Fig. 2, and we plan to investigate this improvement in the future. Our discrete moment system is sparse but also small, so we invert with LU factorization. If we were solving bigger problems we could use a robust iterative solver like algebraic multigrid pre-conditioned conjugate gradient because the system is symmetric positive definite.

We think that the same process which the nuclear engineering community uses for controlling eigenvalue iteration would work well for the iteration depicted in Fig. 1. That is, discard a few "initial" iterates, average subsequent iterates, and terminate the iteration once the average satisfies a convergence criterion involving the averaged iterates. Here, we simply terminate after a fixed number of iterations and use the final iterate as the solution. The fixed number of iterations is an empirical, problem-dependent quantity that we choose to be at least three and never more than twenty.

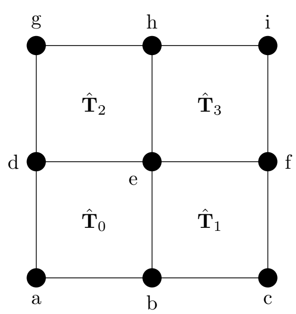

(a) Tallies

$$\begin{matrix} \hat{\mathbf{T}}_2, & \hat{\mathbf{T}}_2, & \hat{\mathbf{T}}_3 \\ \hat{\mathbf{T}}_0, & \hat{\mathbf{T}}_0, & \hat{\mathbf{T}}_1 \\ \hat{\mathbf{T}}_0, & \hat{\mathbf{T}}_0, & \hat{\mathbf{T}}_1 \end{matrix}$$

(b) No averaging

$$\begin{matrix} \hat{\mathbf{T}}_2, & \frac{\hat{\mathbf{T}}_2+\hat{\mathbf{T}}_3}{2}, & \hat{\mathbf{T}}_3 \\ \frac{\hat{\mathbf{T}}_0+\hat{\mathbf{T}}_2}{2}, & \frac{1}{4}\sum_{i=1}^{4}\hat{\mathbf{T}}_i, & \frac{\hat{\mathbf{T}}_1+\hat{\mathbf{T}}_3}{2} \\ \hat{\mathbf{T}}_0, & \frac{\hat{\mathbf{T}}_0+\hat{\mathbf{T}}_1}{2}, & \hat{\mathbf{T}}_1 \end{matrix}$$

(c) Averaging

Figure 2: Interpolation of the zonal MC estimator $\hat{\mathbf{T}}$ to the finite difference nodes

3. RESULTS

We refer to our scheme as Nose to distinguish it from IMC. Nose is an abbreviation of "No scattering events". Nose gives two values for the scalar flux, one from the transport solve and one from the moment solve. We refer to these as the "MC Flux" and the "SMM Flux", respectively. Both should converge to the same solution in the limit as we refine the mesh and increase the number of MC particles. We use uniform meshes for the source mesh, tally mesh, and moment mesh in which the width of each zone is $h = 1/Z$ where Z^{dim} is the number of zones and the spatial domain is the unit cube. We ran 1D calculations with one spatial coordinate and one angular coordinate, and 2D calculations with two spatial coordinates and two angular coordinates.

To exercise the thick diffusion limit, we ran a 1D fixed-source problem with vacuum boundaries and uniform material properties set to $\sigma_t = \frac{1}{\epsilon}$, $\sigma_a = \epsilon$, $\sigma_s = \frac{1}{\epsilon} - \epsilon$, and $q = \epsilon$ with $\epsilon \in (0, 1]$ as in [11]. Figure 3 demonstrates that Nose converges to the diffusion solution as $\epsilon \to 0$ even when the mesh is large relative to the mean free path. Figure 3 also demonstrates how Nose gives two values for the scalar flux, as previously mentioned. Figure 4 shows that Nose runtime is independent of ϵ whereas IMC runtime, which is $O(1/\epsilon^2)$, explodes as $\epsilon \to 0$.

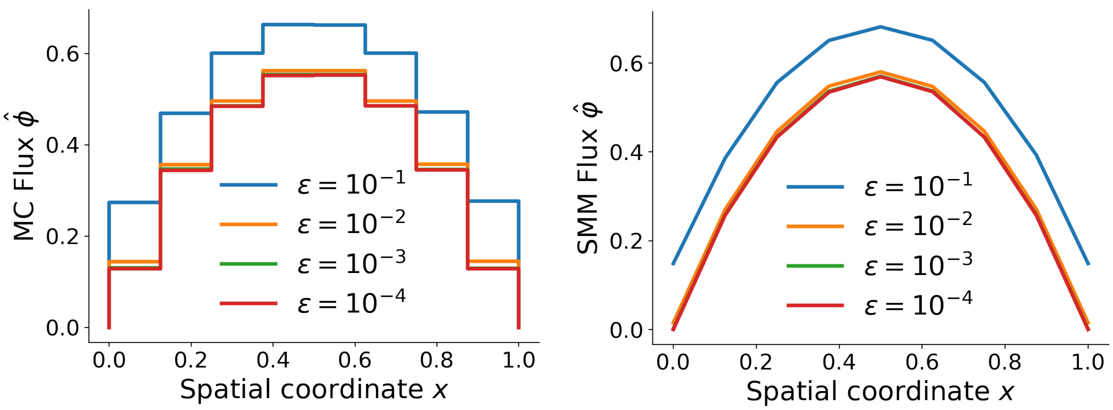

(a) MC Flux $\hat{\phi}(x)$ for $N = 10^6$ **(b) SMM Flux $\hat{\varphi}(x)$ for $N = 10^6$**

