Implicit Opacities in Thermal Radiative Transfer

Samuel Olivier [1,*], Quincy A. Huhn[2], James S. Warsa[1]

[1]Los Alamos National Laboratory, Los Alamos, NM; [2]Texas A&M University, College Station, TX

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

ABSTRACT

Currently, state-of-the-art numerical algorithms for solving thermal radiative transfer use a semi-implicit time integration scheme where opacities are treated explicitly (i.e. computed at the beginning of the time step). This explicit treatment places restrictions on the time step size to ensure stability and accuracy. Here, we investigate the accuracy, stability, and cost of time integration schemes that treat the opacity implicitly within the framework of the Second Moment Method. We demonstrate that the fully implicit time integration schemes lead to significant improvements in accuracy per cost on gray and frequency-dependent Marshak wave problems.

Keywords: implicit opacity; thermal radiative transfer; second moment method

1. INTRODUCTION

Thermal radiative transfer (TRT) describes the exchange of energy between thermal photons and matter. This energy transfer is governed by nuclear data known as opacities that describe the probability per unit distance of photons interacting with a material. Opacities are often strong functions of material temperature and photon energy. In particular, problems of interest typically begin at cold temperatures with large opacities that reduce as the material heats up. This results in severe spatial separation of the probability of interaction where cold parts of the spatial domain are millions or even billions of times more likely to absorb photons than hotter regions. In addition, opacities are typically more transparent to high-frequency radiation, allowing high-frequency photons to travel further and with faster dynamics than lower-frequency photons. Thus, opacities drive the spatial, spectral, and temporal behavior of the solution. Accurately evolving the opacities is then a crucial component of accurate TRT simulations.

The current state-of-the-art in numerical TRT algorithms use semi-implicit backward Euler temporal discretizations where absorption-emission is treated implicitly but the opacities are treated explicitly. In other words, the opacities are computed at the beginning of each time step using the previous time step's temperature solution. The nonlinear absorption-emission physics are then resolved assuming the opacities are constants independent of temperature. While this procedure does produce the fully nonlinear solution in the limit of small time steps, the explicit treatment of the opacities requires restrictions on the time step size to ensure temporal accuracy and stability. In the context of tightly coupled multiphysics simulations, this restriction could mean lowering the time step for the entire multiphysics suite, increasing the cost of every physics component.

In this paper, we investigate numerical algorithms for efficiently treating the opacities implicitly in time. We use a recently developed Second Moment Method (SMM) scheme [1] extended to TRT problems as a vehicle

*solivier@lanl.gov

for investigating the impact of implicitly integrated opacities on accuracy, stability, and cost. We proceed by providing brief background on the SMM algorithm, presenting the extension of SMM to the implicit treatment of opacities, and giving numerical results demonstrating that the fully implicit schemes are more efficient than semi-implicit in accuracy per cost. Finally, we give conclusions and recommendations for future work.

2. SECOND MOMENT METHODS FOR TRT

The TRT system of equations is given by

$$\frac{1}{c}\frac{\partial I_g}{\partial t} + \mathbf{\Omega}\cdot\nabla I_g + \sigma_g I_g = \frac{\sigma_g B_g(T)}{4\pi}\,, \tag{1a}$$

$$C_v\frac{\partial T}{\partial t} = \sum_g c\sigma_g E_g - \sigma_g B_g(T)\,, \tag{1b}$$

where I_g is the radiation intensity in group g, $\mathbf{\Omega}$ the direction of flight, $E_g = \frac{1}{c}\int I_g \,\mathrm{d}\Omega$ the energy density, σ_g the opacity in group g, T the material temperature, B_g the Planck emission function integrated over group g, c the speed of light, and C_v the heat capacity of the material. Here, we have neglected scattering so that there is only a single total interaction opacity and omit discussion of boundary conditions for brevity. We refer to Eq. 1a as the "high-order" (HO) problem due to its high dimensionality and Eq. 1b as the energy balance equation. The HO transport equation is discretized in angle with the Discrete Ordinates (S_N) method. Lumped linear Discontinuous Galerkin (DG) is used to discretize the resulting multigroup, S_N transport and energy balance equations in space.

