M&C 2017 - International Conference on Mathematics & Computational Methods Applied to Nuclear Science & Engineering, Jeju, Korea, April 16-20, 2017, on USB (2017)

Residual Monte Carlo Transport in Time with Consistent Low-Order Acceleration for 1D Thermal Radiative Transfer

Simon R. Bolding,* Jim E. Morel†

*Los Alamos National Laboratory, Los Alamos, NM
†Texas A&M University Nuclear Engineering Department, College Station, TX
sbolding@tamu.edu, morel@tamu.edu

Abstract - *We have extended a high-order low-order (HOLO) algorithm for thermal radiative transfer problems to include Monte Carlo (MC) integration of the time variable. Within each discrete time step, fixed-point iterations are performed between a high-order (HO) exponentially-convergent Monte Carlo (ECMC) solver and a low-order (LO) system of equations. The ECMC algorithm integrates the angular intensity over a time step, and the low-order (LO) radiation equations are closed consistently in the time variable. The time closure increases accuracy in optically-thin problems compared to a backward Euler discretization. The LO system is based on spatial and angular moments of the transport equation and a linear-discontinuous finite-element (LDFE) spatial representation, producing equations similar to the standard S_2 equations. The emission source is fully implicit in time, and Newton iterations efficiently resolve the nonlinear temperature dependence of the LO equations at each time step. The HO solver computes angular and temporal consistency terms that preserve the accuracy of the MC integration in the LO equations. We have implemented the ECMC algorithm with linear and constant, doubly-discontinuous trial spaces. Numerical results demonstrate that the second discontinuity in the time variable is necessary for sufficient consistency to achieve stable convergence, for the chosen closure of the LO equations. Results are compared to an implicit Monte Carlo (IMC) code and the HOLO algorithm with a BE time discretization. One-dimensional, gray test problems were tested for a range of optical thicknesses. The HOLO algorithm is more efficient and accurate than IMC with sufficient mesh resolution and number of particle histories.*

I. INTRODUCTION

Accurate solutions to the thermal radiative transfer (TRT) equations are important for simulations in the high-energy, high-density physics regime, e.g., for inertial confinement fusion and astrophysics. Computational modeling of TRT problems features coupling between a photon radiation field and a high-temperature material, where energy is exchanged through absorption and emission of photons by the material. Typical applications often require solution in a mix of streaming and diffusive regions due to absorption-emission physics and cross sections that are a function of material temperature. In this work, we improve on the time-integration accuracy of a high-order low-order (HOLO) method in optically thin regions where particles stream without undergoing many interactions, while preserving the computational efficiency of a residual MC HO solver in optically thick regions.

Moment-based hybrid Monte Carlo (MC) methods have demonstrated great potential for accelerated solutions to TRT problems [1, 2, 3]. These nonlinear acceleration methods iterate between a high-order (HO) transport equation and a low-order (LO) system formulated with angular moments and a fixed spatial discretization. Physics operators that are expensive for the HO solver to resolve directly in tightly coupled problems, e.g., photon absorption and emission, are moved to the LO system. The lower-rank LO equations can be solved with Newton methods to allow for nonlinearities in the LO equations to be efficiently resolved [4, 2]. The high-order (HO) problem is defined by the radiation transport equation with isotropic sources computed with the previous LO solution. A MC transport solution to the HO problem is used to construct consistency terms that appear in the LO equations. These consistency terms preserve the accuracy of the HO solution in the next LO solve, as the two solutions iteratively converge.

Previously, residual MC methods have been used to provide efficient solution to the HO transport problem [1, 2]; high-fidelity solutions, with minimal statistical noise, have been achieved for problems with optically-thick, diffusive regions that lead to slowly varying solutions. However, the algorithms in previous work have used a backward Euler (BE) discretization for the time variable. The BE discretization can inaccurately disperse radiation wavefronts in optically thin problems, leading to inaccuracies.

We have extended the algorithm in [2] to include higher-accuracy MC treatment of the time variable for the radiation unknowns. The exponentially-convergent Monte Carlo (ECMC) algorithm was modified to include integration of the time variable; this includes the introduction of a step, doubly-discontinuous (SDD) trial space representation in time. We have also investigated a linear, doubly-discontinuous (LDD) and a linear-discontinuous (LD) projection in time. The overall higher-dimensionality of the linear spaces requires development of a modified sampling approach that should be useful for extending the ECMC algorithm to higher spatial dimensions. A new parametric closure of the LO equations, introducing additional time-closure consistency terms, was derived to capture the time accuracy of the HO ECMC simulations. The LO equations can preserve the accuracy of the ECMC radiation transport treatment in time, with the same numerical expense as Backward Euler (BE) time-discretized S_2 equations. We have derived the LO equations directly from the transport equation such that, neglecting spatial discretization

differences, the HO and LO solutions are consistent upon convergence, preserving space-angle-time moments. Herein we briefly describe the algorithm, and we present results for one-dimensional (1D), grey test problems. We compare our method to the implicit MC (IMC) method [5] for accuracy and statistical efficiency for several representative problems.

1. Thermal Radiative Transfer Background and IMC

The continuous 1D, grey TRT equations consist of the radiation and material energy rate equations, i.e.,

$$\frac{1}{c}\frac{\partial I(x,\mu,t)}{\partial t} + \mu\frac{\partial I(x,\mu,t)}{\partial x} + \sigma_a I(x,\mu,t) = \frac{1}{2}\sigma_a acT^4(x,t) \quad (1)$$

$$\rho c_v \frac{\partial T(x,t)}{\partial t} = \sigma_a \phi(x,t) - \sigma_a acT^4(x,t), \quad (2)$$

with appropriate initial and boundary conditions specified. In the above equations, x is the position, t is the time, μ is the x-direction cosine of the angular intensity $I(x,\mu,t)$, σ_a is the macroscopic absorption cross section (cm^{-1}), and a, c, ρ, and c_v are the radiation constant, speed of light, mass density, and specific heat, respectively. Physical scattering could be included in Eq. (1), but it is omitted for brevity and simplicity. The desired transient unknowns are the material temperature $T(x,t)$ and the scalar radiation intensity $\phi(x,t) = \int_{-1}^{1} I(x,\mu,t)\,\mathrm{d}\mu$. The scalar intensity is related to the radiation energy density E_r by the relation $E_r = \phi/c$. The equations can be strongly coupled through the gray Planckian emission source $\sigma_a acT^4$, which is a nonlinear function of temperature, and the absorption term $\sigma_a\phi$. In optically thin problems, with small σ_a, the solution becomes increasingly linear as the emission source becomes negligible.

We will compare results in this work to the implicit Monte Carlo (IMC) method. The IMC method [5] is the standard approach for solution of the TRT equations with Monte Carlo particle transport [6]. The IMC method partially linearizes the system of equations over a discrete time step, with material properties evaluated at the previous-time-step temperature. The linearized system produces a transport equation with an approximate emission source and an effective scattering cross section representing absorption and re-emission of photons over a time step [5]. This transport equation is advanced over a time step via a MC simulation. The MC transport simulation tallies energy absorption over a discretized spatial mesh, which can be used to directly estimate a spatially discretized representation of the end of time step material temperature.

For this work, we are primarily interested in comparing to the time discretization of IMC. The material temperature and emission source are discretized with an implicit time discretization, i.e., a BE discretization. However, because the linearization is approximate, the system is not truly implicit, and there is a limit on the time step size to produce physically accurate results in problems that are tightly coupled and strongly nonlinear [7]. The linearized equations are integrated over the n-th time step defined for $t \in [t^{n-1/2}, t^{n+1/2}]$, with width $\Delta t = t^{n+1/2} - t^{n-1/2}$ and center $t^n = t^{n-1/2} + \Delta t/2$. The radiation equation is solved via MC simulation of particle histories, with the time-averaged energy deposition tallied over the spatial mesh. The time-integrated radiation equation, in nonlinear form, is

$$I^{n+1/2}(x,\mu) - I^{n-1/2}(x,\mu) = \Delta t\left[\sigma_a^{n-1/2}\bar{I}(x,\mu) - \mu\frac{\partial \bar{I}(x,\mu)}{\partial x} + \frac{1}{2}\sigma_a ac\left(T^{n+1/2}\right)^4(x)\right]. \quad (3)$$

The end-of-time-step intensity $I^{n+1/2}(x,\mu) \equiv I(x,\mu,t^{n+1/2})$ is stored as "census" particles that have reached $t^{n+1/2}$, representing a continuous sample of the phase space at that particular time [5], to be used in the next time step. In strongly diffusive regions, the accuracy will be limited to first order by the time discretization of the temperature terms. However, in optically-thin regions, higher-accuracy for the radiation terms is achieved. It is noted that the time-averaged effective scattering source resulting from linearization of the emission source in IMC is treated approximately in the time variable to allow the MC simulation to simulate the isotropic scattering events [6, 5].

