Journal of Computational Physics 229 (2010) 5597–5614

Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

Robust and accurate filtered spherical harmonics expansions for radiative transfer

Ryan G. McClarren [a,*], Cory D. Hauck [b]

[a] Department of Nuclear Engineering, Texas A&M University, College Station, TX 77843-3133, USA
[b] Computational Mathematics Group, Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA

ARTICLE INFO

Article history:
Received 4 December 2009
Received in revised form 29 March 2010
Accepted 29 March 2010
Available online 4 April 2010

Keywords:
Radiative transfer
Spherical harmonics method

ABSTRACT

We present a novel application of filters to the spherical harmonics (P_N) expansion for radiative transfer problems in the high-energy-density regime. The filter we use is based on non-oscillatory spherical splines and a filter strength chosen to (i) preserve the equilibrium diffusion limit and (ii) vanish as the expansion order tends to infinity. Our implementation is based on modified equations that are derived by applying the filter after every time step in a simple first-order time integration scheme. The method is readily applied to existing codes that solve the P_N equations. Numerical results demonstrate that the solution to the filtered P_N equations are (i) more robust and less oscillatory than standard P_N solutions and (ii) more accurate than discrete ordinates solutions of comparable order. In particular, the filtered P_7 solution demonstrates comparable accuracy to an implicit Monte Carlo solution for a benchmark hohlraum problem in 2D Cartesian geometry.

1. Introduction

Radiation plays an important role in several high-energy-density physics contexts such as neutrino transport in core-collapse supernovae [1,2] or thermal radiative transfer (via photons) in inertial confinement fusion [3,4] and terrestrial radiating shock experiments [5–7]. In these examples, radiation is a major conduit for energy exchange between different materials in the system. Consequently, the fidelity of any radiation-hydrodynamics calculation will rely heavily on accurate simulations of time-dependent radiation transport.

A common challenge for thermal radiative transfer computations in high-energy-density problems is the presence of a material medium that is characterized by a wide-ranging optical thicknesses. Such is the case for experimental setups that have regions of vacuum situated between optically thick materials or when there are large spatial variations in the radiation energy-density, since the material opacity is generally smaller for high-energy photons than for low-energy photons. Material interaction is the primary coupling mechanism for radiation energy. Thus in regions where a material is transparent to the radiation, photons move for long distances without their direction of flight being changed through collisions with the material. The lack of redistribution of photon direction leads to so-called "ray effects" in discrete ordinates (S_N) solutions to radiative transfer problems [8,9] in which the angular grid leaves an imprint on the numerical solution. These effects can be partially mitigated by increasing the number of ordinates or by using biased quadrature sets to resolve important areas of the problem [10].

Stochastic methods, such as the implicit Monte Carlo method of Fleck and Cummings [11], are widely considered to be the most accurate method for radiative transfer in the high-energy-density regime. These methods do not suffer from ray effects,

* Corresponding author. Tel.: +1 979 845 8311.
E-mail addresses: rgm@tamu.edu (R.G. McClarren), hauckc@ornl.gov (C.D. Hauck).

0021-9991/$ - see front matter © 2010 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2010.03.043

but they are not without non-physical artifacts. Monte Carlo solutions are characterized by noise due to a finite sampling of the phase space. Moreover, the implicit Monte Carlo method can allow non-physical overheating of the material medium when the radiation is strongly coupled to the material through absorption [12,13].

A third path for radiative transfer simulations is the spherical harmonics (P_N) method. This is a spectral method, based on a linear expansion of the radiation intensity in angle using spherical harmonic basis functions. The method yields a hyperbolic system of partial differential equations for the expansion coefficients, which are just the angular moments with respect to the basis (up to some normalization constants).

The spherical harmonic expansion has formal spectral convergence to the solution of the full radiative transfer equation and preserves the property of rotational invariance.[1] This contrasts with the discrete ordinates method, where a lack of rotational invariance manifests itself in the aforementioned ray effects. As a hyperbolic system, the P_N equations approximate the movement of radiation as a series of waves. The speed of these waves is bounded by the speed of light. This restriction is physically reasonable and consistent with the transfer equations, unlike flux-limited diffusion methods where photons are allowed to travel at unbounded speeds.

The spherical harmonic method does have its own shortcomings. For example, in streaming regions, where interactions with the material are rare, solutions to the hyperbolic system have non-physical oscillations. These oscillations are numerical artifacts – a by-product of approximating a non-smooth function with a smooth basis. A catastrophic consequence of these oscillations is that they can cause the radiation energy-density to become negative which, when the radiative transfer is coupled to a material equation [14,15], can cause the material temperature to become negative. Worse still, it has been shown that these negative solutions can arise in any finite-order spherical harmonics approximation [14].

There has been much previous work to address negative radiation energy densities associated with the spherical harmonics method. Several closures have been proposed to truncate the series in a more physical or better-behaved way [16–21]. One currently fashionable approach is to use maximum entropy methods which generate *nonlinear* expansions with spherical harmonic functions. However, most of these treatments are designed to close low order expansions and may not provide a straightforward means to correct higher-order spherical harmonics expansions, or if they do extend to general P_N expansions, deriving and solving the resulting equations for a general expansion is onerous and may be computationally unstable [22–24]. Other attempts to correct the negativity in the P_N equations include adding artificial scattering terms to make the intensity more isotropic [15] or introducing a floor for the radiation energy density. A general closure for the P_N equations based on solving a local optimization problem is the subject of concurrent work of the authors as well [25]. However, this closure is also computationally intensive.

Thus, in this work, and in our preliminary results presented at a recent conference [26], we detail a method for improving spherical harmonic solutions that is simple to extend to high-order expansions, gives formal convergence to the underlying transport solution, and preserves the equilibrium diffusion limit. Using our method, we are able to solve problems of radiative transfer that could not be solved with standard P_N methods because they caused the material temperature to become negative. An attractive feature of our method is that it can be implemented in existing codes that solve the P_N equations.

Our approach to making spherical harmonics methods more robust is to use a filter to remove oscillations from the reconstruction of the intensity. By removing these oscillations, we are able to avoid negative solutions while maintaining angular accuracy. We discuss general properties that a filter should possess and then detail a specific example. The filter that we present is by no means the only possibility, nor do we claim that it is optimal. Nevertheless, we believe that the idea of applying filters to the spherical harmonics expansion is a efficient and effective way to make the method more robust and useful for real-world, predictive simulations.

2. The spherical harmonics expansion

Before presenting filtered spherical harmonics, we recall briefly the equations of radiative transfer and the standard spherical harmonics expansion. The equations we are ultimately interested in solving are the radiative transfer equation [27]

$$\frac{1}{c}\frac{\partial I}{\partial t} + \Omega \cdot \nabla_{\mathbf{x}} I + \sigma_{\mathrm{t}} I = \sigma_a \frac{ac}{4\pi} T^4 + \frac{\sigma_{\mathrm{s}}}{4\pi}\langle I\rangle \tag{1a}$$

and an energy equation for a stationary material

$$\frac{\partial}{\partial t} e(T) = \sigma_a(\langle I\rangle - acT^4). \tag{1b}$$

In these equations, $I(\mathbf{x},\Omega,t)$ is the specific intensity, $e(T)$ is the internal energy-density, $T(\mathbf{x},t)$ is the material temperature, and c is the speed of light. The radiation constant a is related to the Stefan–Boltzmann constant σ_{sb} by the relation $a = 4\sigma_{sb}/c$. The material medium is characterized by absorption, scattering, and total opacities (σ_a, σ_s, and $\sigma_{\mathrm{t}} = \sigma_a + \sigma_s$ respectively) and by an equation of state $C_v(T) = e'(T)$. The independent variables are the spatial variable $\mathbf{x} = (x,y,z) \in \mathbb{R}^3$, the angular variable

[1] Roughly speaking, this property says that solutions to the transfer equation are invariant under a common orthogonal transformation of the spatial and angular components of phase space.