The SMM algorithm couples a reduced-dimensional, diffusion-based, low-order (LO) model to the energy balance equation as a means to iteratively accelerate the computation of the nonlinear absorption-emission physics. Here, we use a gray LO system to further reduce the expense of solving the LO system. The gray LO problem takes the form

$$\frac{\partial E}{\partial t} + \nabla\cdot\boldsymbol{F} + c\sigma_E E = \sigma_P B(T)\,, \tag{2a}$$

$$\frac{1}{c}\frac{\partial \boldsymbol{F}}{\partial t} + \frac{c}{3}\nabla E + \sigma_R \boldsymbol{F} = -\nabla\cdot\mathbf{T} + \sum_g \int \mathbf{\Omega}\left(\sigma_R - \sigma_g\right) I_g \,\mathrm{d}\Omega\,, \tag{2b}$$

where E and $\boldsymbol{F}$ are the LO gray energy density and radiative flux, respectively, and $\mathbf{T}$ is the SMM correction tensor

$$\mathbf{T} = \sum_g \int \left(\mathbf{\Omega}\otimes\mathbf{\Omega} - \frac{1}{3}\mathbf{I}\right) I_g \,\mathrm{d}\Omega\,. \tag{3}$$

We define the gray opacities σ_E, σ_R, and σ_P as the energy density, Rosseland, and Planck-weighted opacities, respectively, as follows

$$\sigma_E = \frac{\sum_g \sigma_g E_g}{\sum_g E_g}\,, \quad \sigma_R = \frac{\sum_g \sigma_g \frac{\partial B_g}{\partial T}}{\sum_g \frac{\partial B_g}{\partial T}}\,, \quad \sigma_P = \frac{\sum_g \sigma_g B_g}{\sum_g B_g}\,. \tag{4}$$

In addition, the gray Planck emission function $B(T) = \sum_g B_g(T)$ is used in the LO and energy balance systems. Note that Eq. 2 represents a diffusion system with extra source terms such that the LO and HO solutions match. In particular, $-\nabla\cdot\mathbf{T}$ and $\sum_g \int \mathbf{\Omega}\left(\sigma_R - \sigma_g\right) I_g \,\mathrm{d}\Omega$ correct for the diffusion approximation and the use of σ_R in the first moment, respectively. The LO system is discretized in space using a recently

developed consistent Local Discontinuous Galerkin (LDG) finite element technique [1]. This discretization is designed so that the discrete LO and HO solutions match to iterative tolerances while allowing the use of scalable linear solvers. Details on the derivation of SMM and the multigroup aspects of moment acceleration schemes applied to TRT can be found in [2, 3].

The SMM algorithm simultaneously solves

$$\frac{1}{c}\frac{\partial I_g}{\partial t} + \mathbf{\Omega}\cdot\nabla I_g + \sigma_g I_g = \frac{\sigma_g B_g(T)}{4\pi}\,, \tag{5a}$$

$$\frac{\partial E}{\partial t} + \nabla\cdot\boldsymbol{F} + c\sigma_E E = \sigma_P B(T)\,, \tag{5b}$$

$$\frac{1}{c}\frac{\partial \boldsymbol{F}}{\partial t} + \frac{c}{3}\nabla E + \sigma_R \boldsymbol{F} = -\nabla\cdot\mathbf{T} + \sum_g \int \mathbf{\Omega}\left(\sigma_R - \sigma_g\right) I_g \,\mathrm{d}\Omega\,, \tag{5c}$$

$$C_v\frac{\partial T}{\partial t} = c\sigma_E E - \sigma_P B(T)\,. \tag{5d}$$

Note that the energy balance equation has been modified to use the energy density computed from the LO system. The semi-implicit solution algorithm at each time step proceeds as follows. Opacities are computed from the initial or previous time step's temperature and remain fixed until the next time step. The HO problem is solved given an emission source. The SMM sources are computed. The LO and energy balance equations are then simultaneously solved using Newton iteration with nonlinear elimination of the temperature. That is, Eq. 5d is nonlinearly solved with Newton iteration each time a new LO energy density is computed. The temperature from the LO solve is used to form a new emission source for the HO problem. The outer and inner iterations are terminated when successive temperature and gray energy density solutions meet a residual-based relative stopping criterion.

3. EXTENSION TO IMPLICIT OPACITIES

Letting superscript denote time step index, the semi-implicit time integration scheme evaluates all physics at I_g^{n+1}, E^{n+1}, $\boldsymbol{F}^{n+1}$, and T^{n+1} except the multigroup opacity which is evaluated at T^n. The fully implicit algorithm modifies the semi-implicit algorithm by additionally evaluating the opacities at T^{n+1}. The implicit algorithm's additional nonlinearity in $\sigma_g(T^{n+1})$ must be resolved with iteration. We assume that we do not have access to the derivative of the opacity with respect to temperature. Thus, we use $\sigma_g(T_k^{n+1})$ where T_k^{n+1} is the temperature from a previous nonlinear iteration. As mentioned above, the SMM algorithm has an outer iteration where the HO system is solved and a nested, inner iteration where the LO system is solved given the solution from the outer iteration. Thus, there are two possible places for the opacity to be updated: at the end of each outer iteration or the end of each nested, inner iteration. We refer to these methods as Implicit-HO and Implicit-LO, respectively, in reference to where the opacity nonlinearity is resolved.