2. The High-Order Low-Order Algorithm

Previously, we have developed a HOLO algorithm for 1D TRT problems, based on BE time-discretized HO and LO equations [2]. In the time-discrete HOLO algorithm, the LO solver resolves the time-discrete material temperature spatial distribution $T^{n+1/2}(x)$ over each time step, whereas the HO solver computes weighted angular integrals of the intensity. The HOLO formulation has several desirable properties. In particular, the LO solver can efficiently converge nonlinearities in diffusive systems, without the need to solve the nonlinear equations with MC simulation. Because the nonlinearities are converged, the temperature and emission source have a truly implicit discretization, preserving the discrete maximum principle [8]. Additionally, by using the ECMC HO solver, solutions with minimal statistical noise can be achieved efficiently, preventing instability issues that may be introduced through noise in the consistency terms.

To achieve temporal accuracy similar to IMC, we compute weighted temporal integrals of the intensity with the HO solver, used for computing additional consistency terms. We must assume a time discretization for the temperature field to produce a linear HO transport problem with closable LO equations. As in the IMC method, a BE time discretization is applied to emission source throughout, but the radiation variables are left in terms of time-averaged and end-of-time-step unknowns. Currently, our residual formulation requires a space-angle LDFE projection of the solution in order to estimate $I(x,\mu,t^{n+1})$, rather than the continuous sample represented by the census in IMC. This projection can be inaccurate with insufficient mesh resolution in near-void problems. However, the LDFE projection of the solution, estimated with MC inversion of the linear transport operator, will greatly increase the accuracy over a standard finite-difference discretization of the radiation equation. The HOLO algorithm should still demonstrate improvement over IMC in efficiency and accuracy in problems with intermediate and large optical thickness.

The fully-discrete LO equations are based on space-time-angle moments of the TRT equations, formed over a spatial

finite-element (FE) mesh. Angularly, the LO radiation equations are similar to S_2 equations, with element-averaged consistency parameters that are time-averaged, intensity-weighted averages of μ. The angular treatment is analogous to the hybrid method in [9]. A lumped LDFE spatial discretization (e.g., see [10]) is used to close the system spatially. Additional consistency parameters must be introduced to the LO equations to eliminate the auxiliary time-unknowns from the LO radiation equations. The additional time consistency terms are based on parametric modifications to a standard time discretization. Once closed, a system of equations is formed for the primary moment unknowns. If the angular and time consistency parameters were exact, then the LO equations would produce the exact moments of the solution, neglecting spatial discretization differences between the two systems. The HO consistency parameters are lagged in each LO solve. The LO equations always conserve energy, independent of the accuracy of the consistency terms.

The solution to the LO system is used to construct a spatially LDFE, and temporally constant, representation of the emission source on the right hand side of Eq. (1). This defines a fixed-source, pure absorber transport problem for the HO operator. This HO transport problem is solved with the ECMC algorithm. The HO transport problem can be viewed as a characteristic method, where we are using ECMC to invert the continuous-streaming, time-derivative, and removal operators [2]. The ECMC algorithm is an iterative residual MC method that uses batches of MC histories to estimate the error in the current trial-space estimate of $I(x,\mu,t)$. It is noted that because we are not using mesh adaptation in this work, exponential convergence in iterations cannot generally be maintained, but reduced variance overall can still be achieved. The initial guess for each solve is based on the solution from the previous time step, which allows for efficient reduction of statistical noise in problems with minimal change over the time step. The output from ECMC is a projection $\tilde{I}(x,\mu,t)$ of the intensity onto the chosen finite-element trial space, i.e., the functional representation of the intensity. Once computed, $\tilde{I}(x,\mu)$ is used to directly evaluate the necessary LO angular and time-closure consistency parameters. The HO solution is not used to directly estimate a new temperature at the end of the time step, which eliminates the need to linearize the emission source for stability.

Iterations between the HO and LO solves can increase accuracy in strongly nonlinear problems. However, for the problems tested here, only a single HO solve is performed during each time step. Thus, the HOLO algorithm, for the n-th time step, is

1. Perform a LO solve to produce an initial guess for $T_{LO}^{n+1/2}(x)$ and $\phi_{LO}^{n+1/2}(x)$, based on angular consistency terms estimated with $\tilde{I}^{n-1/2}(x,\mu)$ and a BE time discretization.

2. Solve the HO system for $\tilde{I}_{HO}(x,\mu,t)$ using ECMC, based on the current LO estimate of the emission source.

3. Compute LO angular and time-closure consistency parameters with $\tilde{I}_{HO}(x,\mu,t)$.

4. Solve the LO system using HO consistency parameters to produce a new estimate of $\phi_{LO}^{n+1/2}$ and $T_{LO}^{n+1/2}$.

5. Store $\tilde{I}^{n+1/2}(x,\mu) \rightarrow \tilde{I}^{n-1/2}(x,\mu)$, and move to the next time step.

II. THE LO SYSTEM

We will define the LO equations and closure before detailing the HO solver that is used to compute consistency terms present in the LO equations. To derive the LO equations, we reduce the dimensionality of Eq. (1) and Eq. (2) by taking spatial, angular, and temporal integrals. We will then introduce approximations to close the system, while being as consistent with the HO solver as possible.

The spatial domain is divided into N_c uniform spatial cells. The spatial moments are taken over each spatial cell i: $x \in [x_{i-1/2}, x_{i+1/2}]$, weighted with the standard linear FE basis functions. For example, the left moment operator is defined by

$$\langle\cdot\rangle_{L,i} = \frac{2}{h_i}\int_{x_{i-1/2}}^{x_{i+1/2}} b_{L,i}(x)(\cdot)\mathrm{d}x, \tag{4}$$

where $h_i = x_{i+1/2} - x_{i-1/2}$ is the width of the spatial element and $b_{L,i}(x) = (x_{i+1/2} - x)/h_i$ is the basis function corresponding to position $x_{i-1/2}$. The right moment is defined with basis function $b_{R,i}(x) = (x - x_{i-1/2})/h_i$. Angularly, the equations are integrated over the positive and negative half ranges. The angular integrals of the intensity are defined as $\phi^{\pm}(x) = \pm 2\pi\int_0^{\pm 1} I(x,\mu)\mathrm{d}\mu$. Finally, the equations are integrated over the n-th time step defined for $t \in [t^{n-1/2}, t^{n+1/2}]$ with width $\Delta t = t^{n+1/2} - t^{n-1/2}$ and center t^n.

The L and R moments and + and − half-range integrals are applied in pair-wise combination to Eq. (1), followed by integration over the time step. After algebraic manipulation, this ultimately produces 4 moment equations per spatial element. The streaming terms in the resulting equations are manipulated to form averages of μ, weighted with basis functions and the time-averaged intensity, analogous to previous work [2, 9]. The emission source and temperature-dependent cross sections are approximated with a BE discretization to help close the system. For example, application of the $\langle\cdot\rangle_{L,i}$ moment with the positive half-range integral to Eq. (1) ultimately yields

$$\frac{\langle\phi\rangle_{L,i}^{+,n+1/2} - \langle\phi\rangle_{L,i}^{+,n-1/2}}{c\Delta t} - 2\overline{\mu}_{i-1/2}^{+}\overline{\phi}_{i-1/2}^{+} + \overline{\{\mu\}}_{L,i}^{+}\langle\overline{\phi}\rangle_{L,i}^{+} + \overline{\{\mu\}}_{R,i}^{+}\langle\overline{\phi}\rangle_{R,i}^{+} + \sigma_{a,i}^{n+1/2}h_i\langle\overline{\phi}\rangle_{L,i}^{n+1/2,+} = \frac{h_i}{2}\langle\sigma_a^{n+1/2}acT^{n+1/2,4}\rangle_{L,i}, \tag{5}$$

where over-barred quantities represent the exact averaging over the time step. A more thorough derivation and definitions for all of the moment equations can be found in [11].

At this point, the only approximation has been the BE time-discretization of the emission source in both governing equations. The face- and volume-averaged angular consistency terms, e.g., $\overline{\mu}_{i-1.2}^{+}$, were formed only through algebraic manipulation. They are approximated with angular intensity from the

previous HO solve. For example, the L and + time-averaged consistency term is

$$\overline{\{\mu\}}^{+}_{L,i,HO} \simeq \frac{\frac{2}{h_i\Delta t}\int_{t^{n-1/2}}^{t^{n+1/2}}\int_{0}^{1}\int_{x_{i-1/2}}^{x_{i+1/2}} \mu\, b_{L,i}(x)\tilde{I}_{HO}(x,\mu,t)\mathrm{d}x\mathrm{d}\mu\mathrm{d}t}{\frac{2}{h_i\Delta t}\int_{t^{n-1/2}}^{t^{n+1/2}}\int_{0}^{1}\int_{x_{i-1/2}}^{x_{i+1/2}} b_{L,i}(x)\tilde{I}_{HO}(x,\mu,t)\mathrm{d}x\mathrm{d}\mu\mathrm{d}t}, \quad (6)$$

where $\tilde{I}_{HO}(x,\mu,t)$ is a space-angle-time finite element projection of the HO intensity, to be later defined. It is noted that this consistency term contains no division by σ_a, so these equations are directly valid in a void.