$\Omega \in \mathbb{S}^2$ (that is, Ω is a point on the unit sphere), and the temporal variable $t \in [0,\infty)$. In spherical coordinates, Ω is specified by a polar angle $\theta \in [0,\pi]$ and an azimuthal angle $\varphi \in [0,2\pi)$. For notational convenience we have written integration over the unit sphere as

$$\langle I \rangle \equiv \int_{\mathbb{S}^2} I\, d\Omega = \int_0^{2\pi} \int_{-1}^{1} I\, d\mu d\varphi,$$

where $\mu := \cos\theta \in [-1,1]$.

Eq. (1) does not take into account photon frequency, non-isotropic scattering such as Compton scattering, or material motion. These additional complications can be handled in our method, but the additional details required would obfuscate the main thrust of this paper. In the remainder of the paper, we will refer to exact solutions of Eq. (1) as the *transport solution*.

The Nth-order spherical harmonic expansion of I is given by

$$I \approx \sum_{l=0}^{N} \sum_{m=-l}^{l} Y_l^m(\mu,\varphi) I_l^m, \tag{2}$$

where, for integers l and m, Y_l^m is the spherical harmonic function of order l and degree m, given by

$$Y_l^m = (-1)^m \sqrt{\frac{2l+1}{4\pi}\frac{(l-m)!}{(l+m)!}} P_l^m(\mu) e^{im\varphi}, \qquad 0 \leqslant l, \quad 0 \leqslant m \leqslant l$$

and P_l^m is an associated Legendre function. For $m < 0$, $\overline{Y}_l^m = (-1)^m Y_l^{-m}$. The spherical harmonic moments of I are defined $\overline{Y}$ as

$$I_l^m := \langle \overline{Y}_l^m I \rangle,$$

where the overbar denotes the complex conjugate.

By taking spherical harmonic moments of Eq. (1a) and using recursion relations for the spherical harmonics, one can write the P_N equations as [27,28]

$$\begin{aligned} &\frac{1}{c}\frac{\partial}{\partial t} I_l^m + \frac{1}{2}\frac{\partial}{\partial x}\left(-C_{l-1}^{m-1} I_{l-1}^{m-1} + D_{l+1}^{m-1} I_{l+1}^{m-1} + E_{l-1}^{m+1} I_{l-1}^{m+1} - F_{l+1}^{m+1} I_{l+1}^{m+1}\right) + \frac{i}{2}\frac{\partial}{\partial y}\left(C_{l-1}^{m-1} I_{l-1}^{m-1} - D_{l+1}^{m-1} I_{l+1}^{m-1} + E_{l-1}^{m+1} I_{l-1}^{m+1} - F_{l+1}^{m+1} I_{l+1}^{m+1}\right) \\ &\quad + \frac{\partial}{\partial z}\left(A_{l-1}^m I_{l-1}^m + B_{l+1}^m I_{l+1}^m\right) + \sigma_t I_l^m = \delta_{l0}\delta_{m0}\left(\sigma_a \frac{ac}{\sqrt{4\pi}} T^4 + \sigma_s I_0^0\right) \end{aligned} \tag{3}$$

for $l \in [0,N]$ and $m \in [-l,l]$. The constants in Eq. (3) are given by

$$\begin{aligned} A_l^m &= \sqrt{\frac{(l-m+1)(l+m+1)}{(2l+3)(2l+1)}}, \quad B_l^m = \sqrt{\frac{(l-m)(l+m)}{(2l+1)(2l-1)}}, \\ C_l^m &= \sqrt{\frac{(l+m+1)(l+m+2)}{(2l+3)(2l+1)}}, \quad D_l^m = \sqrt{\frac{(l-m)(l-m-1)}{(2l+1)(2l-1)}}, \\ E_l^m &= \sqrt{\frac{(l-m+1)(l-m+2)}{(2l+3)(2l+1)}}, \quad F_l^m = \sqrt{\frac{(l+m)(l+m-1)}{(2l+1)(2l-1)}}. \end{aligned} \tag{4}$$

The material energy equation is only coupled to the I_0^0 moment:

$$\frac{\partial}{\partial t} e(T) = \sigma_a\left(\sqrt{4\pi} I_0^0 - acT^4\right). \tag{5}$$

Though the P_N equations are written in complex form, an equivalent real-valued system can be derived by rewriting the exponential in the spherical harmonics expansion in term of sines and cosines or by adding and subtracting I_l^m and I_l^{-m} to define new moments. In this work, we only solve the equations in two-dimensional Cartesian geometries in the x–z plane. The assumption of no y-dependence implies that the real and imaginary components in Eq. (3) are decoupled.

The P_N equations are formally exact in the limit $N \to \infty$; nevertheless, we must use a finite value of N to do numerical computations. Thus the system must be closed since the equation for I_l^m includes moments of order $l+1$. A common approach is truncation: simply set $I_{N+1}^m = 0$. This closure is known as the "triangular truncation" in the spectral methods literature [29] and simply "the P_N closure" in the particle transport literature [30]. The P_N closure is motivated by the fact that as the number of collisions between the photons and the material increases, the higher moments of I become small because the intensity relaxes to an isotropic distribution.

Unfortunately, the truncation closure allows large oscillations in the reconstruction. One way to understand this fact is to observe that the expansion in Eq. (2) can be derived by finding coefficients b_l^m (which depend on $\mathbf{x}$) that minimize the cost functional [29]

$$\mathcal{J}(b_l^m) = \int_{\mathbb{S}^2} d\Omega \left(I - I_N(b_l^m)\right)^2, \tag{6}$$

where

$$I_N(b_l^m) = \sum_{l=0}^{N} \sum_{m=-l}^{l} Y_l^m(\mu,\varphi) b_l^m.$$

The solution to this minimization problem is $b_l^m = I_l^m$; that is, given N, $I_N(I_l^m)$ is the best approximation of I in the L^2 norm. Yet, the functional $\mathcal{J}$ does not account for gradients with respect to Ω. As a result, the expansion can be highly oscillatory. In regions of free-streaming transport, the P_N equations will advect these oscillations through space.

We exhibit the oscillatory nature of the P_N reconstruction in two important examples. In Fig. 1 we plot the P_N reconstruction of a delta function at $\mu = \frac{1}{2}$ and a Heaviside (step) function with jump at $\mu = \frac{1}{2}$. Both of these distributions arise in real problems: the delta function represents a beam and the Heaviside function corresponds to the edge of a shadow. Oscillations in the P_N expansions are evident in both cases and become worse as the order of the expansion increases. Moreover, in the reconstruction of the beam, the P_{29} reconstruction has the largest negative values. In the next section, we add a term to the functional $\mathcal{J}$ that penalizes large derivatives in angle to derive a *filtered* spherical harmonics expansion.

3. Filtered spherical harmonic expansions

We now derive an expansion of the intensity that is based on spherical harmonics, but does not allow large oscillations in the solution. The modified expansion is termed a *filtered spherical harmonics expansion* in that each term of the original spherical harmonics expansion is multiplied by an order-dependent coefficient and thus acts as a high-frequency filter. Following

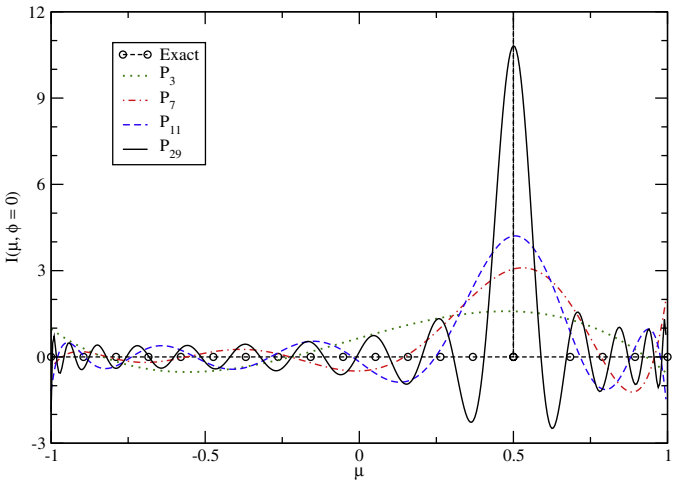

(a) Beam: $I(\mu^c \varphi) = \delta(\mu - \frac{1}{2})$.