For Implicit-HO, the opacity is updated at the end of each outer, HO iteration. This acts to fold the resolution of the temperature-dependence of the opacity into the HO solve, resulting in a fixed-point iteration where both the opacity and the HO problem are converged simultaneously. Since the nonlinearity of the fixed-point problem has increased, we expect Implicit-HO to require more HO iterations than the semi-implicit scheme.

For Implicit-LO, the LO and energy balance iteration is combined with the opacity iteration. Note that the inner solve utilizes a Newton iteration to accelerate its convergence. Since we do not use the derivative of the opacity, this combined opacity and absorption-emission iteration is effectively a modified Newton algorithm where absorption-emission is linearized but the opacity is not. Since the opacity is resolved in the LO solve, we expect the LO solve to accelerate both the absorption-emission and opacity nonlinearities, resulting in

more LO iterations but reduced HO iterations compared to Implicit-HO. This acceleration comes at the cost of significantly more calls to evaluate the opacities than Implicit-HO as Implicit-LO updates the opacities inside a nested loop.

Note that the gray opacities σ_E, σ_R, and σ_P have weight functions that *are* treated implicitly regardless of whether σ_g is. All methods update σ_E at each outer iteration to ensure the absorption rates between HO and LO match. In the explicit and Implicit-HO methods, σ_R and σ_P are updated before each inner solve and remain iteratively fixed until the next HO iteration. For Implicit-HO, this ensures the LO system is not iterating on the Planck and Rosseland collapsing spectra, improving LO efficiency. In contrast, the Implicit-LO method updates σ_E, σ_R, and σ_P in the LO iteration each time σ_g is updated. The details of these methods are codified in Algorithm 1.

Algorithm 1 SMM algorithm with Explicit, Implicit-HO, or Implicit-LO opacity treatment.

```
while outer not converged do
  Compute emission given previous temperature iterate
  Solve HO problem with transport sweep on emission source
  Compute gray absorption opacity, σ_E
  if Explicit or Implicit-HO then
    Compute σ_R and σ_P
  end if
  Compute SMM sources
  while inner not converged do
    Solve linearized diffusion system
    Solve energy balance nonlinearly given absorption source
    if Implicit-LO then
      Compute σ_g, σ_E, σ_R, and σ_P
    end if
  end while
  if Implicit-HO then
    Compute σ_g
  end if
end while
if Explicit then
  Compute σ_g
end if
```

4. NUMERICAL RESULTS

We now compare the performance of the Explicit, Implicit-HO, and Implicit-LO opacity treatments on two one-dimensional benchmark problems with analytic opacities. Accuracy is assessed on fixed spatial, angular, and energy grids by comparing the solutions generated with a range of time step sizes to a reference solution generated with a much smaller time step size. The methods were implemented using the MFEM finite element framework [4]. Algebraic multigrid-preconditioned conjugate gradient is used to solve the LO diffusion system. Both the outer and inner iterations are solved to a residual-based relative tolerance of 10^{-3}.

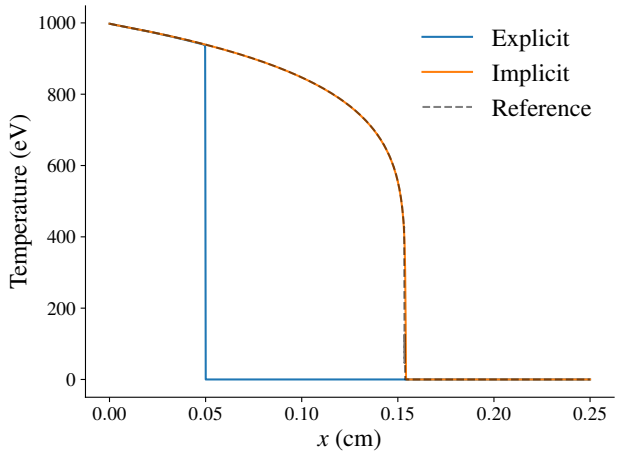

Figure 1. Temperature at $t = 1$ sh computed using explicit and implicit opacity.