For simplicity, the face terms $\phi^{\pm}_{i\pm1/2}$ are eliminated from the system using a lumped LDFE spatial approximation, with standard upwinding [2]. The emission source is also represented with a lumped LDFE interpolant. There is some inconsistency introduced in the lumped LDFE spatial approximations. Assuming iterative convergence of consistency terms, the LO solution and projection of the HO solution may differ for any given spatial mesh, but the two solutions will converge as the mesh is refined. This approximation has proven stable for problems tested and demonstrates preservation of the equilibrium diffusion limit [11]. Boundary conditions are incorporated through upwinding and the face term resulting from integration of the streaming operator.

The material energy equations are similarly integrated in space and time. The lumped LDFE approximation is introduce for $T(x)$ and $T^4(x)$ to close the equation spatially, along with the BE time discretization for the emission source. The L moment temperature equation is

$$\frac{\rho_i c_{v,i}}{\Delta t}\left[T^{n+1/2}_{L,i} - T^{n-1/2}_{L,i}\right] + \sigma^{n+1/2}_{a,i}\left(\langle\overline{\phi}\rangle^{+}_{L,i} + \langle\overline{\phi}\rangle^{-}_{L,i}\right) = \sigma^{n+1/2}_{a,i} ac\left(T^{n+1/2}_{L,i}\right)^4, \quad (7)$$

where cross sections have been evaluated at $t^{n+1/2}$ and $T_{L,i}$ and $T_{R,i}$ are the LD edge values of the temperature, e.g., $T(x) = b_{L,i}(x)T_{L,i} + b_{R,i}(x)T_{R,i}$ for $x \in (x_{i-1/2}, x_{i+1/2})$.

1. Parametric Time Closure with HO information

At this point, there is still too many unknowns in the LO equations. Quantities at $t^{n-1/2}$ are known from the previous time step or an initial condition, but a a relation is needed between the time-averaged radiation quantities and their corresponding values at $t^{n+1/2}$. The closure of each LO equation must account for inconsistencies in the time-discretization of the two solvers. Previous work, applied to radiation-hydrodynamics problems, has enforced consistency in time by adding a local artificial source to the time-discretized LO equations in each cell [3]. This source was based HO estimate of the difference in the integral treatments of the time derivative between the HO and LO systems. The advantage of this form is that the LO solver exclusively deals in time-averaged unknowns for the radiation terms in the equations. Alternatively, we will use a local, parametric closure to directly eliminate the auxiliary temporal radiation unknowns, introducing additional consistency terms.

Equation (5) will only contain time-averaged radiation unknowns if $\langle\phi\rangle^{n+1/2}_{L,i}$ is eliminated from the system. The simplest closure is a weighted average

$$\langle\phi\rangle^{+,n+1/2}_{L,i} \approx \gamma^{+}_{L,i,HO}\langle\overline{\phi}\rangle^{+}_{L,i}, \quad (8)$$

where $\gamma^{+}_{L,i,HO}$ is a time-closure consistency parameter. The consistency parameter can be determined from Eq. (8) by using moments of $\bar{I}_{HO}(x,\mu)$ and $I^{n+1/2}_{HO}(x,\mu)$, i.e.,

$$\gamma^{+}_{L,i,HO} = \frac{\langle\phi_{HO}\rangle^{+,n+1/2}_{L,i}}{\overline{\phi}^{+}_{L,i,HO}}. \quad (9)$$

Because the time-closures account for the different spatial moment equations, there is four per spatial cell. For a linear problem, as long as the HO solution used to compute this closure satisfies the same moment equations as the LO system, or is at least *an accurate approximation to the moment equations*, then the closure relation will stably provide consistency.

The unknowns of interest are $\langle\overline{\phi}\rangle^{\pm}_{L,i}$, $\langle\overline{\phi}\rangle^{+}_{R,i}$, $T^{n+1/2}_{L,i}$, and $T^{n+1/2}_{R,i}$. The four spatially-closed radiation moment equations per cell, the HO approximation of the angular consistency terms, the two temperature moment equations, and the parametric time closures (e.g., Eqs. (5), (6), (7), and (8)) provide sufficient equations to solve for these unknowns. Summation of the moment equations over all cells and application of boundary conditions defines a global, nonlinear LO system of equations. This discrete system of equations is solved using a hybrid Newton-Picard method, as in previous work [2]. The linearized equations produce scattering terms that couple the two directions together, which can be directly inverted in 1D. The LO system is fully converged within each solve. Once time-averaged unknowns have been calculated, the local time closures provide $\phi^{n+1/2}_{LO}(x)$ for the next time step.

For the initial LO solve, within a time step, the angular parameters are calculated based on the $\tilde{I}^{n-1/2}_{HO}(x,\mu)$ and all γ values are set to unity, producing a BE discretization. Other closures, e.g., a modified Crank-Nicolson, have been explored. In optically thin problems, the problem is nearly linear, and the choice of this closure has minimal effect on results because all other auxiliary unknowns have been consistently eliminated from the system with HO information (with the exception of the spatial closure). However, for optically thick problems with high statistical noise, the Crank-Nicolson introduced some instabilities.

III. THE RESIDUAL MC HIGH ORDER SOLVER

1. Trial Space Representation

To apply the ECMC algorithm [12, 2], it is necessary to have a functional representation of the intensity for all phase space variables so a residual can be evaluated. A finite element representation is formed in x, μ, and t. The domain is divided into a uniform grid, where the element with the i-th spatial, j-th angular, and n-th temporal indices spans the domain $\mathcal{D}_{ijn} : x_{i-1/2} < x < x_{i+1/2} \times \mu_{j-1/2} \leq \mu \leq \mu_{j+1/2} \times t^{n-1/2} <$

$t \leq t^{n+1/2}$. We have implemented three different trial space representations for the intensity in t. However, in x and μ, the intensity is always represented with an LDFE projection, which we will denote $\tilde{I}(x,\mu)$. The LDFE projection preserves the zeroth and first moments in x and μ of the intensity. Standard upwinding [10, 2] is used to define the solution on faces for evaluating terms resulting from the spatial derivative in the streaming term. The LDFE projection is not guaranteed to be positive. Before computing consistency terms, any detected negative values for $\tilde{I}^{n+1/2}(x,\mu)$ or $\tilde{I}^{n}(x,\mu)$ are made positive by uniformly decreasing the slopes or setting the average to the floor temperature in some cases (see [11] for more details). In time, values at $t^{n-1/2}$ are upwinded from the previous time step for both trial spaces.

The first time space is a step, doubly-discontinuous (SDD) trial space, with the time variable illustrated in Fig. 1b. The SDD trial space representation for $I(x,\mu,t)$ is

$$\tilde{I}(x,\mu,t) = \begin{cases} \tilde{I}^{n-1/2}(x,\mu) & t = t^{n-1/2} \\ \tilde{I}^{n}(x,\mu) & t \in (t^{n-1/2}, t^{n+1/2}) \\ \tilde{I}^{n+1/2}(x,\mu) & t = t^{n+1/2} \end{cases} \tag{10}$$

where we have used $\tilde{I}^{n}$ to denote the time-averaged LDFE projection in x and μ of the intensity over the interior of the time step; the LDFE projections at $t^{n-1/2}$ and $t^{n+1/2}$ are denoted $\tilde{I}^{n-1/2}$ and $\tilde{I}^{n+1/2}$, respectively. The SDD trial space provides a projection for all the desired unknowns that result from time integration of the transport equation; it provides sufficient information to close the LO equations and evaluate the temporal consistency terms. Another benefit of this trial space is it allows for the residual sampling infrastructure from the time-discrete formulation of this algorithm to be used with minor modifications.

The second trial space is the LDD trial space in time, as illustrated in Fig. 1b, with an LDFE representation in x and μ. For a particular space-angle-time element, this trial space is defined as

$$\tilde{I}(x,\mu,t) = \begin{cases} \tilde{I}_{ij}^{n-1/2}(x,\mu) & t = t^{n-1/2}, \\ \tilde{I}_{ij}^{n}(x,\mu) + \frac{2}{\Delta t} I_{t,ij}^{n}(t-t^{n}), & t \in (t^{n-1/2}, t^{n+1/2}), \\ \tilde{I}_{ij}^{n+1/2}(x,\mu) & t = t^{n+1/2} \end{cases} \tag{11}$$

where $\tilde{I}_{ij}^{n}(x,\mu)$ is the time-averaged LDFE projection in x and μ over $\mathcal{D}_{ijn}$ and $I_{t,ij}^{n}$ is the finite-element slope of $I(x,\mu,t)$ averaged over $\mathcal{D}_{ijn}$, i.e.,

$$I_{ij}^{n} = \frac{6}{\Delta t} \iiint_{\mathcal{D}_{ijn}} \left(\frac{t-t^{n}}{\Delta t} \right) I(x,\mu,t)\, \mathrm{d}x\mathrm{d}\mu\mathrm{d}t. \tag{12}$$

Thus, there is a unique time slope for each element.