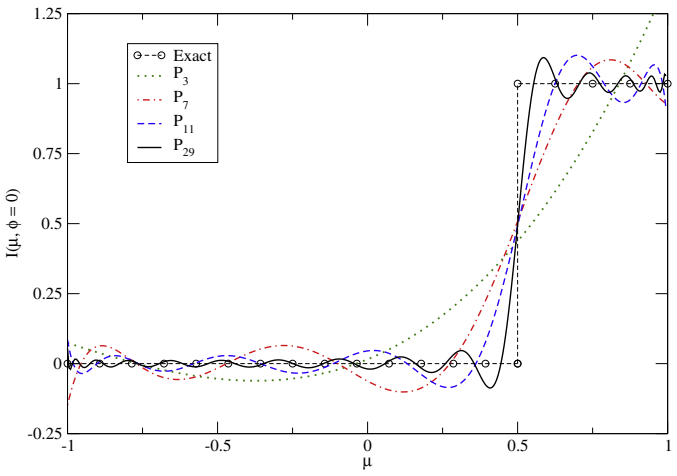

(b) Shadow: $I(\mu^c \varphi) = H(\mu - \frac{1}{2})$, where H is the Heaviside function.

Fig. 1. The P_N reconstruction of the intensity for several values of N, plotted as a function of μ, with $\varphi = 0$ fixed.

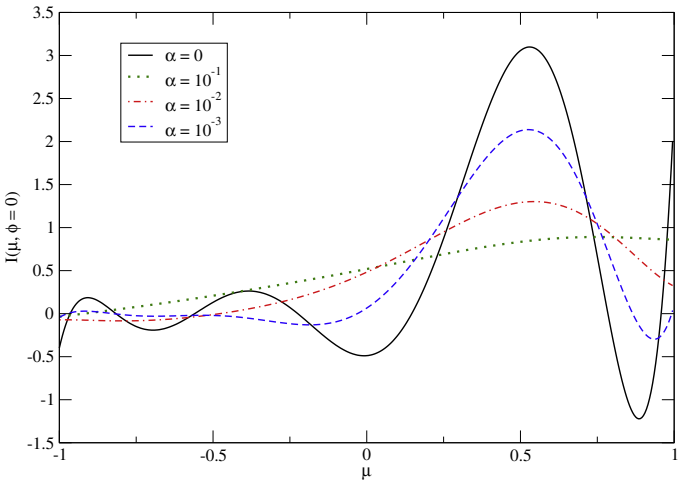

(a) FP_7 reconstructions of the beam

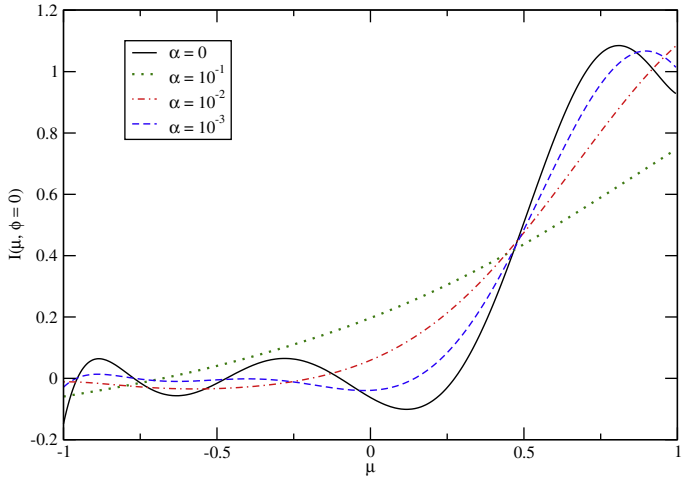

(b) FP_7 reconstructions of the shadow

Fig. 2. The filtered P_N reconstructions for several values of α, plotted as a function of μ, with $\varphi = 0$ fixed.

the colloquial terminology, we will also refer this expansion as the *filtered* P_N or FP_N expansion and the resulting equations as the *filtered* P_N or FP_N equations.

The specific type of filter we apply is the spherical spline expansion discussed in Boyd's monograph [29]. To derive this filtered expansion, we add a penalty term to the functional $\mathcal{J}$ from the previous section to generate the new cost functional

$$\mathcal{J}_\alpha = \int_{\mathbb{S}^2} d\Omega (I - I_N)^2 + \alpha \int_{\mathbb{S}^2} d\Omega \left(\nabla_\Omega^2 I_N\right)^2. \tag{7}$$

This new functional penalizes the second derivatives of the reconstruction.[2] The parameter $\alpha > 0$ tunes the strength of this penalty, and obviously $\alpha = 0$ recovers the original P_N expansion.

The minimizer of $\mathcal{J}_\alpha$ can be found in a simple closed form. We recall that the spherical harmonic functions are eigenfunctions of the spherical Laplacian:

$$\nabla_\Omega^2 Y_l^m = -l(l+1) Y_l^m.$$

Using this relation, it can be shown [29] that the minimizer of $\mathcal{J}_\alpha$ is $I_N\left(\widehat{I}_l^m\right)$, where

$$\widehat{I}_l^m := \frac{I_l^m}{1 + \alpha l^2 (l+1)^2}. \tag{8}$$

[2] In a more general setting, the penalty terms may include arbitrary even derivatives.

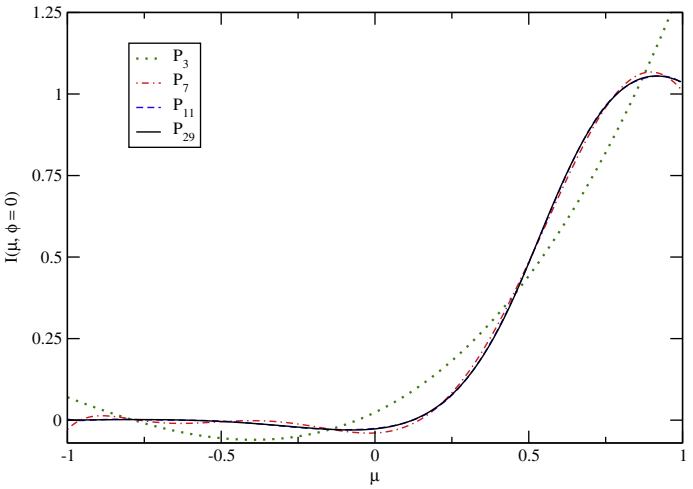

Fig. 3. Filtered P_N reconstructions at $\alpha = 10^{-3}$ fixed. Different expansion orders are plotted as a function of μ at $\varphi = 0$. Note the P_{11} and P_{29} curves coincide.

Therefore, our filtered expansion is simply

$$I \approx \sum_{l=0}^{N} \sum_{m=-l}^{l} \frac{I_l^m Y_l^m(\mu,\varphi)}{1+\alpha l^2(l+1)^2}. \tag{9}$$

The filtered expansion preserves many of the fundamental properties of the P_N expansion. First and foremost, the P_N expansion is rotationally invariant.[3] This property is not changed by the filtered expansion, because the filter affects all m values for a given l level in the same way. Second, linearity of the P_N equations is preserved by an appropriate choice of α. Finally, because $\alpha > 0$, the equations are hyperbolic – that is, all the eigenvalues of the filtered P_N equations are real, as we will discuss later.

The effects of the filter strength α are shown in Fig. 2. Here we take the filtered P_7 reconstruction of the beam and the shadow from Fig. 1. For the largest value of α shown, $\alpha = 0.1$, the filtered expansion grossly smears out the original intensity. As α is decreased, the filtered expansion reduces to the original P_N expansion. In the reconstruction of the shadow, Fig. 2(b), we can readily see the effect of the filter. As the value of α is increased from zero, the step at $\mu = \frac{1}{2}$ is smoothed out and the oscillations are damped considerably. The filtered expansions can still produce negative values, but the magnitude of these values is greatly reduced compared with the standard expansion. In cases where a guaranteed positive solution is essential, other methods are available. Specifically, the positive-P_N closure in Ref. [25] uses a constrained quadratic optimization problem to ensure a positive reconstruction.