4.1. Gray Marshak Wave

The first problem is a gray Marshak wave with strongly temperature-dependent opacity of the form

$$\sigma(T) = 10^{12}/T^{3.2}\,. \tag{6}$$

The exponent of 3.2 is chosen so that the product of the opacity and the Planck emission function remains nonlinear in temperature. The heat capacity is $C_v = 3 \times 10^{12}\,\mathrm{erg\,cm^{-3}\,eV^{-1}}$. The radiation field and material temperature are initially in thermal equilibrium at $T = 0.025\,\mathrm{eV}$. The domain is $[0, 0.25\,\mathrm{cm}]$ with a boundary source of $T = 1\,\mathrm{keV}$ impinging on the domain at $x = 0$. The domain is meshed into 1000 equal sized elements resulting in $h = 2.5 \times 10^{-4}\,\mathrm{cm}$. S_6 Gauss-Legendre angular quadrature is used. The final time is $t = 1\,\mathrm{sh}$. In the results that follow, the reference solution is generated with the Implicit-HO time integration scheme using a time step of size $\Delta t = 2 \times 10^{-5}\,\mathrm{sh}$. On this problem, the initial opacity is $\sigma(0.025) = 1.33 \times 10^{17}\,\mathrm{cm^{-1}}$ and is thus extremely opaque to radiation. However, as the material heats, the opacity quickly drops allowing radiation to penetrate further into the domain, resulting in a sharp wave traveling through the domain from left to right. Figure 1 plots the temperature at the final time of $t = 1\,\mathrm{sh}$ using a time step of size $\Delta t = 2 \times 10^{-3}\,\mathrm{sh}$. The Implicit-HO time integration scheme is used for the implicit opacity solution. With this time step size, the explicit treatment of the opacity results in stagnation of the wave. Compared to the reference, the explicit opacity solution has a steeper profile and a slower wave speed. This is due to the opacity being frozen within each time step, causing the front of the wave to see a colder, more opaque opacity is it travels through each cell. By contrast, the implicit opacity scheme results in a solution visibly matching the reference solution despite the large time step size.

Figure 2a depicts the error with respect to the reference solution over a range of time step sizes. For small enough time steps, all methods converge with the expected first-order accuracy. Data points beyond $\Delta t = 2 \times 10^{-3}\,\mathrm{sh}$ could not be obtained for the explicit opacity scheme due to temporal discretization instability. The Implicit-HO and Implicit-LO schemes converge identically for all time steps. For the smallest time step size, implicit opacities results in an 8x more accurate solution. The implicit methods are able to still produce an accurate answer even when the explicit opacity solution is stagnating, resulting in greater than 50x improvements in accuracy at large time step sizes.

Figure 2b shows the iterative performance of the three schemes. The average number of sweeps per time step is shown in solid lines and the average number of inner LO solves per outer iteration is shown in dashed lines.