We will also test problems for an LD trial space in time. The only difference from the LDD trial space is that there is no discontinuity at $t^{n+1/2}$. To compute consistency terms and advance to the next time step, the LD approximation of $\tilde{I}(x,\mu,t)$ is simply evaluated at $t^{n+1/2}$, extrapolating to the end of the time step. This introduces an additional approximation error for $I^{n+1/2}(x,\mu)$, and leads to a HO solution that does not as closely satisfy the exact moment equations (e.g., Eq. (5)), as accurately as the LDD trial space. This inconsistency can lead to instabilities in the LO equations for rapidly varying solutions.

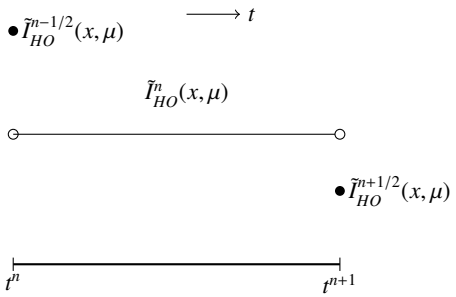

(a) SDD trial space

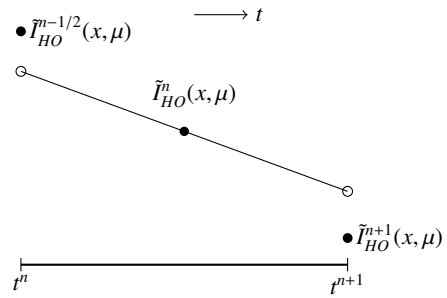

(b) LDD trial space

Fig. 1: Illustration of the time variable for $\tilde{I}_{HO}(x,\mu,t)$ with two unique trial spaces.

2. The Algorithm

The transport equation to be solved by ECMC is given by Eq. (1), but with a fixed LDFE Planckian emission source that is estimated by the previous LO solve. We write the equation in operator notation as

$$\mathbf{L}I(x,\mu,t) = q_{LO}(x) \tag{13}$$

where $q_{LO} = \sigma_a ac \left(T^{n+1/2}\right)_{LO}^{4}/2$ denotes the latest estimate of the emission source, and remains constant for the entire HO solve. The *continuous* linear transport operator $\mathbf{L}$ is

$$\mathbf{L}I(x,\mu,t) \equiv \left[\frac{1}{c}\frac{\partial}{\partial t} + \mu\frac{\partial}{\partial x} + \sigma_a \right] I(x,\mu,t). \tag{14}$$

The m-th approximate solution to Eq. (13) is $\tilde{I}^{(m)}(x,\mu,t)$, where m identifies the MC batch. The m-th residual is

$$r^{(m)} = q - \mathbf{L}\tilde{I}^{(m)}. \tag{15}$$

Addition of the residual equation to Eq. (13) gives the error equation

$$\mathbf{L}(I - \tilde{I}^{(m)}) = \mathbf{L}\epsilon^{(m)} = r^{(m)}, \tag{16}$$

where $I(x,\mu,t)$ is the exact solution to Eq. (13) (which contains approximation error from the representation of $\tilde{I}^{n-1/2}(x,\mu)$ and q_{LO}), and $\epsilon^{(m)}$ is the error in $\tilde{I}^{(m)}$.

The inverse of $\mathbf{L}$ in Eq. (16) is estimated via MC simulation without discretization error. This is a standard MC simulation, where particle histories are tracked in space, angle, and time, e.g., in IMC [5, 6, 11]. Particle histories are sampled from the source $r^{(m)}(x,\mu,t)$, as explained below. Tallies of the error particles estimate moments of $\epsilon^{(m)}$, which are added to the moments used to construct the finite-element representation $\tilde{I}^{(m)}$. In operator notation, we denote this as $\tilde{\epsilon}^{(m)} = \mathbf{L}^{-1} r^{(m)}$. The LDFE projections of the error $\overline{\epsilon}$ and $\epsilon^{n+1/2}$ are computed using generalizations of volumetric path-length and particle density estimators. The estimators are weighted by appropriate basis functions over each element. For the algorithm with the SDD trial space, particles are allowed to stream without interaction, and the tallies are adjusted accordingly [2]. The details of the tallies specific to this work are given in Sec. V.

The ECMC algorithm is

1. Initialize $\tilde{I}^{(0)}(x,\mu,t)$ with $\tilde{I}^{n-1/2}(x,\mu)$.
2. Compute $r^{(m)}$.
3. Estimate $\tilde{\epsilon}^{(m)} = \mathbf{L}^{-1} r^{(m)}$ with N Monte Carlo histories.
4. Compute $\tilde{I}^{(m+1)} = \tilde{I}^{(m)} + \tilde{\epsilon}^{(m)}$
5. Optionally repeat 2 – 4 for desired number of batches.

The use of $\tilde{I}^{n-1/2}(x,\mu)$ as the initial guess greatly increases statistical efficiency in regions of the problem where the solution is slowly varying. For the LDD and LD trial spaces, we also initialize the time slopes with the corresponding values from the previous time step. If the error is sufficiently estimated each batch, both statistically and with the projected trial-space representation, then the overall error in the solution can converge at an exponential rate. However, eventually the estimated projection of the error is not sufficiently accurate and adaptive refinement would be necessary to continue convergence. It is not clear what the best approach to adapt the solution in time is, or if it is practical. Thus, we are primarily only gaining the residual benefit for the algorithm in this work, although in some cases multiple batches can improve overall efficiency over a single batch.

A drawback of this HO algorithm is that a truncation error occurs by keeping only the LDFE projection of the intensity between time steps, which is not present in IMC. Adaptive mesh refinement is likely necessary to efficiently capture rapidly-varying solutions, but this was not done here for simplicity. Adaptive refinement in space and angle could be included in the iterative algorithm in future work, which has been demonstrated for the time-discrete algorithm previously in [2]. This would also help minimize memory requirements.

3. Sampling from the Residual

A. The SDD Trial Space

Computing and sampling from the residual defined by Eq. (15) is similar to the sampling algorithms for a steady-state transport equation [12, 13, 2]. The discontinuities in Eq. (10) introduce δ-function sources at $t^{n-1/2}$ and $t^{n+1/2}$ because of the time derivative. Additionally, the residual has a spatial δ-function source on the upwind face of each element (resulting from the spatial derivative in the streaming term), and a 2D linear, interior volumetric source. The contribution from the δ-function source at $t^{n+1/2}$ can be analytically determined because all particles born immediately reach census [11]. Thus, it is never sampled, and the contribution is added in at the end of the simulation.

Because the residual can be negative, particles can be sampled with both negative and positive weights. The particles are sampled from $|r(x,\mu,t)|$ using rejection sampling over each element. The weights are modified to be negative if $r(x,\mu,t) < 0$ for the sampled phase-space position. Starting particle weights are normalized to have a magnitude of unity. The final tallies are then multiplied by $\|r(x,\mu,t)\|_1$, the L_1 norm of the residual over the entire sampling domain. Because of the choice of the SDD trial space, the most complex L_1 integral is the two-dimensional integral of a linear function. Thus, the L_1 norm over all sampling space can be analytically evaluated, as in previous work [12]. To reduce variance in optically thick regions, systematic sampling [14] is performed, with particles placed proportional to the magnitude of the residual over each element, as in [2]. Then the choice of a volumetric or either δ-function source within the element is discretely sampled, and the corresponding probability distribution function (PDF) sampled with rejection.

B. The LD and LDD Trial Spaces

Sec. V. provides a definition of the residual for the LDD trial space. The only difference between the residual for the LDD and LD trial spaces is the discontinuity source at $t^{n+1/2}$, which is analytically treated as for the SDD trial space. Unlike the SDD trial space, we can not evaluate the L_1 norm of the residual analytically. Additionally, the higher-dimensional residual terms will generally be less efficient to sample with rejection, at least for certain elements. Alternatively, we can use importance sampling [14] with unnormalized particle weights to estimate the magnitude of the residual. Previous work on higher-dimensional residual MC has applied a similar approach for a continuous global polynomial expansion trial space [15]. Because the solution is continuous, except for at the boundary, a uniform sampling can be performed over the entire domain and boundary, with weights that correct for the bias and estimate the magnitude of the residual. Because our finite element space contains spatial and temporal discontinuities for each element, particles should be distributed more closely to the true residual. Additionally, because $\tilde{I}^{n-1}(x,\mu,t)$ is typically a good approximation to $I(x,\mu,t^n)$, uniform sampling of the domain is very inefficient for thermal radiative transfer problems.