In Fig. 3 the effect of the filter is shown when α is fixed, but the expansion order increases. Here we see that for a given value of α, the accuracy of the reconstruction stagnates: the P_7 solution is much better (and less negative) than the P_3 reconstruction, but beyond P_7, the benefits of going to higher-order diminish significantly. This highlights the fact that the filter strength must decrease as the expansion order increases in order to see the benefits of higher-order reconstructions.

4. Choosing the filter strength

So far, we have shown how the filtered expansion can improve the reconstruction of the radiation intensity by removing oscillations. In this section, we discuss our approach to selecting α. Our method is, by no means, the only way to choose α; but it does provide a recipe that for a fixed mesh spacing, preserves the equilibrium diffusion limit, maintains rotational invariance, improves robustness, and formally converges to the transport solution.

Before presenting a formula for α, we first delineate the properties we would like this formula to have.

- *Vanishing filter strength as $N \to \infty$.* To ensure convergence to the transport solution (at least formally), the value of α should go to zero as the order of the expansion increases. This property is motivated by the results we saw in Fig. 3: if α remains constant as N increases, the quality of the reconstruction stagnates.
- *The equilibrium diffusion limit.* The radiative transfer/material energy system in Eqs. 1a and 1b has the important limit that as σ_a becomes large, the radiation and blackbody source come into equilibrium and the two equations become a single

[3] In abstract form, the radiative transfer equation can be written $\mathcal{T} I = 0$. It is straightforward to show that the operator $\mathcal{T}$ commutes with the rotation operator $\mathcal{R} : f(\mathbf{x}, \Omega, t) \mapsto f(O^T\mathbf{x}, O^T\Omega, t)$, for any orthogonal matrix O. The P_N approximation maintains this property. In particular, if $\mathcal{P}$ is the projection onto the expansion defined in (2), then the operator $\mathcal{P} \circ \mathcal{T}$ also commutes with $\mathcal{R}$.

nonlinear diffusion equation for the material temperature [31]. To derive this equilibrium diffusion limit, one introduces a small parameter $\epsilon > 0$ and rescales the quantities in (1) as follows:

$$\sigma_a \to \frac{\sigma_a}{\epsilon}, \quad \sigma_s \to \epsilon\sigma_s, \quad c \to \frac{c}{\epsilon}, \quad C_v \to \epsilon C_v. \tag{10}$$

Following the standard asymptotic analysis in Refs. [31,32], it can easily be shown that if $\alpha = O(\epsilon^2)$, then the asymptotic diffusion limit of the radiative transfer/material energy system is preserved by the FP_N system through $O(\epsilon)$. Therefore, α should be dependent on σ_t in a certain way.

- *Rotational invariance.* The choice of α should preserve the rotational invariance of the transport equation. This will be the case if α does not depend on Ω. In a preliminary study [26], we presented a means for choosing α based on keeping the P_1 expansion positive. However, the way in which positivity was enforced made α dependent on the solution in particular directions. The loss of rotational invariance was obvious in our results, and we have since abandoned this approach.
- *Small α.* We desire α to be as small as possible – that is, just large enough to reduce negative solutions and oscillations in the solution to acceptable levels. For instance, a problem driven largely by free-streaming would potentially need a larger value of α to eliminate negative solutions, compared to a problem where scattering was the dominant process.

Taking all of these items into account, we propose to calculate α as

$$\alpha = \frac{\omega}{N^2}\frac{1}{(\sigma_t L + N)^2}. \tag{11}$$

Here, L is characteristic length scale used to make α non-dimensional, and ω is a user defined constant. This form for α has the characteristics that we laid out above. As $N \to \infty$ the magnitude of α goes to zero, and when σ_t is an $O(1/\epsilon)$ quantity, $\alpha = O(\epsilon^2)$. Also, α is not dependent on the solution in any particular direction. It only depends on the material properties in the problem and the order of the expansion.

The presence of the problem dependent parameter ω, as well as the power to which N is raised, allows the strength of the filter to be adjusted for a given problem, while maintaining the desired properties of α. We understand that introducing such a "tunable" parameter, is, on the surface, distasteful. Nonetheless, we also assert that the form we have given for α controls the impact of the tunability: no matter what ω is chosen, α will still go to zero as N increases (as long as N is raised to a positive power). The value of ω behaves in a similar manner to the implicitness parameter in implicit Monte Carlo simulations [11] in the sense that it does not affect the limiting properties of the solution. Furthermore, the numerical experiments shown below demonstrate that one can use a constant value of ω across different problems.

Though in this work we solve the filtered P_N equations with a semi-implicit scheme where we can evaluate the filter terms explicitly, an implicit time integration scheme could have issues treating this term because of a potential temperature dependence of σ_t. We believe that in such an instance treating the filter strength using a lagged value of σ_t would be sufficient to ensure the radiation intensity is filtered where it should be without adding additional computational cost to the implicit solve. In principle, one could treat the filter strength fully implicitly and for large time step sizes this might be necessary.

We also would like to point out that the strength of α could be made time-dependent. In many problems the large oscillations in the P_N solutions appear in the transient regimes, soon after a source is turned on. In such cases, ω could be made a function of time so that the filter is initially strong and then subsides at later times as the system relaxes. We have not pursued such time-dependent filters, but we believe they could be useful.

5. Implementation and modified equations

The filtering procedure can be used in combination with any existing P_N code with minor alterations. For this work we use a semi-implicit, linear, discontinuous Galerkin method [32,33]. The method employs a double minmod slope limiter to eliminate the artificial oscillations that can arise from our spatial scheme while still preserving the equilibrium diffusion limit [34].

Our implementation is based on a set modified equations that are derived from the filtering procedure. For compactness, we first rewrite the original P_N system (3) in the following abstract form

$$\frac{1}{c}\frac{\partial \mathbf{u}}{\partial t} + \nabla_{\mathbf{x}} \cdot (\mathbf{J}\mathbf{u}) + \sigma_t \mathbf{u} = \mathbf{L}(\sigma_a \mathbf{q} + \sigma_s \mathbf{u}), \tag{12a}$$

$$\frac{\partial e}{\partial t} = -\mathbf{e}_0^T(\sigma_a \mathbf{q} - \sigma_a \mathbf{u}). \tag{12b}$$

Here the vector $\mathbf{u}$ contains the components of I_l^m; we assume its first component is I_0^0. The matrix $\mathbf{J}$ is a constant matrix whose elements are components of the matrices A, B, C, D, E, and F that are defined in (4). The matrix $\mathbf{L}$ corresponds to the Kronecker delta pair in (3). Thus all of its entries are zero, except for a one in the first diagonal entry that multiplies I_0^0. The vector $\mathbf{e}_0$ is a unit vector whose first component is one and

$$\mathbf{q} = \frac{ac}{4\pi}T^4\mathbf{e}_0. \tag{13}$$

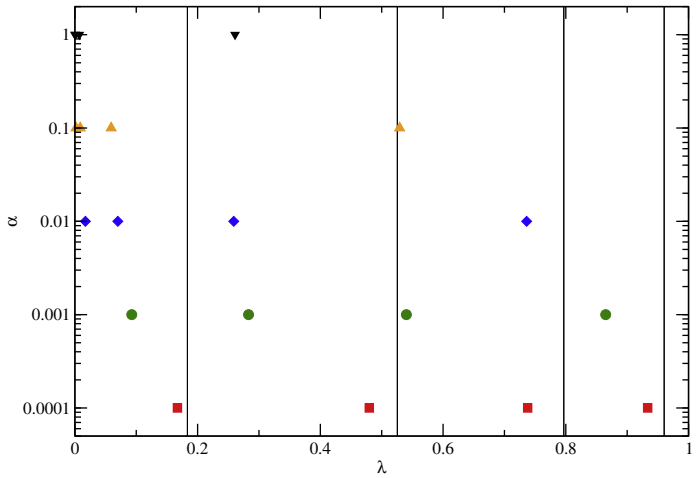

Fig. 4. Positive eigenvalues of the filtered P_7 equations for several values of α. The solid lines are the eigenvalues for the standard P_7 equations.