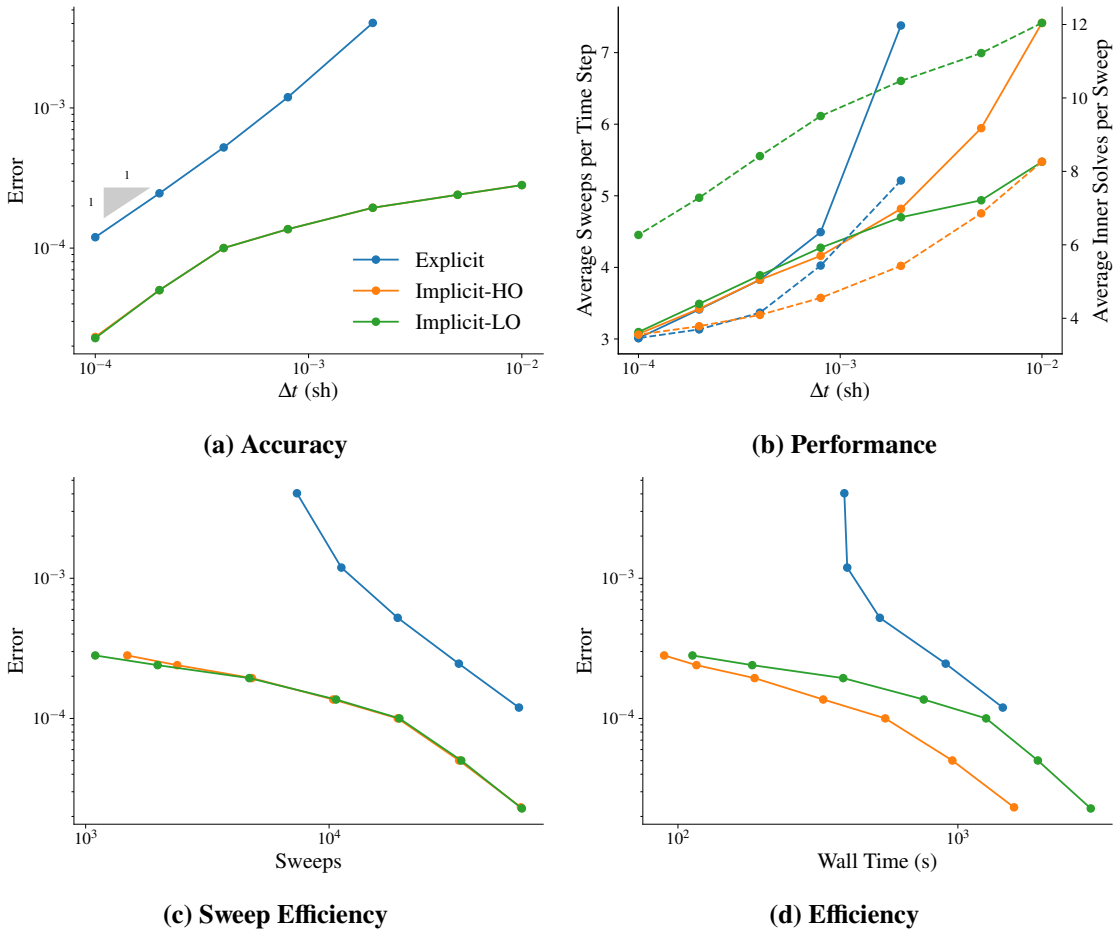

(a) Accuracy **(b) Performance**

(c) Sweep Efficiency **(d) Efficiency**

Figure 2. Accuracy, iterative performance, and efficiencies on the gray Marshak problem.

For the explicit method, both sweeps and LO solves increase as the time step becomes large. This is due to the steepening of the profile resulting in non-physically stiffer absorption-emission physics at the wave front. The placement of the opacity update results in the Implicit-HO method requiring more sweeps and fewer inner solves than Implicit-LO at large time steps while Implicit-LO has more LO solves and fewer sweeps than Implicit-HO, as expected. Accuracy per sweep and accuracy per runtime are presented in Figs. 2c and 2d. On all time steps, the implicit opacity methods are more efficient than the explicit opacity method. In terms of runtime, Implicit-HO is the most efficient. Implicit-LO updates the opacity more than Implicit-HO and thus incurs more cost, resulting in lower efficiency. However, in terms of sweep efficiency, Implicit-LO has a slight advantage for large time steps. Note that the sweep has lower scalability in higher dimensions than the element-local opacity update operations, especially with spatial decomposition in parallel. We thus expect Implicit-LO to be more competitive in efficiency on 2/3D calculations where the sweep cost dominates the cost of all other components of the algorithm.

4.2. Frequency-Dependent Marshak Wave

The second test problem is a Marshak wave problem with frequency-dependent opacity and Planck emission. As with the gray problem, there is a boundary source at $T = 1\,\text{keV}$, the radiation and material are initially