To apply the importance sampling algorithm, we sample from a simpler distribution function that can correctly be normalized to a PDF. The chosen PDF represents a decent approximation to the residual. For each element, the new PDF to be sampled from is a piece-wise constant function, spanning the same domain as the true residual, i.e., including two δ functions, with their corresponding subsection of the domain, and the full domain $\mathcal{D}_{ijn}$ (for the interior source). The probability of sampling a particular constant source is proportional to an approximation of the L_1 norm of the residual over that element.

The L_1 norm is approximated with product 2-point Gaussian quadratures over each piece of the residual domain. Thus, the PDF for an element becomes

$$\begin{aligned} p^n_{ij}(x,\mu,t) &= \frac{\|r\|^n_{1,i\pm1/2\,j}}{\|r\|_1 \Delta t \Delta\mu}\delta^{\mp}\left(x-(x_i \pm \frac{h_i}{2})\right) \\ &+ \frac{\|r\|^{n-1/2}_{1,ij}}{\|r\|_1 \Delta x \Delta\mu}\delta^{+}\left(t-t^{n-1/2}\right) \\ &+ \frac{\|r\|^n_{1,ij}}{\|r\|_1 \Delta t \Delta\mu \Delta x}, \quad (x,\mu,t)\in\mathcal{D}_{ijn}, \end{aligned} \tag{17}$$

where $\|r\|_1$ is the L_1 norm of $r(x,\mu,t)$ over the entire domain, $\|r\|^n_{1,i\pm1/2\,j}$ is the norm of the spatial δ-function portion of the element residual (where the $\pm$ sign corresponds to the direction of μ for the element), $\|r\|^{n-1/2}_{1,ij}$ is the norm of the temporal δ-function portion of the residual, and$\|r\|^n_{1,ij}$ is the norm of the residual over the interior of the element domain; all of these norms are approximated with quadrature. Particles are trivially sampled from $p(x,\mu,t)$ and particle weights are initialized as

$$w(x,\mu,t) = \frac{r(x,\mu,t)}{p(x,\mu,t)}. \tag{18}$$

Although the quadrature approximation may be poor in regions of the domain where zero-crossings of the residual occur, the overall sampling algorithm is unbiased. We expect that for reasonably fine meshes the locations of particle origins are sufficiently proportional to the magnitude of the residual. High relative variance in weights of particles can also lead to large variances in the tallies [14]. However, the slope tallies and extrapolated solution should help produce less noise than the census tallies of the SDD trial space overall. As before, we stratify based on $p(x,\mu,t)$ to place particles proportional to the total probability of sampling from each element and adjust the weights to account for sampling of integer number of histories.

IV. RESULTS AND ANALYSIS

We have simulated three 1D, grey test problems to demonstrate the efficacy of our HOLO algorithm: a near-void problem, an optically thin problem, and a standard Marshak wave problem. We will compare sample statistics and accuracy of solutions to an IMC code with source tilting [16] and the HOLO algorithm with BE time-discretized equations [2]. Throughout this section, results that use the BE time discretization are indicated with HOLO-BE, and results with the MC-based time closure are indicated with HOLO-LD, HOLO-LDD, or HOLO-SDD, corresponding to the chosen temporal trial space. All plotted results depict the LO solution from the HOLO algorithm. Some figures depict an effective radiation temperature defined as $T_r(x) = (\phi(x)/ac)^{1/4}$

To provide a quantitative measure of statistical efficiency, a spatially integrated measure of variance in cell-averaged radiation energy densities $E^{n+1/2}_{r,i}$ was computed, based on the solution at the simulation end time. To form sample estimates of the variance in the cell-averaged values, twenty independent simulations for each particular result and method were performed.The sample variance for a particular cell-averaged $E^{n+1/2}_{r,i}$ is

$$S^2_i = \frac{20}{20-1}\sum_{l=1}^{20}\left(\overline{E^{n+1/2}_i} - E^{n+1/2,(l)}_{r,i}\right)^2, \tag{19}$$

where $E^{n+1/2,(l)}_i$ is the cell-averaged scalar intensity for cell i from the l-th of 20 independent simulations, and $\overline{E^{n+1/2}_i}$ is the corresponding sample mean, from the 20 simulations. To provide a normalized, spatially-integrated result, we form a norm over cells as

$$\|s\| = \left(\frac{\sum_{i=1}^{N_c} S^2_i}{\sum_{i=1}^{N_c}\left(\overline{E^{n+1/2}_{r,i}}\right)^2}\right)^{1/2}, \tag{20}$$

where N_c is the number of spatial elements. We use this statistic to form a figure of merit (FOM) to demonstrate how statistical accuracy scales with the number of histories performed. Our FOM is defined as

$$\text{FOM} = \frac{1}{N_{\text{tot}}\|s\|^2} \tag{21}$$

where N_{tot} is the total number of histories performed during a particular simulation. A larger value of the FOM indicates that the method produced less variance in the solution, per history performed, for a given problem. This form of the FOM is typically chosen because the variance is expected to reduce inversely proportional to N_{tot}, so for standard MC simulations the FOM becomes, on average, independent of N_{tot} [14]. The FOM is not necessarily expected to be independent of N_{tot} for the HOLO method because of correlations in the ECMC solver. It is difficult to compare computational times with the IMC method because they are implemented in different code infrastructures. Additionally, the HOLO algorithm is converging nonlinearities, whereas the IMC method is only taking a single linearized step, for each time step. Thus, we only focus on the reduction in overall variance, per particle history, as a measure of efficiency.

We have also compared against IMC for accuracy for the optically thin problem. An integral error in cell-averaged energy densities for the l-th simulation, with $N^{(l)}_c$ spatial cells, is computed as

$$\|e\|^{(l)}_{a,rel} = \left(\frac{\sum_{i=1}^{N^{(l)}_c}\left(E^{n+1/2,(l)}_{r,i} - E^{n+1/2,ref}_{r,i}\right)^2}{\sum_{i=1}^{N^{(l)}_c}\left(E^{n+1/2,ref}_{r,i}\right)^2}\right)^{1/2}, \tag{22}$$

where $E^{n+1/2,ref}_{r,i}$ is computed by spatially averaging the reference IMC solution over the i-th coarse spatial cell, with the assumption of uniform mesh spacing.

1. Near-Void Problem

For the first problem, the material properties are uniform throughout a 2.0-cm wide domain with $\rho c_v = 0.01374$ GJ cm^{-3} keV^{-1}, $\sigma_a = 10^{-6}$ cm^{-1}. The near-void cross section essentially uncouples the material and radiation. The material and radiation are initially in equilibrium at a temperature of

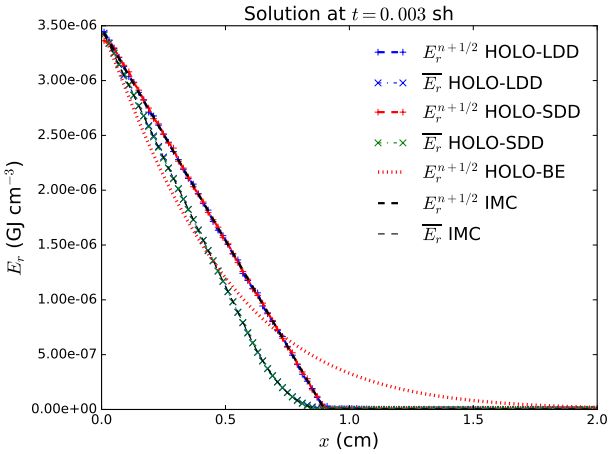

(a) Radiation energy densities $E_{r,i}$ with $\Delta t = 0.001$ sh.

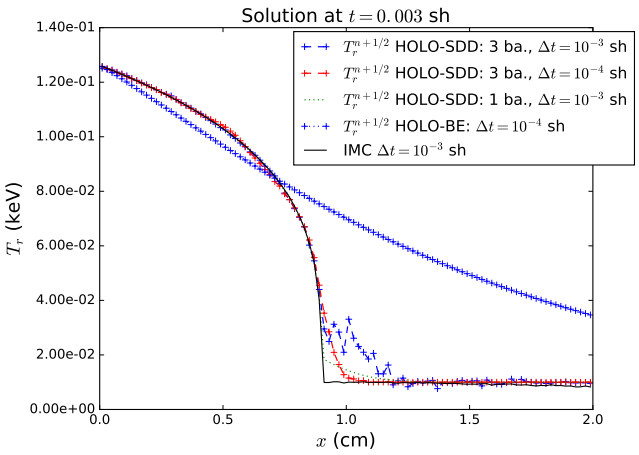

(b) Radiation temperatures $T_{r,i}$ for different time step sizes and numbers of batches.