For the purposes of the derivation that follows, we ignore the exact form of the spatial discretization, including any stability issues with DG. We also ignore the temperature equation, which is unaffected by the filter. A first-order, semi-implicit, temporal discretization of the original P_N system takes the form

$$\frac{\mathbf{u}^{n+1}-\mathbf{u}^n}{c\Delta t}+\nabla_{\mathbf{x}}\cdot(\mathbf{J}\mathbf{u}^n)+\sigma_t\mathbf{u}^{n+1}=\mathbf{L}\left(\sigma_a\mathbf{q}^{n+1}+\sigma_s\mathbf{u}^{n+1}\right). \tag{14}$$

We apply the filter to the current value of $\mathbf{u}$, via a diagonal matrix $\mathbf{F}$. This gives

$$\frac{\mathbf{u}^{n+1}-\mathbf{F}\mathbf{u}^n}{c\Delta t}+\nabla_{\mathbf{x}}\cdot(\mathbf{J}\mathbf{F}\mathbf{u}^n)+\sigma_t\mathbf{u}^{n+1}=\mathbf{L}\left(\sigma_a\mathbf{q}^{n+1}+\sigma_s\mathbf{u}^{n+1}\right) \tag{15}$$

which we rewrite as

$$\frac{\mathbf{u}^{n+1}-\mathbf{u}^n}{c\Delta t}+\nabla_{\mathbf{x}}\cdot(\mathbf{J}\mathbf{F}\mathbf{u}^{n+1})+\sigma_t\mathbf{u}^{n+1}+\mathbf{S}\mathbf{u}^n=\mathbf{L}\left(\sigma_a\mathbf{q}^{n+1}+\sigma_s\mathbf{u}^{n+1}\right). \tag{16}$$

where the diagonal matrix $\mathbf{S}$ is given by

$$\mathbf{S}:=\frac{\mathbf{Id}-\mathbf{F}}{c\Delta t}. \tag{17}$$

Thus the filtering process has effectively introduced a new source and modified the fluxes by a spatially dependent factor. We note that the modified source term is similar in spirit to the recent work in Ref. [15]. However, in that work, the additional source term is isotropic, whereas the source term in Eq. (16) dampens higher-order moments more strongly, and therefore, acts as a forward–peaked scattering term. We call this new system the *filtered* P_N (or FP_N) equations. Several remarks about the FP_N system are in order.

- The filter matrix $\mathbf{F}$ introduces an explicit spatial dependence into the flux of the FP_N system, via the cross-section in the definition of α in (11). From a physical point of view, it is desirable to have information about the material medium embedded into the closure. Be that as it may, a complete mathematical understanding of this system is more difficult than for the original P_N system. In particular, discontinuities in σ_t will lead to additional discontinuities in the FP_N solution, so that interface conditions must be applied when solving for fluxes between computational cells. In the DG method, it is assumed that the flux crossing an interface is continuous at that interface. For the original P_N system, which is linear, this implies that characteristic waves crossing an interface are also continuous at that interface. This property is consistent with the common assumption that the radiation intensity in directions crossing an interface is continuous across the interface. However, the explicit spatial dependence of the FP_N flux means that the characteristic waves will be discontinuous at material interfaces, unless the DG method is modified. The implications of this fact and whether such a modification is appropriate should be explored in more detail.
- Unlike the original P_N system, there is no simple relation for the eigenvalues of the FP_N equations.[4] However, the effect of α on the eigenvalues for the case $N = 7$ is shown in Fig. 4. As α is increased, the eigenvalues are reduced and begin to cluster

[4] The eigenvalues of the free-streaming P_N equations are the values for μ for which $Y^m_{N+1}=0$ for $m\leqslant N+1$ [28]. This can be easily derived for the eigenvalues in the z-direction and extended to the x and y directions by invoking the rotational invariance of the spherical harmonics expansion.

near zero. Thus, in additional to damping the oscillations in the P_N reconstruction, the filter slows down the speed at which information is propagated.

- The source term in (16) is a diagonal matrix. For each fixed l, its effect is to multiply components I_l^m, $|m| \leqslant l$, by the factor

$$s_l = \frac{1}{c\Delta t}\frac{\alpha l^{2k}(l+1)^{2k}}{1+\alpha l^{2k}(l+1)^{2k}}. \tag{18}$$

For the form of α given in (11) and with $\omega = \frac{c\Delta t}{\Delta x}$,

$$s_l = \frac{cl^2(l+1)^2}{\mathrm{d}xN^2(\sigma tL+N)^2 + c\mathrm{d}tl^2(l+1)^2} \tag{19}$$

We note that these terms blow up as $\Delta x, \Delta t \to 0$ together. This is a direct result of our choice of ω: had we not made that parameter $c\Delta t/\Delta x$, the value of s_l would be singular if $\Delta t \to 0$ independently. While such singular behavior is unwanted, it should also be noted that we do not wish to recover the original P_N system in the continuum limit either. Indeed, at this point we are unsure of what the proper continuum limit should be. Even so, we have found the filter gives excellent results for a relatively wide range of mesh parameters. Further investigation of this issue is ongoing.