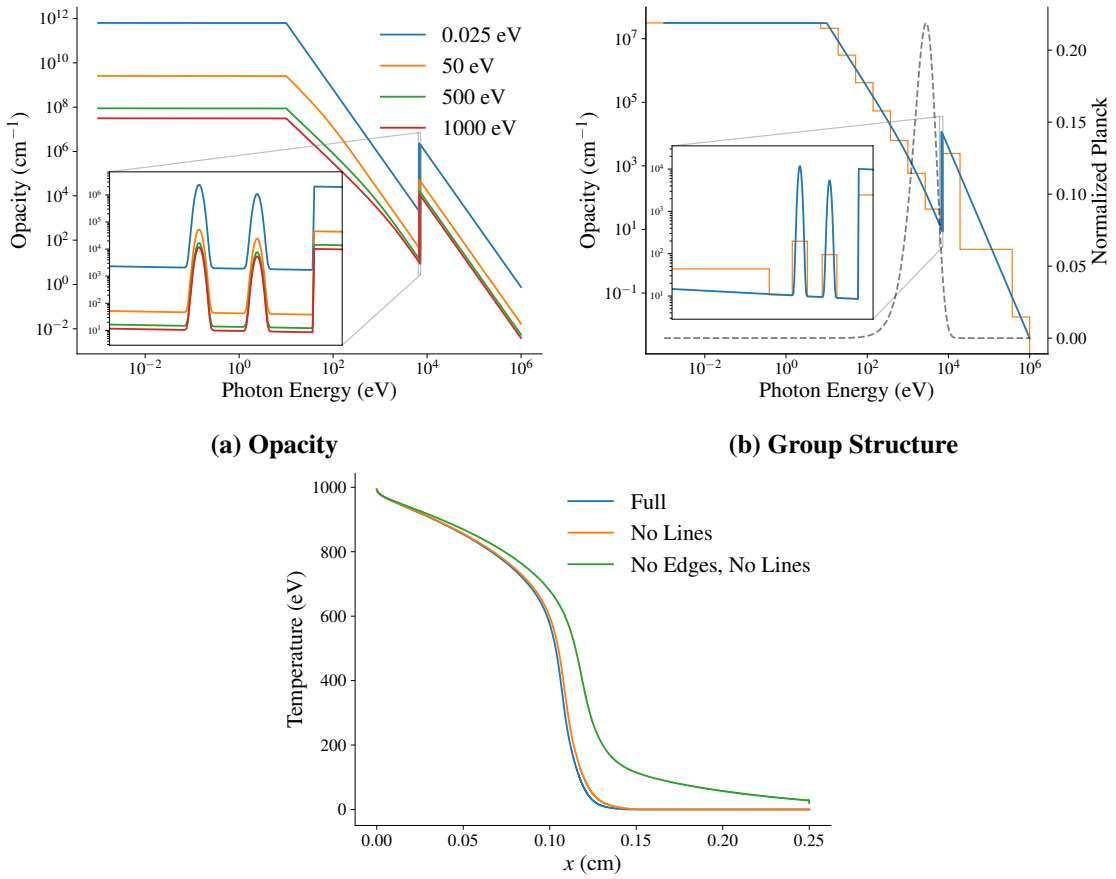

(a) Opacity **(b) Group Structure**

(c) Solution

Figure 3. Plots of the opacity, group structure, and resulting solutions for the multi-group Marshak problem.

in equilibrium at $T = 0.025\,\text{eV}$, S_6 quadrature is used, the domain is $[0, 0.25\,\text{cm}]$, and the mesh has 1000 elements. However, now the final time is $t = 3\,\text{sh}$.

Following Brunner (2023) [5], the opacity is defined to mimic the presence of the "edges" and "lines" seen in real opacity data. An edge is modeled as a step function that increases the opacity by a fixed amount past a certain photon energy, ϵ_{edge}. Lines are modeled using a series of Gaussians centered at fixed intervals before an edge. In addition, [5] also proposes limiting the opacity at lower energies to prevent the opacity becoming unbounded at small photon energies. The analytic form of the opacity is

$$\sigma(\epsilon, T) = \frac{C_0}{\sqrt{T}\hat{\epsilon}^3}\left(1 - e^{\hat{\epsilon}/T}\right)\left[1 + C_1 H(\hat{\epsilon} - \epsilon_{\text{edge}}) + C_2 \sum_{\ell=0}^{N_l-1} \frac{1}{N_l - 1} \exp\left(-\frac{1}{2}\left(\frac{\hat{\epsilon} - (\epsilon_{\text{edge}} - (\ell+1)\delta_s)}{\delta_w}\right)^2\right)\right], \tag{7a}$$

$$\hat{\epsilon} = \max(\epsilon_{\min}, \epsilon)\,, \tag{7b}$$

where $\epsilon = h\nu$ is the photon energy, $H(x)$ the Heaviside function that is zero for $x < 0$ and unity for $x \geq 0$, ϵ_{edge} the location of the edge, $\epsilon_{\min}$ the minimum energy, N_l the number of lines, and δ_s and δ_w control the spacing and spread of the lines, respectively. Here, we use $C_0 = 1 \times 10^{14}\,\text{eV}^{3.5}/\text{cm}$, $C_1 = C_2 = 1200$,