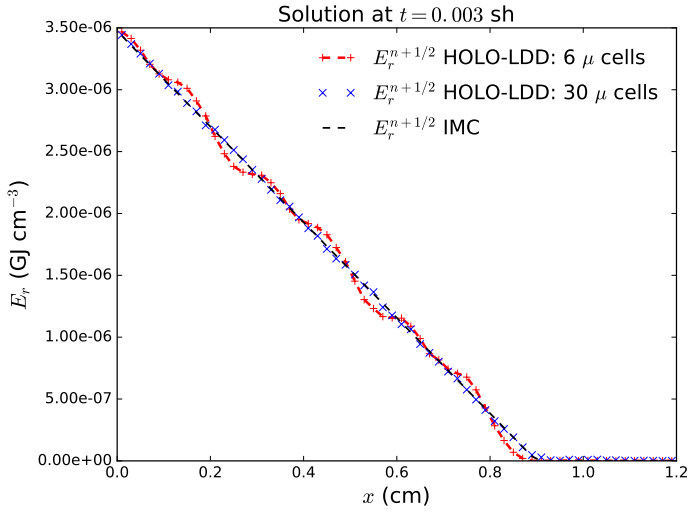

(c) Census radiation energy densities $E_{r,i}^{n+1/2}$ for the HOLO-LDD method with different numbers of μ cells and $\Delta t = 0.001$ sh.

Fig. 2: Comparisons between IMC and the HOLO method with different time discretizations and settings.

0.01 keV. An isotropic incident intensity with $T_r = 0.150$ keV is applied at $x = 0$ for $t > 0$; the incident intensity on the right boundary is 0.01 keV. The simulation end time is $t = 0.003$ sh.

A comparison of the cell-averaged radiation energy densities E_r for IMC and the HOLO method with different time closures is depicted in Fig. 2a; values for the time-averaged and $t^{n+1/2}$ energy densities, from the final time step, are given. The end of time step value for the HOLO method with a BE discretization is also depicted. Three relatively large time steps are taken with $\Delta t = 0.001$ sh. For the HOLO results, three ECMC batches were performed with a total of 3×10^6 histories per time step, and the IMC results were generated with 3×10^6 histories per time step. The spatial meshes had 100 spatial cells and all HOLO results used 30 μ cells. The MC treatment of the time variable and the closure of the LO equations allow the LO results to correctly reconstruct the wave-front location of IMC, whereas the BE discretization artificially propagates energy. This problem is nearly linear due to the small cross sections, so the HO time closures are stable and correctly reproduce the time-averaged and $t^{n+1/2}$ HO moments, even for the large time steps of this problem. This problem represents a limiting case where IMC is more efficient and accurate. Because the IMC particles are only streaming, the space of the initial source is all that must be sampled, whereas the ECMC algorithm must resample the distribution and residual space each time step.

A comparison of HOLO-SDD solutions with variable numbers of batches and time step sizes, but the same total 9×10^6 histories for each simulation, is given in Fig. 2b. Results are depicted as radiation temperatures. By plotting proportional to the fourth-root of the radiation energy density, the noise at low magnitudes past the wave-front are more apparent

for the 3 batches and $\Delta t = 0.001$ sh case. This noise is minimal relative to the scale of E_r, but it demonstrates a deficiency of the residual method. The representation over the time step leads to particles sampled near the wave-front with a time near t^n that travel into the equilibrium region. This is not a bias, but rather an under-sampling of the phase space; if sufficient histories were performed there would be negative particles that canceled out this error. The ECMC iterations can also lead to small negative averages for the HO solution in the equilibrium region. In such cells, the average was set to the floor value and slopes to zero. For the case of a single batch, there is less noise past the wavefront because the choice of $\tilde{I}^{n-1/2}(x,\mu)$ as an initial guess for $\tilde{I}^n(x,\mu,t)$ prevents most particles from traveling past what the physical transport should allow.

The noise in the equilibrium region is significantly reduced by taking a smaller time step, although the mesh-imprinting error is increased. The mesh-imprinting error is a discretization effect resulting from only storing a projection between time steps. Fig. 2c illustrates the convergence of this effect as a function of the number of angular elements used, for the HOLO-LDD method. At coarser mesh sizes, the imprinting of the mesh is visible in the location of the wavefront, although the location of the wavefront is accurate. Smaller time step sizes can increase the mesh imprinting because a more difficult to resolve solution is being projected more often. The error is reduced as the mesh is refined. Generally, the imprinting is reduced as σ_a is increased and absorption-emission events smooth the angular intensity across each time step.

2. Optically-Thin Problem

We modify the near-void problem by increasing the absorption cross section to 0.2 cm^{-1}; all other problem parameters are the same. Radiation temperatures at the end of the last time step are compared in Fig. 3 for the IMC and various HOLO methods. The HOLO results were generated with 30 μ cells, and all spatial meshes used 100 cells. All results used 3×10^6 histories per time step. There is good agreement between the HOLO-LDD and HOLO-SDD results with IMC, except some dispersion past the wavefront. This dispersion is caused by the spatial discretization inconsistency between the LDFE HO projection and the lumped LDFE LO equations; this dispersion is not present in the HO solution. As in the previous problem, the HOLO-BE results are very inaccurate at capturing the wavefront location. IMC demonstrates substantial statistical noise in the equilibrium region. The HOLO-LD method is unstable for this problem. This is caused by inconsistency between the HO and LO space-angle-time moment equations. In particular, the approximation of $I^{n+1/2}(x,\mu)$ via extrapolation based on the interior slope leads to an unstable closure relation. At lower history counts, the other HOLO methods can become unstable as well.

The value of $\|e\|_{a,rel}$ was computed by averaging the computed value of Eq. (22) from 20 simulations. The IMC method and various HOLO methods are compared against a reference IMC solution. Because the material is loosely coupled in this problem, we expect IMC to be accurate with sufficient particle histories. The reference solution is the average of 20 IMC simulations of 20×10^6 histories, each with $\Delta t = 10^{-4}$ sh. The estimated value of $\|s\|$ for the reference solution is 0.025%. Sample statistics are also compared via the FOM from Eq. (21).

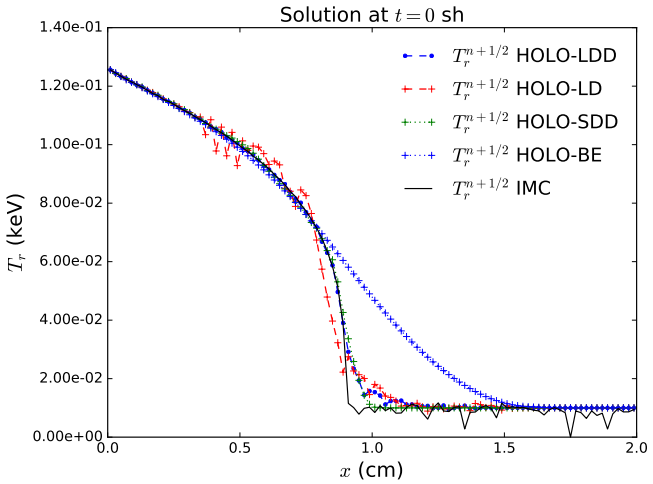

Fig. 3: Comparison of radiation temperatures of IMC and the HOLO method for different time step sizes and numbers of batches, for optically thin problem.

Table. I compares values of the FOM and $\|e\|_{a,rel}$ for $\Delta t = 10^{-4}$ sh. The results in the table are for 200 spatial cells, computed against a reference solution with 100 spatial cells. The HOLO results were generated with a single batch per time step, although similar accuracy was found for the case of two batches (with the same number of histories per time step). All FOM values are normalized to the IMC result with 200 x cells and 30,000 histories.

The HOLO method, with sufficient histories to prevent instabilities, is more accurate and substantially more efficient. For reference, statistics were measured for the HOLO-BE method with two batches of 150,000 histories per time step, producing $\|e\|_{a,rel} = 10.5\%$ and FOM = 3100, demonstrating substantial inaccuracy but improved efficiency. For 100 x cells, the projection error limits accuracy and the IMC method becomes more accurate.