Figure 3: Flux in the thick diffusion limit for 10^6 photons and 8 spatial zones

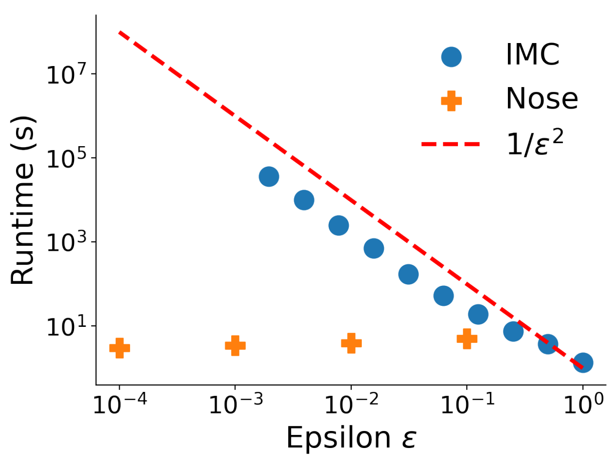

Figure 4: Runtime for thick diffusion limit calculation

To verify our method, we use the method of manufactured solutions (MMS). In 1D, we set the MMS solution to $\psi = (1 + \mu + \mu^2)\sin \pi x$, and in 2D we set it to

$$\psi = \alpha + \mathbf{\Omega} \cdot \boldsymbol{\beta} + \mathbf{\Omega} \otimes \mathbf{\Omega} : \mathbf{\Theta}, \quad \alpha = \sin \pi x \sin \pi y, \quad \boldsymbol{\beta} = \begin{pmatrix} \alpha \\ \alpha \end{pmatrix}, \quad \mathbf{\Theta} = \begin{pmatrix} \alpha & \alpha \\ \alpha & \alpha \end{pmatrix}. \tag{10}$$

We set $\sigma_a = \sigma_s = 1$. We seek the hypothesized error of $\frac{C_0}{\sqrt{N}} + C_1 h$. We test the first term by fixing h small and increasing N. We test the second term by fixing N large and decreasing h. Figure 5 shows that the error has the hypothesized dependence on N and h in 1D but not in 2D. The 1D calculation error in Fig. 5 (a) flatlines after 4^{10} because N is big enough that $\frac{C_0}{\sqrt{N}} < C_1 h$.