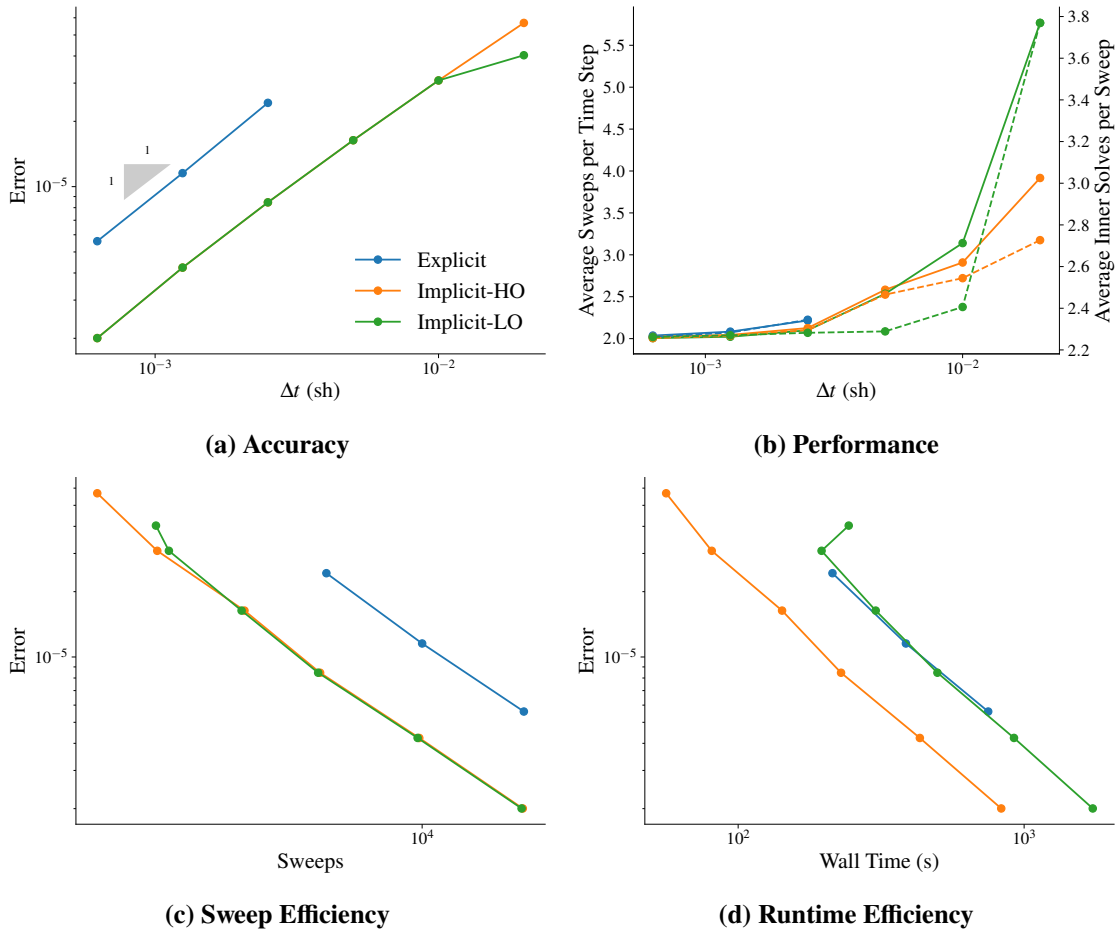

(a) Accuracy **(b) Performance**

(c) Sweep Efficiency **(d) Runtime Efficiency**

Figure 4. Accuracy, iterative performance, and efficiency on the multi-group Marshak problem.

$N_l = 2$, $\epsilon_{\text{min}} = 10\,\text{eV}$, $\epsilon_{\text{edge}} = 7196\,\text{eV}$, $\delta_s = 200\,\text{eV}$, and $\delta_w = 10\,\text{eV}$. The opacity is plotted at a range of relevant temperatures spanning from the initial equilibrium temperature to the boundary source temperature in Fig. 3a. An inset is provided to zoom into the line features.

Figure 3b shows the opacity at $T = 1\,\text{keV}$ and the group structure used on this problem. The opacity in each group is computed with three-point Gauss-Legendre quadrature. The depicted 29-group energy grid is composed of a base grid of 20 logarithmically spaced groups between 1×10^{-3} eV and 1×10^{6} eV with one group added to represent the low-frequency interval $[0, 1 \times 10^{-3}\,\text{eV}]$ and 8 groups added to resolve the edge and lines. The Planck emission spectrum at $T = 1\,\text{keV}$ is also included to show that our placement of the edge and lines overlaps with the tail end of the boundary source's emission spectrum, which will result in more high-frequency absorption. The effects of these opacity features on the solution are shown in Fig. 3c. Here, "Full" refers to the opacity as specified in Eq. 7, "No Lines" to the case of $C_1 = 0$, and "No Edges, No Lines" to $C_1 = C_2 = 0$. Observe that the inclusion of the opacity edge steepens the Marshak wave profile by making the material more opaque to photons with energy above ϵ_{edge}. The inclusion of lines has a less pronounced effect but also results in a steeper wave due to analogous arguments.