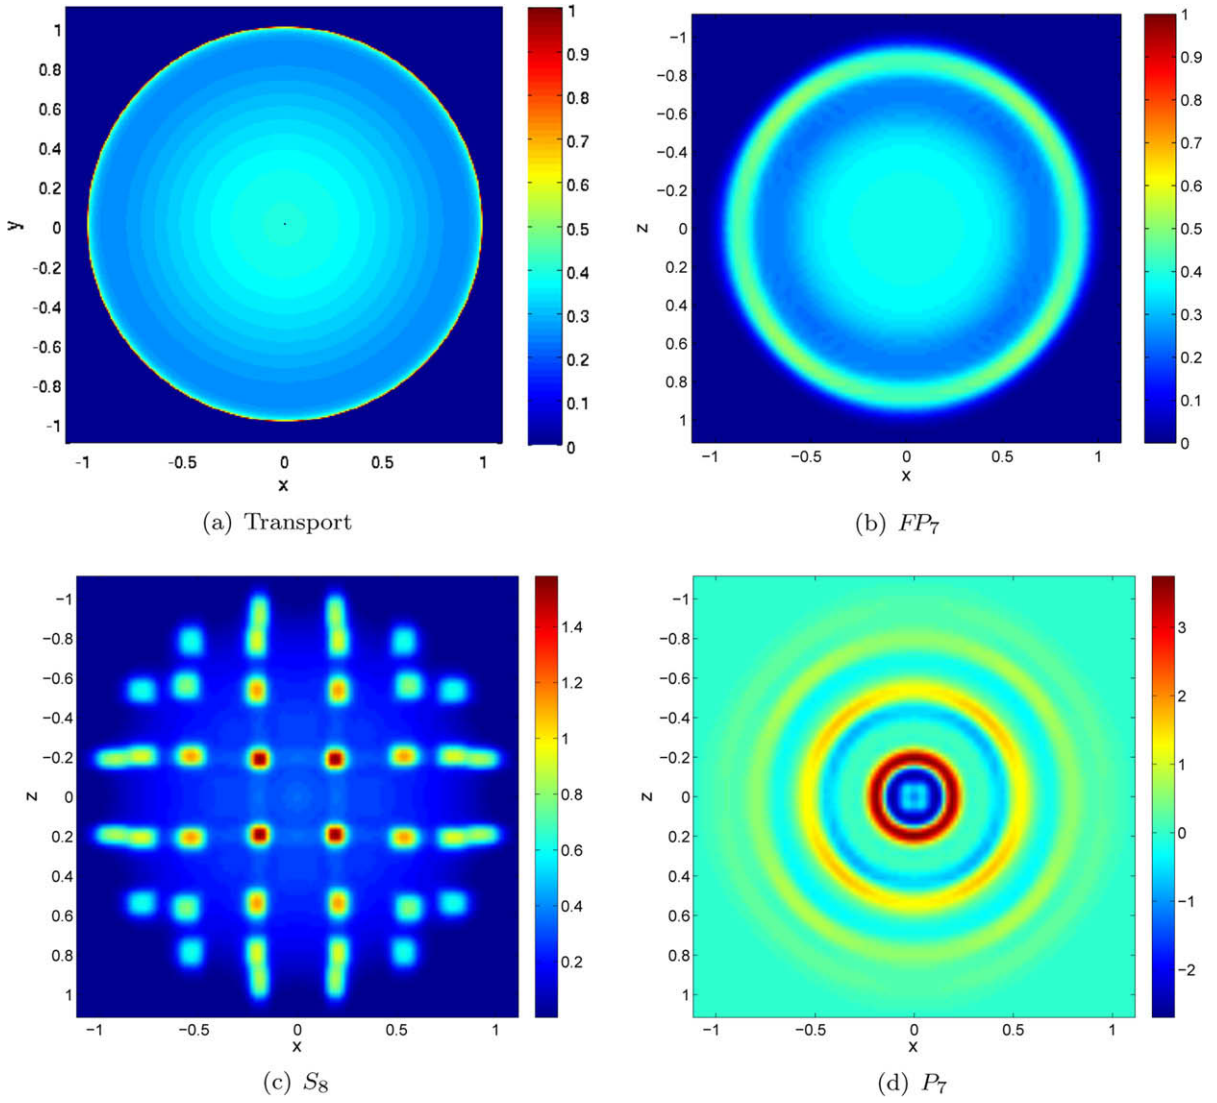

Fig. 5. Scalar density ϕ for the line-source problem when $ct = 1$. Note the changes in color scale to accommodate the different ranges of the S_8 and P_7solutions.

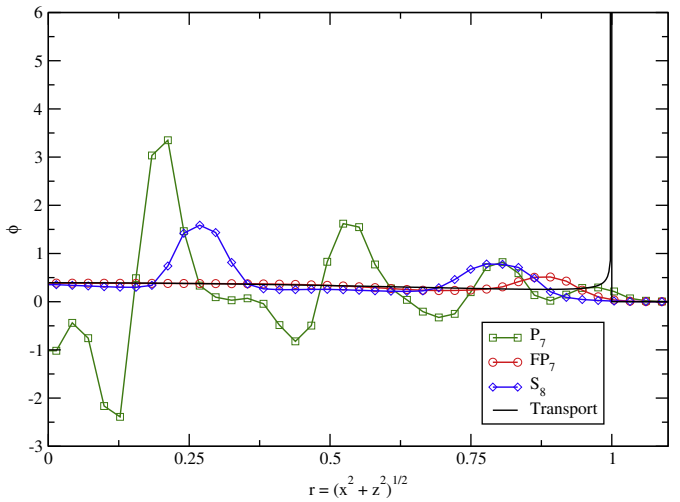

Fig. 6. Scalar density ϕ for the line-source problem along the diagonal line $x = z$ when $ct = 1$.

We now apply a predictor–corrector time integrator to the FP_N system, treating the new source term **Su** *explicitly*. Including the temperature equation, we have the following[5]:

$$\frac{\mathbf{u}^{n+1/2}-\mathbf{u}^{n}}{c\Delta t/2}+\nabla_{\mathbf{x}}\cdot(\mathbf{J}\mathbf{F}\mathbf{u}^{n})+\mathbf{S}\mathbf{u}^{n}+\sigma_{\mathrm{t}}\mathbf{u}^{n+1/2}=\mathbf{L}(\sigma_{a}\mathbf{q}^{n+1/2}+\sigma_{\mathrm{s}}\mathbf{u}^{n+1/2}), \tag{20a}$$

$$\frac{e^{n+1/2}-e^{n}}{\Delta t/2}=-\sigma_{a}\mathbf{e}_{0}^{T}(\mathbf{q}^{n+1/2}-\mathbf{u}^{n+1/2}), \tag{20b}$$

$$\frac{\mathbf{u}^{n+1}-\mathbf{u}^{n}}{c\Delta t}+\nabla_{\mathbf{x}}\cdot(\mathbf{J}\mathbf{F}\mathbf{u}^{n+1/2})+\mathbf{S}\mathbf{u}^{n+1/2}+\sigma_{\mathrm{t}}\mathbf{u}^{n+1}=\mathbf{L}(\sigma_{a}\mathbf{q}^{n+1}+\sigma_{\mathrm{s}}\mathbf{u}^{n+1}), \tag{20c}$$

$$\frac{e^{n+1}-e^{n}}{\Delta t}=-\sigma_{a}\mathbf{e}_{0}^{T}(\mathbf{q}^{n+1}-\mathbf{u}^{n+1}). \tag{20d}$$

In the streaming limit ($\sigma_{\mathrm{t}}\to 0$), this is a second order approximation of the modified Eq. (16) and the temperature equation. It can be understood in terms of the original filter **F** through a rearrangement of terms in the time differencing:

$$\frac{\mathbf{u}^{n+1/2}-\frac{1}{2}(\mathbf{F}+\mathbf{Id})\mathbf{u}^{n}}{c\Delta t/2}+\nabla_{\mathbf{x}}\cdot(\mathbf{J}\mathbf{F}\mathbf{u}^{n})+\sigma_{\mathrm{t}}\mathbf{u}^{n+1/2}=\mathbf{L}(\sigma_{a}\mathbf{q}^{n+1/2}+\sigma_{\mathrm{s}}\mathbf{u}^{n+1/2}), \tag{21a}$$

$$\frac{e^{n+1/2}-e^{n}}{\Delta t/2}=-\sigma_{a}\mathbf{e}_{0}^{T}(\mathbf{q}^{n+1/2}-\mathbf{u}^{n+1/2}), \tag{21b}$$

$$\frac{(\mathbf{u}^{n+1}-\mathbf{F}\mathbf{u}^{n+1/2})+(\mathbf{u}^{n+1/2}-\mathbf{u}^{n})}{c\Delta t}+\nabla_{\mathbf{x}}\cdot(\mathbf{J}\mathbf{F}\mathbf{u}^{n+1/2})+\sigma_{\mathrm{t}}\mathbf{u}^{n+1}=\mathbf{L}(\sigma_{a}\mathbf{q}^{n+1}+\sigma_{\mathrm{s}}\mathbf{u}^{n+1}), \tag{21c}$$

$$\frac{e^{n+1}-e^{n}}{\Delta t}=-\sigma_{a}\mathbf{e}_{0}^{T}(\mathbf{q}^{n+1}-\mathbf{u}^{n+1}). \tag{21d}$$

This form makes it clear how the algorithm combines filter and non-filtered components of **u**.

6. Numerical results

Our numerical results demonstrate that the filtered P_N expansions do indeed enhance the robustness and quality of spherical harmonics solutions. In all of the results that follow, we use $\omega=\frac{c\Delta t}{\Delta x}$ and $L = 1$ cm. The problems we solve are 2D Cartesian problems, and to solve P_N equations in this geometry we set $\frac{\partial l}{\partial y}=0$.

6.1. The line-source problem

The first problem we solve is the so-called "line-source problem". Here we have an initial condition of

$$I(x,z,\Omega,t)=\frac{1}{4\pi}\delta(x)\delta(z),$$

[5] It should be noted that the predictor value $e^{n+1/2}$ is not needed in the corrector step and therefore, not explicitly computed.