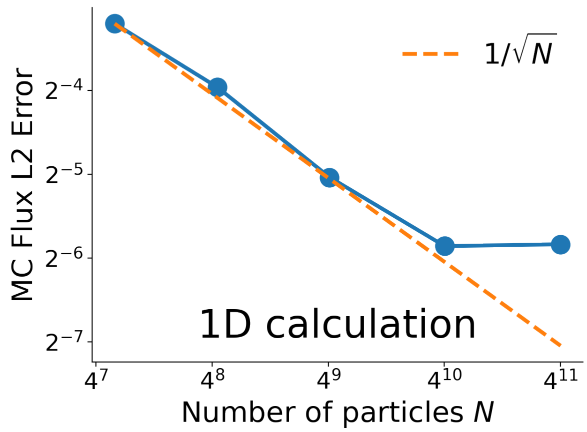

(a) Error as N increases in 1D

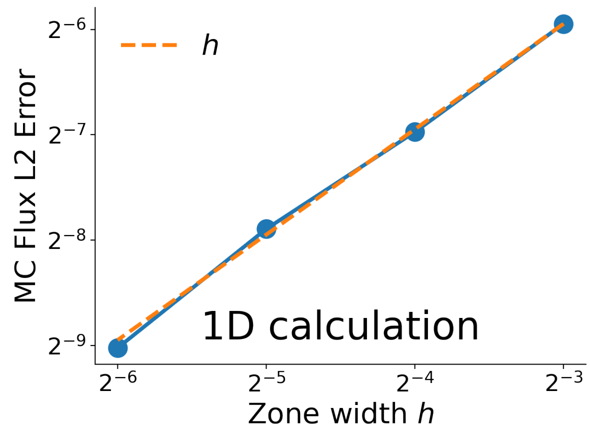

(b) Error as h decreases in 1D

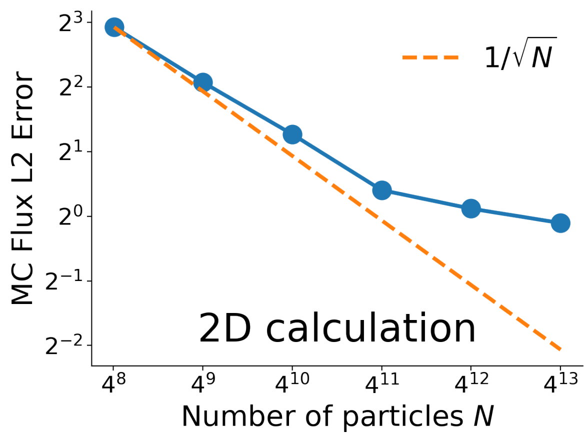

(c) Error as N increases in 2D

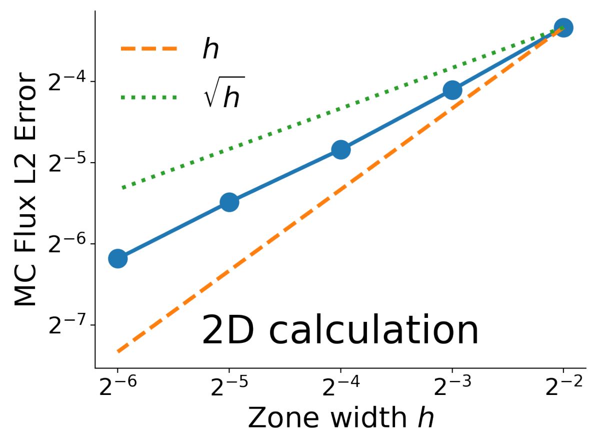

(d) Error as h decreases in 2D

Figure 5: The MC Flux error as N increases or h decreases

We think that the noise in $\hat{\mathbf{T}}$ is causing the 2D errors to exceed the hypothesized values. As previously mentioned, how to avoid the amplification of noise in $\hat{\mathbf{T}}$ caused by differentiation is a research question which we are investigating. Why is 1D unaffected? We attribute the resilience of 1D to operator symmetry. The div-grad and div-div operators on the left and right sides of Eq. (8), respectively, are different operators in 2D. In 1D they are both second-order ordinary derivatives. We think that this symmetry may avert the noise amplification.

4. CONCLUSIONS AND FUTURE WORK

We implemented a method for accelerating MC transport that unconditionally eliminated scattering events everywhere in phase space which we call Nose ("No scattering events"). We presented a thick diffusion limit study showing that Nose runtime is independent of ϵ whereas IMC runtime is $O(1/\epsilon^2)$, which explodes as $\epsilon \to 0$. We presented a verification study showing that Nose converges to an MMS solution with the hypothesized error of $\frac{C_0}{\sqrt{N}} + C_1 h$ in 1D, whereas the error in 2D exceeds the hypothesis. We are investigating the effect of noise on the 2D solution and testing ways to ameliorate it. We are also pursuing mathematical arguments for statements 1, 2, and 3 from the introduction and we are testing more sophisticated moment system discretizations. Future work includes demonstrating boundary sources and heterogeneity, and perhaps anisotropic scattering. Finally, we are interested in comparisons to RW as well as other acceleration schemes like IMD.

ACKNOWLEDGEMENTS

We thank Anil Prinja and especially Sam Olivier for productive discussions on Monte Carlo and moment methods. This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

REFERENCES

[1] 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).

[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 Comput Phys*, **volume 54**, pp. 508–523 (1984).

[3] 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).

[4] E. W. Larsen. "Diffusion-synthetic acceleration methods for discrete-ordinates problems." *Transport Theory and Statistical Physics*, **volume 13**(1-2), pp. 107–126 (1984).

[5] V. Y. Gol'din. "A quasi-diffusion method of solving the kinetic equation." *USSR Computational Mathematics and Mathematical Physics*, **volume 4**, pp. 136–149 (1964).

[6] E. Lewis and W. Miller. "A comparison of p1 synthetic acceleration techniques." *Transactions of the American Nuclear Society*, **volume 23** (1976).

[7] L. Chacón, G. Chen, D. A. Knoll, C. Newman, H. Park, W. Taitano, J. A. Willert, and G. Womeldorff. "Multiscale high-order/low-order (HOLO) algorithms and applications." *Journal of Computational Physics*, **volume 330**, pp. 21–45 (2017). URL https://doi-org.libproxy.berkeley.edu/10.1016/j.jcp.2016.10.069.

[8] H. Hammer, H. Park, and L. Chacón. "A multi-dimensional, moment-accelerated deterministic particle method for time-dependent, multi-frequency thermal radiative transfer problems." *Journal of Computational Physics*, **volume 386**, pp. 653–674 (2019). URL https://www.sciencedirect.com/science/article/pii/S0021999119301524.

[9] M. A. Cooper and E. W. Larsen. "Automated Weight Windows for Global Monte Carlo Particle Transport Calculations." *Nuclear Science and Engineering*, **volume 137**, pp. 1–13 (2001).

[10] S. S. Olivier. *High-Order Moment Methods for Thermal Radiative Transfer*. Ph.D. thesis, University of California, Berkeley (2022).

[11] 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). URL https://www.sciencedirect.com/science/article/pii/0021999187901707.