The accuracy over a range of time step sizes is plotted in Fig. 4a. The reference solution is taken to be Implicit-HO with a time step of size $\Delta t = 1 \times 10^{-4}\,\text{sh}$. First-order accuracy is seen for all methods and

time step sizes aside from Implicit-LO at the largest time step size where the SMM algorithm struggled to converge in the first few time steps, reducing the accuracy of the method. On this frequency-dependent problem, implicit opacities results in a smaller 3x improvement in accuracy over explicit opacities. Note that stagnation of the solution was not observed. Instead, the explicit scheme became immediately unstable for time steps sizes larger than 2.5×10^{-3} sh. This difference in behavior is likely due to the use of a *gray* LO system. For large time steps, the gray LO system cannot effectively accelerate the spectrum of the HO system, resulting in reduced robustness. On the gray problem, the gray LO system fully accelerates the HO system.

The average numbers of sweeps per time step and LO solves per outer iteration are depicted in Fig. 4b where, again, solid lines represent sweeps and dashed lines LO solves. For small time steps the methods are indistinguishable, taking similar numbers of sweeps and LO solves. For larger time step sizes, Implicit-LO no longer accelerates the HO iteration requiring a similar number of sweeps per time step as Implicit-HO, a marked change from the performance on the gray Marshak problem. This difference in behavior is likely due to the fact that the opacity is being updated but the spectrum is fixed from the most recent HO solve. Implicit-LO's reduced effectiveness is apparent in the sweep efficiency shown in Fig. 4c where the largest time step shows Implicit-HO being more efficient than Implicit-LO. However, both implicit opacity treatments are more efficient in sweeps than explicitly treating the opacities. Finally, efficiency in total runtime is shown in Fig. 4d. Here, Implicit-HO is the most efficient with Implicit-LO being about as efficient as explicit opacities. Implicit-LO is more impacted by the increased cost of computing multigroup opacities since Implicit-LO updates the opacity at each LO iteration. In 2/3D, we expect this cost to be less pronounced with respect to the sweep, resulting in Implicit-LO being more efficient than explicit.

5. CONCLUSIONS

We presented an extension of the Second Moment Method (SMM) applied to thermal radiative transfer (TRT) problems that allows treating the opacity implicitly, resulting in two novel schemes that are fully implicit in time. The schemes differ only in where the opacity is updated in the SMM solution procedure. Implicit-HO updates the opacity in the outer, high-order iteration whereas Implicit-LO updates the opacity inside the nested, low-order iteration. Implicitly treating the opacity resulted in 8-50x improvement in accuracy on a gray Marshak problem and 3x improvement on a frequency-dependent Marshak problem. In both problems, implicitly treating the opacity significantly extended the region of stability, allowing SMM to produce accurate answers at very large time steps where explicit treatment was unusably inaccurate or unstable. On both problems, Implicit-HO was the most efficient of the three schemes in terms of accuracy per runtime, meaning the improvements in accuracy from implicitly treating the opacity were greater than the associated costs. The effectiveness of Implicit-HO, which simply uses a fixed-point iteration on the opacity, suggests that acceleration may not be required and thus any TRT algorithm could incorporate implicit opacities in an analogously simple manner and potentially see similar improvements in efficiency compared to a semi-implicit method. We plan to extend this work to include real, tabular opacities on 2D multi-material problems and to include comparisons to other TRT algorithms.

ACKNOWLEDGEMENTS

This research used resources provided by the Darwin testbed at Los Alamos National Laboratory (LANL) which is funded by the Computational Systems and Software Environments subprogram of LANL's Advanced Simulation and Computing program. S.O. was supported by the U.S. Department of Energy as a Nicholas C. Metropolis Fellow under the Laboratory Directed Research and Development program of the Los Alamos National Laboratory. LANL report number LA-UR-24-31438. LANL is operated by Triad National Security, LLC, for the National Nuclear Security Administration of the U.S. Department of Energy (Contract

No. 89233218CNA000001).

REFERENCES

[1] S. Olivier, B. S. Southworth, J. S. Warsa, and H. Park. "Consistent Second Moment Methods with Scalable Linear Solvers for Radiation Transport." (2024). URL https://arxiv.org/abs/2404.17473.

[2] S. Olivier and T. S. Haut. "High-Order Finite Element Second Moment Methods for Linear Transport." *Nuclear Science and Engineering*, **volume 0**(0), pp. 1–36 (2023).

[3] D. Y. Anistratov. "Stability analysis of a multilevel quasidiffusion method for thermal radiative transfer problems." *Journal of Computational Physics*, **volume 376**, pp. 186–209 (2019).

[4] "MFEM: Modular Finite Element Methods [Software]." https://mfem.org (2010).

[5] T. A. Brunner. "A Family of Multi-Dimensional Thermal Radiative Transfer Test Problems." URL https://www.osti.gov/biblio/2280904.