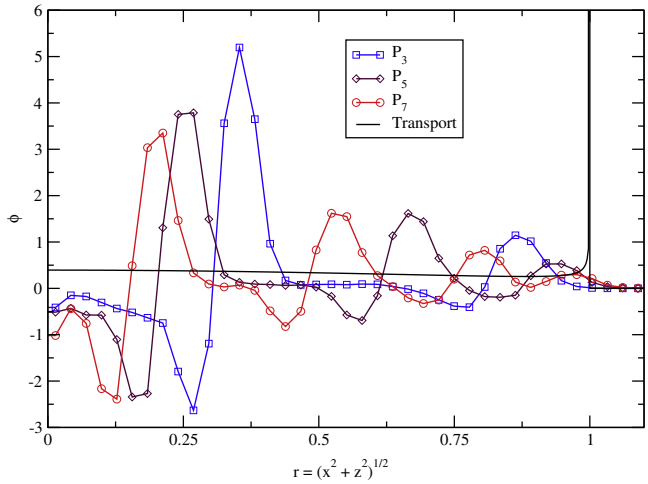

(a) P_N

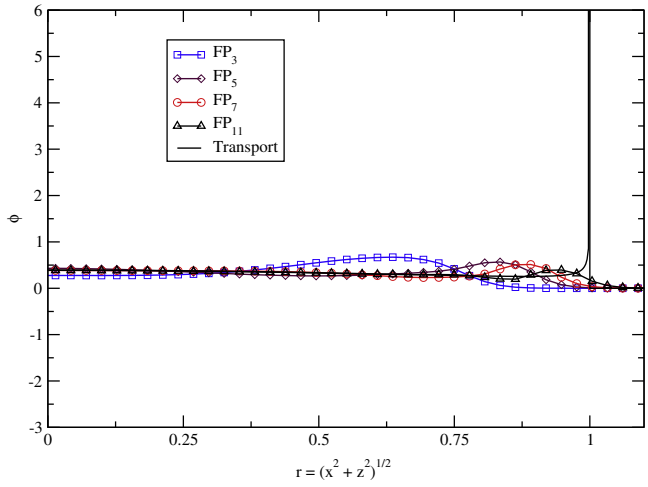

(b) FP_N

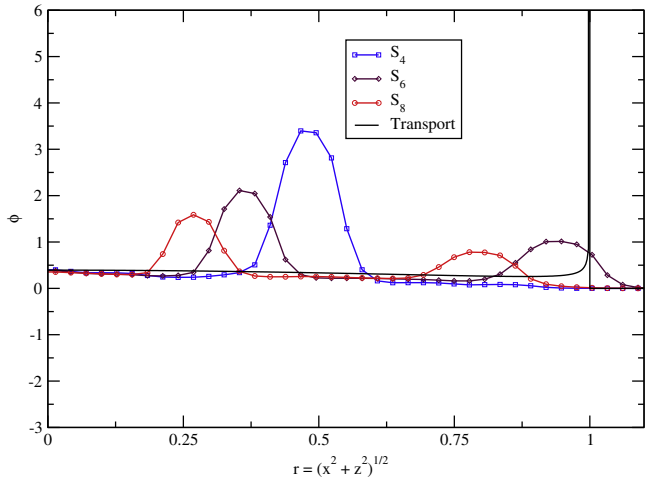

(c) 19 S_N

Fig. 7. Scalar density ϕ for the line-source problem along the diagonal line $x = z$ when $ct = 1$.

and $\sigma_t = \sigma_s = 1.0\ \mathrm{cm}^{-1}$. We display the solution to this problem in terms of the scalar density:

$$\phi := \langle I \rangle \equiv \sqrt{4\pi} I_0^0$$

which, from a physical perspective, is always non-negative. There is a semi-analytic solution to the radiative transfer equation for this problem given by Ganapol [35]. This solution, which we refer to as the transport solution, has a delta function of uncollided photons supported on the circle $\sqrt{x^2+z^2}=ct$ and a smooth region of collided photons in the disk $\sqrt{x^2+z^2}<ct$. There is no radiation beyond the delta function, because photons cannot travel faster than the speed of light. Though this problem has cylindrical symmetry, we solve it in a 2D Cartesian domain because we are ultimately interested in the behavior of transport methods in 2- and 3D.

This is a particularly vexing problem for spherical harmonics methods. Indeed the P_1 solution to this problem has a *negative* singularity in the solution [36]. The uncollided particles in this problem have a very forward–peaked, beam-like distribution, and as our analysis above showed (c.f. Fig. 1), such an intensity will lead to a negative P_N reconstruction. Previous works have used this problem as test bed for schemes that have been devised to improve the robustness of the P_N equations and for comparisons between methods [37,25,21,26]. See Ref. [15] for a discussion of negative P_N solutions to a similar problem, which therein is a called a radiation blast wave.

In Fig. 5 we show the solution for several methods at time t when $ct = 1$ – that is, one mean-free time after the pulse. The P_N and S_N solutions were obtained using $N_x = N_z = 110$ on a domain of 2.2 cm $\times$ 2.2 cm, and $\Delta t = 0.24\Delta x/c$. We used an equal-weight quadrature set in the S_N computation. The transport solution in Fig. 5(a) has the color range limited to a maximum of 1 because the magnitude of the delta function at $\sqrt{x^2+z^2}=1$ would mask the collided radiation on the full scale. We can see that behind the uncollided radiation, there is a smooth, nonzero region. Comparing the transport solution with filtered P_7 solution, we see that the solution in this collided region is comparable to the transport solution, especially near the origin. The wavefront of uncollided radiation in the filtered P_7 solution does not agree with the transport solution: it is not as sharp as the transport solution and radiation not traveled as far from the origin as it has in the transport solution. Yet, the filtered P_7 solution is the only solution in Fig. 5 that has a similar character to the transport solution: a smooth region near the origin with a spike of uncollided radiation at the outer edge. The standard P_7 solution in Fig. 5(d) is a series of rings moving out from the origin with a negative region behind the inner two rings. The S_8 solution in Fig. 5(c) is a collection of dots corresponding to the discrete directions along which the S_N equations track the flow of radiation. These dots are ray effects.

We can get a more detailed view of the solutions by looking along the diagonal line $x = z$ in the first quadrant. The solution along this line is shown in Fig. 6, where oscillations in the P_N solution are clear. There are also oscillations in the S_8 solution due to ray effects, although these oscillations do not cause the solution to go negative. We also note that picking a different line through the center of the domain will give a different shape for the S_8 solution due to its lack of rotational invariance. The filtered P_7 solution has a hump in the solution where it is trying to resolve the delta function of uncollided radiation. Behind this it agrees quite well with the transport solution.

We can also look at how the different methods behave as the angular resolution is increased. Fig. 7 shows the solution using several angular approximations for each method. The standard P_N solutions in Fig. 7(a) demonstrate no discernible convergence when going from P_3 to P_5 to P_7 other than the oscillations slightly decreasing in magnitude. This contrasts with the filtered P_N solutions, Fig. 7(b), where each successive increase in the P_N order narrows the leading hump and moves it closer to the transport wavefront. In the S_N solutions in Fig. 7(c), the oscillations also dampen as N increases. Note, however, that the lack of rotational invariance in the S_N solution plays a role. Along this particular line, the S_6 solution has moved farther from the origin than the S_8 solution; along a different line, the S_8 solution might lead the S_6 solution.

Before leaving the line-source problem, we again look at the solution along the diagonal $x = z$, but now when $ct = 10$ cm. In these numerical solutions, we used $N_x = N_z = 75$ and a domain of 15 cm $\times$ 15 cm. From Fig. 8, we see that all methods considered are in close agreement with the transport solution, except (i) at the leading edge of the transport solution, where

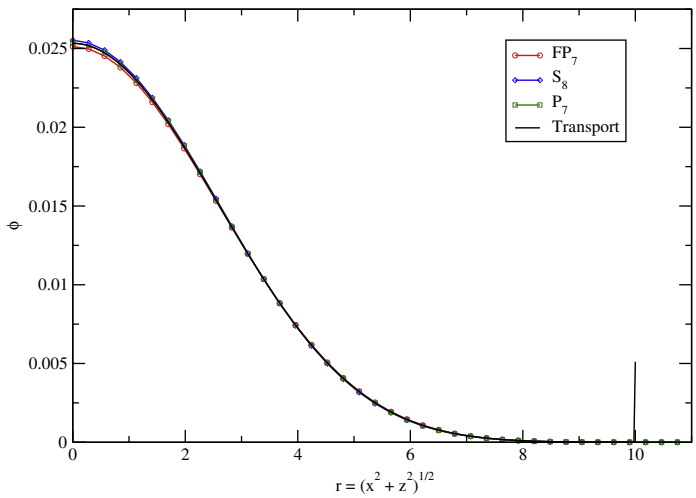

Fig. 8. Comparison of scalar density ϕ for the line-source problem along the diagonal line $x = z$ when $ct = 10$.

some photons are still uncollided and (ii) near the origin, where all methods show some small discrepancy with the transport solution.