TABLE I: Comparison of $\|e\|_{a,rel}$ and FOM values for the end of time step radiation energy densities, of the last time step, for the optically thin problem, 200 x cells, and $\Delta t = 1 \times 10^{-4}$ sh. The reference IMC result has 100 x cells. Simulation end time is $t = 0.003$ sh. Fractional error in all results below 0.01

$\|\mathbf{e}\|_{\mathbf{a,rel}}$			
hists./step	IMC	HOLO-SDD	HOLO-LDD
30,000	2.93%	14.00%	14.50%
300,000	0.99%	0.37%	0.46%
1,000,000	0.49%	0.18%	0.19%
FOM			
hists./step	IMC	HOLO-SDD	HOLO-LDD
30,000	1	0.11	0.10
300,000	0.90	14.24	8.61
1,000,000	1.01	81.71	71.36

3. Marshak Wave Problem

It is important to demonstrate that the time closures are stable in a mix of optically-thick and optically-thin regions and that the ECMC method is still efficient for such problems. This problem has the same material properties as the optically thin problem except for a temperature-dependent cross-section with $\sigma_a(T) = 0.001T^{-3}$ and the initial temperature is 2.5×10^{-5} keV. The time step size is linearly increased from 0.001 sh to a maximum step of 0.01 sh over the first 10 time steps. It was found for this problem that it was necessary to use more than one batch for the HOLO-SDD and HOLO-LDD methods and large numbers of histories to stably converge, as it is difficult for particles to reach the end of the time step. A second batch results in a better initial approximation for $I^{n+1/2}$, resulting from the extrapolated interior solution of the first batch. Thus, all HOLO results in this section use two batches.

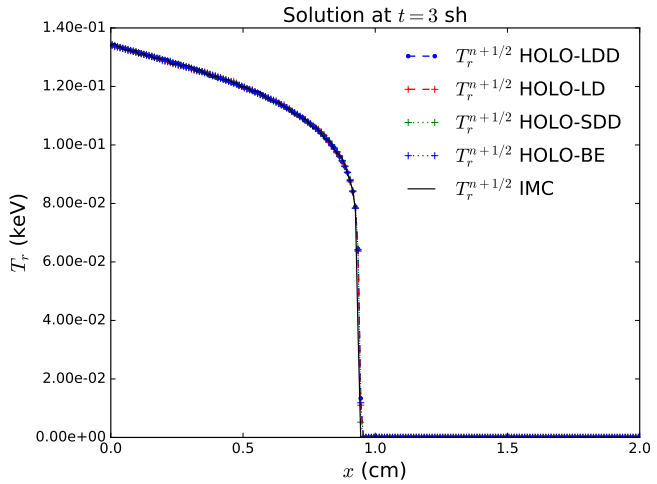

Fig. 4: Comparison of HOLO-TC, HOLO-BE, and IMC methods for the Marshak Wave problem, with 10^6 histories per time step.

Figure 4 compares the accuracy of IMC and the various HOLO methods, showing agreement between the solutions. There were 10^6 histories per time step for all simulations. This slowly varying problem can be accurately modeled with the BE time discretization, but the MC time closures are stable. Because the problem is slowly-varying, the HOLO-LD method was found to stably converge, without the need for the second discontinuity. The HOLO-LD method is sufficiently consistent because for this problem $\tilde{I}^n(x,\mu) \simeq I^{n+1/2}(x,\mu)$.

Table II compares sample statistics for IMC, the HOLO-LD, HOLO-SDD, and the HOLO-BE method. The FOM values are relative to IMC with 300,000 histories per time step. The HOLO-BE method is significantly more efficient for comparable accuracy in this problem. The HOLO-LD method is much more efficient than the HOLO-SDD method because the LD trial space does not require error estimates at t^{n+1}. The doubly-discontinuous trial spaces were found to be exceptionally noisy for this problem, requiring 10^6 histories per time step to remove visible noise, although they were still more efficient than standard IMC, as demonstrated by the FOM. At 300,000 histories per time step the SDD trial space converges, but has large noise and inaccuracies due to oscillations introduced through noise in the time closure.

TABLE II: Comparison of FOM for the Marshak wave problem. Simulation end time is **t = 3.0** sh.

	FOM			
hists./step	IMC	HOLO-SDD	HOLO-LD	HOLO-BE
50,000	0.92	–	96.2	1708
300,000	1.00	0.43	200	2050
1,000,000	0.94	15.95	201	1806

V. CONCLUSIONS

We have extended a HOLO method to include a residual MC treatment of the time variable, producing accurate solutions to simplified TRT problems.. Results demonstrate that the inversion of the transport operator in the residual MC algorithm allows for an accurate projection of the solution across time steps, increasing accuracy in optically thin regions compared to a BE discretization. For most cases tested, a single batch simplifies the sampling space and leads to less variance overall than iterating on the residual. For optically thin problems, moments of the LO solution were found to agree with IMC when used with the LDD and SDD HO trial spaces; the HOLO method was more accurate and efficient than IMC, although sufficient mesh resolution is needed to limit projection error between time steps. For the Marshak wave problem, our ECMC algorithm is more statistically efficient than IMC, although the solution is not guaranteed to converge at lower particle history counts. We have demonstrated a new approach to closing the LO equations in time that produces accurate solutions with a HO solver that is sufficiently consistent. For the Marshak wave problem, the LD trial space in time is significantly more efficient than the doubly-discontinuous trial spaces, but inconsistencies lead to instabilities for thinner problems and finer meshes. Overall, each of the trial spaces has desirable and negative properties for a certain set of problems.

Future investigation of time-integration for this algorithm should include investigation of convergence in outer iterations when incorporating a time-averaged projection of the physical scattering source into the LO solver. Also, adaptive refinement in space and angle should be used to more accurately capture transient solutions. Investigation of the practicality and approach for adapting in time is necessary for obtaining a true ECMC algorithm with time-dependent transport equations. Second-order time-integration methods for the temperature could be incorporated. Although they will require solving a more difficult set of LO equations, the cost of the HO solver with the LDD trial space should be unaffected. Long-term, an ideal algorithm would require no projection of the HO solution between time steps, as in IMC, which will likely require an operator split on the radiation time derivative and a method to map the projected error back on to the census particles. Finally, the new sampling strategy for the LDD and LD trial spaces should be applied to higher-dimensional problems. Because of the difficulties with consistency and exceptionally difficult to resolve gradients in TRT problems, steady-state neutronics

problems should be investigated first.

APPENDIX

The Residual for the LDD Trial Space

The LDFE representation for the element with center (x_i, μ_j, t^n), spanning the interior of the domain $\mathcal{D}_{ijn}$, is

$$\tilde{I}(x,\mu,t) = I^n_{a,ij} + \frac{2}{h_i} I^n_{x,ij}\,(x - x_i) + \frac{2}{\Delta\mu} I^n_{\mu,ij}\left(\mu - \mu_j\right) + \frac{2}{\Delta t} I^n_{t,ij}\,(t - t^n) \tag{A.1}$$

where $\Delta\mu = \mu_{j+1/2} - \mu_{j-1/2}$ and $I^n_{a,ij}$, $I^n_{x,ij}$, $I^n_{\mu,ij}$, and $I^n_{t,ij}$ are the average, x slope, μ slope, and t slope moments, respectively [11]. We substitute Eq. (A.1) into Eq. (15) and analytically apply the continuous transport operator to determine the residual. For each cell, the residual is the sum of four components. The first is the interior volumetric source:

$$r^n_{ij}(x,\mu,t) = r^n_{a,ij} + \frac{2}{h_i} r^n_{x,ij}\,(x - x_i) + \frac{2}{\Delta\mu} r^n_{\mu,ij}\left(\mu - \mu_j\right) + \frac{2}{\Delta t} r^n_{t,ij}\,(t - t^n) \tag{A.2}$$

where

$$r^n_{a,ij} = q_{a,i} - \sigma_{a,i} I^n_{a,ij} - \frac{2}{c\Delta t} I^n_{t,ij} - \mu_j \frac{2}{\Delta\mu} I^n_{x,ij}, \tag{A.3}$$

$$r^n_{x,ij} = q_{x,i} - \sigma_{a,i} I^n_{x,ij}, \tag{A.4}$$

$$r^n_{\mu,ij} = -\sigma_{a,i} I^n_{\mu,ij} - \frac{\Delta\mu}{h_i} I^n_{x,ij}, \tag{A.5}$$

$$r^n_{t,ij} = -\sigma_{a,i} I^n_{t,ij} \tag{A.6}$$

where $q_{a,i}$ and $q_{x,i}$ are the zeroth and first moments of the LO LDFE emission source over the i-th spatial element ($q_{LO}(x)$ does not have a first moment in μ or t). The spatial derivative produces a δ-function source due to the discontinuities in the trial space. Upwinding is used to define the intensity on the inflowing face. The face source component, for $\mu > 0$, becomes

$$r^n_{i-1/2\,j} = \delta^+(x - x_{i-1/2})\mu\left[r^{n,f}_{a,ij} + \frac{2}{\Delta\mu} r^{n,f}_{\mu,ij}\left(\mu - \mu_j\right) + \frac{2}{\Delta t} r^{n,f}_{t,ij}\,(t - t^n)\right] \tag{A.7}$$

where

$$r^{n,f}_{a,ij} = \left(I^n_{a,i-1,j} + I^n_{x,i-1,j}\right) - \left(I^n_{a,ij} - I^n_{x,ij}\right) \tag{A.8}$$

$$r^{n,f}_{\mu,ij} = I^n_{\mu,i-1,j} - I^n_{\mu,ij}, \tag{A.9}$$

$$r^{n,f}_{t,ij} = I^n_{t,i-1,j} - I^n_{t,ij}. \tag{A.10}$$

The face sources for elements with $\mu < 0$ are similarly defined. The time-derivative jump source at $t^{n-1/2}$ is

$$r^{n-1/2}_{ij} = \delta^+(t - t^{n-1/2})\frac{1}{c}\left[r^{n-1/2}_{a,ij} + \frac{2}{h_i} r^{n-1/2}_{x,ij}\,(x - x_i) + \frac{2}{\Delta\mu} r^{n-1/2}_{\mu,ij}\left(\mu - \mu_j\right)\right] \tag{A.11}$$

where

$$r^{n-1/2}_{a,ij} = I^{n-1/2}_{a,ij} - \left(I^n_{a,ij} - I^n_{t,ij}\right), \tag{A.12}$$

$$r^{n-1/2}_{x,ij} = I^{n-1/2}_{x,ij} - I^n_{x,ij}, \tag{A.13}$$

$$r^{n-1/2}_{\mu,ij} = I^{n-1/2}_{\mu,ij} - I^n_{\mu,ij}. \tag{A.14}$$

Finally, the time-derivative component at $t^{n+1/2}$ is

$$r^{n-1/2}_{ij} = \delta^-(t - t^{n+1/2})\frac{1}{c}\left[r^{n-1/2}_{a,ij} + \frac{2}{h_i} r^{n-1/2}_{x,ij}\,(x - x_i) + \frac{2}{\Delta\mu} r^{n-1/2}_{\mu,ij}\left(\mu - \mu_j\right)\right] \tag{A.15}$$

where

$$r^{n+1/2}_{a,ij} = \left(I^n_{a,ij} - I^n_{t,ij}\right) - I^{n-1/2}_{a,ij}, \tag{A.16}$$

$$r^{n+1/2}_{x,ij} = I^{n-1/2}_{x,ij} I^n_{x,ij}, \tag{A.17}$$

$$r^{n+1/2}_{\mu,ij} = I^{n-1/2}_{\mu,ij} I^n_{\mu,ij}. \tag{A.18}$$

Particles sampled from this source would immediately reach census and terminate, thus there is no need to sample it. The analytic contribution from this source can be added, as detailed in [11], reducing variance.

Tallies for the Error

Tallies compute weighted moments of the error, averaged over the phase-space volume of each element. For both trial spaces, the time-averaged LDFE projection $\tilde{I}^n(x,\mu)$ is computed. This requires tallies for the average, x, and μ moments of the error. The tally for analog path-length sampling and the x moment of the error is

$$\hat{\epsilon}^n_{x,ij} = \frac{1}{N}\frac{6}{\Delta t h_i}\sum_{m=1}^{N_{score}} \frac{s_m}{h_i \Delta\mu} w_m\left(x^m_c - x_i\right), \tag{A.19}$$

where x^m_c is the center x-coordinate of the path length for the m-th particle history with length s_m and constant weight w_m, and N_{score} is the number of particles that have traversed the domain $\mathcal{D}_{ijn}$.. There are similar definitions for the average and μ moment. The tallies are derived by integrating the contribution of a differential path length, to the moment of interest, over the entire path length [11]. As in previous work, these tallies can be modified to allow for particles that stream without absorption and have exponentially attenuated weights [2].

For the SDD and LDD trial spaces, moments of $\epsilon(x,\mu,t^{n+1/2})$ must be estimated to produce a projection of the intensity at the end of the time step. For example, the x moment of the error at the end of time step is

$$\epsilon^{n+1/2}_{x,ij} = \frac{6}{h_i \Delta\mu \Delta t}\iiint_{\mathcal{D}_{ij}} \left(\frac{x - x_i}{h_i}\right)\epsilon(x,\mu,t^{n+1/2})\mathrm{d}x\mathrm{d}\mu \tag{A.20}$$

The estimators for these moments are a generalization of the census tallies used in IMC [6, 5]. The census estimator for the x moment is

$$\hat{\epsilon}^{n+1/2}_{x,ij} = \frac{1}{N}\frac{6}{\Delta\mu h_i^2}\sum_{n=1}^{N_{score}} c w_m\left(x_m - x_i\right), \tag{A.21}$$

where x_m is the coordinate of the m-th particle that has reached the end of the time step. Similar tallies are defined for the other space-angle moments. These tallies can be exceptionally noisy in optically thick cells because only particles that reach the end of the time step contribute.

For the LDD and LD time spaces, it is necessary to tally the slope of the intensity in t over each element, i.e., estimate Eq. (12). The estimator, for analog path-length sampling, is

$$\hat{\epsilon}^n_{t,ij} = \frac{1}{N}\frac{6}{h_i \Delta\mu \Delta t}\sum_{m=1}^{N_{score}}\left(\frac{t^m_c - t^n}{\Delta t}\right) w_m s_m \tag{A.22}$$

where t^m_c is the time of the particle at the center of the path length.

ACKNOWLEDGMENTS

This research was supported with funding received from the DOE National Nuclear Security Administration, under Award Number(s) DE-NA0002376.

REFERENCES

1. J. WILLERT and H. PARK, "Residual Monte Carlo High-order Solver for Moment-Based Accelerated Thermal Radiative Transfer Equations," *Journal of Computational Physics*, **276**, 405–421 (2014).
2. S. BOLDING, M. CLEVELAND, and J. MOREL, "A High-Order Low-Order Algorithm with Exponentially-Convergent Monte Carlo for Thermal Radiative Transfer," *Nuclear Science & Engineering: M&C 2015 Special Issue* (Jan. 2017).
3. A. B. WOLLABER, H. PARK, R. LOWRIE, R. RAUENZAHN, and M. CLEVELAND, "Multigroup Radiation Hydrodynamics with a High-Order, Low-Order Method," *Nuclear Science & Engineering: M&C 2015 Special Issue* (Jan. 2017).
4. J. WILLERT, C. KELLY, D. KNOLL, and H. PARK, "A Hybrid Approach to the Neutron Transport k-Eigenvalue Problem using NDA-based Algorithms," M&C. Sun Valley, ID (2013).
5. J. A. FLECK, JR. and J. D. CUMMINGS, JR., "An Implicit Monte Carlo Scheme for Calculating Time and Frequency Dependent Nonlinear Radiation Transport," *J. Comput. Phys.*, **8**, *3*, 313–342 (Dec. 1971).
6. A. B. WOLLABER, "Four decades of implicit Monte Carlo," *Journal of Computational and Theoretical Transport*, **45**, *1-2*, 1–70 (2016).
7. A. B. WOLLABER, E. W. LARSEN, and J. D. DENSMORE, "A Discrete Maximum Principle for the Implicit Monte Carlo Equations," *Nuclear Science and Engineering*, **173**, *3*, 259–275 (2013).
8. E. W. LARSEN, A. KUMAR, and J. E. MOREL, "Properties of the Implicitly Time-differenced Equations of Thermal Radiation Transport," *J. Comput. Phys.*, **238**, 82–96 (Apr. 2013).
9. E. WOLTERS, *Hybrid Monte Carlo - Deterministic Neutron Transport Methods Using Nonlinear Functionals*, Ph.D. thesis, Michigan (2011).
10. J. MOREL, T. WAREING, and K. SMITH, "Linear-Discontinuous Spatial Differencing Scheme for S_n Radiative Transfer Calculations," *Journal of Computational Physics*, **128**, 445–462 (1996).
11. S. BOLDING, *A High-Order Low-Order Algorithm with Exponentially-Convergent Monte Carlo for Thermal Radiative Transfer*, Ph.D. thesis, Texas A&M (2017).
12. J. PETERSON, J. MOREL, and J. RAGUSA, "Residual Monte Carlo for the One-Dimensional Neutron Transport Equation," *SIAM Journal on Scientific Computing*, accepted for publication.
13. J. PETERSON, *Exponentially Convergent Monte Carlo for the 1-D Transport Equation*, Master's thesis, Texas A&M (2014).
14. J. SHULTIS and W. DUNN, *Exploring Monte Carlo Methods*, Academic Press, Burlington, MA 01803 (2012).
15. J. A. FAVORITE and H. LICHTENSTEIN, "Exponential Monte Carlo convergence of a three-dimensional discrete ordinates solution," Tech. rep., Los Alamos National Lab., NM (US) (1999).
16. T. URBATSCH and T. EVANS, "Milagro Version 2: An Implicit Monte Carlo Code for Thermal Radiative Transfer: Capabilities, Development, and Usage," (2006), los Alamos National Laboratory Report LA-14195-MS.