Finally, we also note that the though the FP_N and P_N solutions to the line-source problem should be rotationally invariant, the introduction of a Cartesian grid does destroy true rotational invariance. This can be seen in Fig. 5: along the diagonals $x = z$ or $-x = z$, the solution is higher at the wavefronts than along the lines $x = 0$ or $z = 0$. We would expect that the loss of rotational invariance would be the same in terms of relative magnitude for both FP_N and P_N because this is a spatial discretization effect and both methods use the same spatial discretization. Comparing line outs of the solutions we found this to be the case.

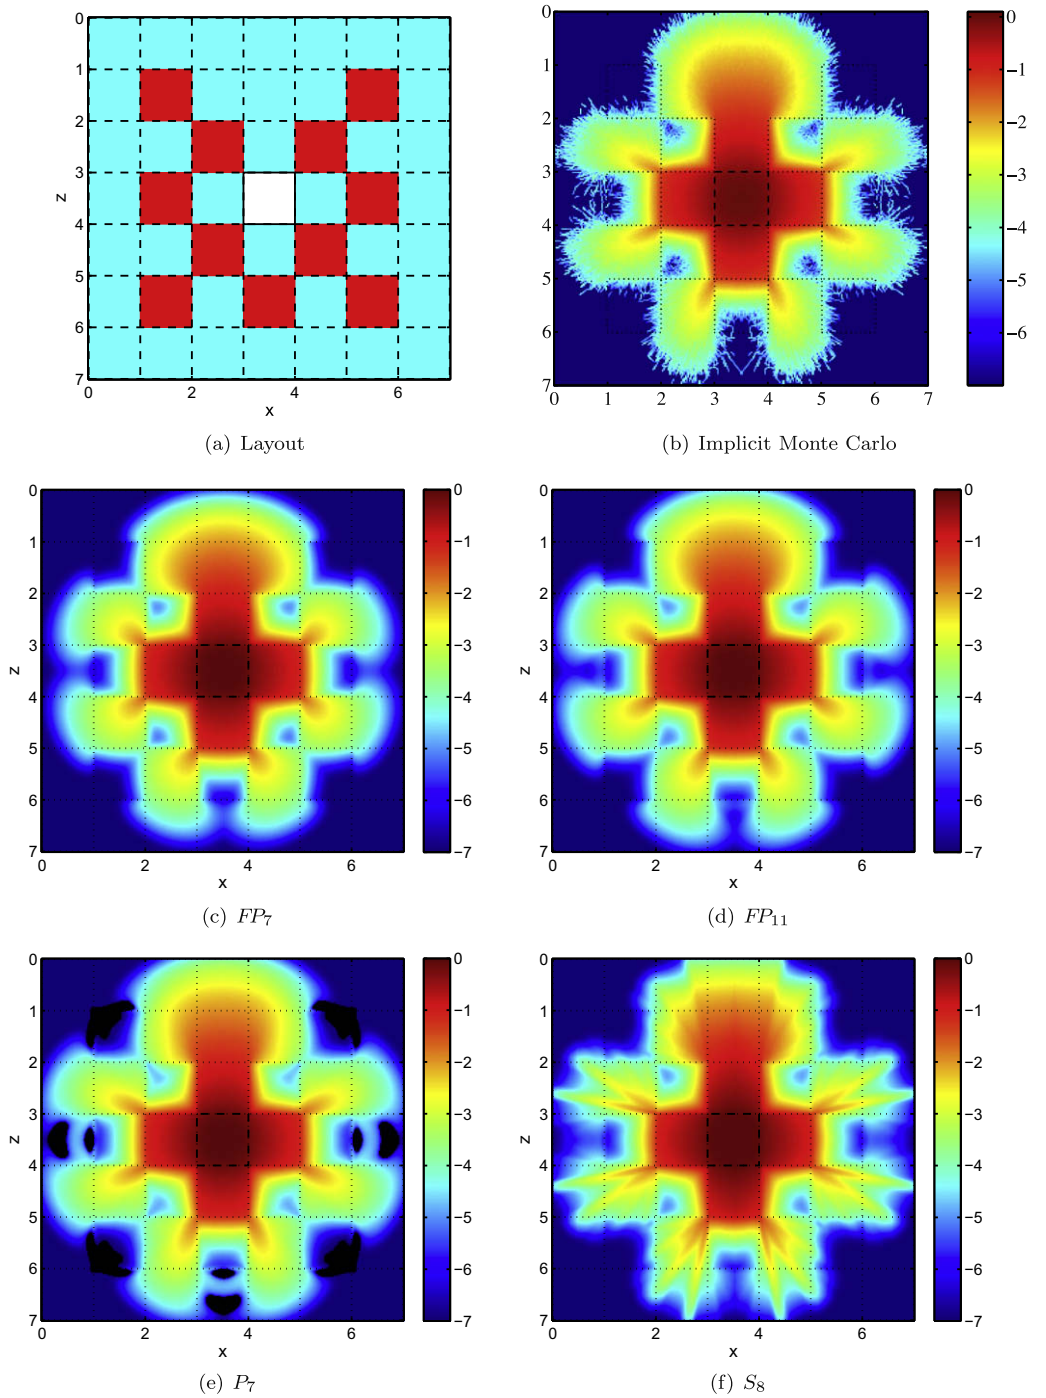

(a) Layout (b) Implicit Monte Carlo

(c) FP_7 (d) FP_{11}

(e) P_7 (f) S_8

Fig. 9. Solution for $\log\phi$ to the lattice problem at $t = 1.07$ ns.

6.2. A lattice problem

In this problem we add the complication of multiple materials: a checkerboard of strongly absorbing regions embedded into a purely scattering medium [38,37]. Here coupling to the material is neglected and only the transport of radiation is considered. (This is equivalent to having an infinitely large heat capacity and zero initial temperature.) The layout for this problem is shown in Fig. 9(a). The problem is defined on a 7×7 cm spatial domain with vacuum boundaries. The light blue regions and the white region are purely scattering with $\sigma_t = \sigma_s = 1\ \text{cm}^{-1}$; the red regions are pure absorbers with $\sigma_t = \sigma_a = 10\ \text{cm}^{-1}$; the white region also has an isotropic source $Q = 1\ \text{cm}^{-3}\ \text{s}^{-1}$. Initially, there is no radiation in the problem.

In Fig. 9(b)–(f) we compare several methods for solving this problem. The IMC solution in Fig. 9(b) was reproduced from Ref. [37]. The other solutions (FP_n, P_n, and S_n) were obtained using $\Delta x = \Delta z = 0.025$ cm. The plots show $\log\phi$ at $t = 1.07$ ns (i.e. $ct = 3.2$ cm). At this early time particles have had just enough time to reach the domain boundary. A noticeable feature of these solutions is that the beam of particles emerging from the regions where the corners of two absorbers meet. Also, the absorbers create shadows, especially near the corners of the problem.

For this problem, the lack of material coupling means that the IMC solution is exact in the limit of an infinite number of particles. Nevertheless, simulations with finite particle numbers are plagued by noise in regions where ϕ is small. The FP_{11} solution agrees well with the IMC solution in regions where noise is not a factor. There is a difference between the FP_{11} and FP_7 solutions, though the two solutions are in qualitative agreement. The standard P_7 solution has oscillations that cause the solution to go negative; regions where this occurs are colored in black. The FP_7 solution did not have these oscillations. Removing these artifacts is significant; indeed, oscillations in the P_n solution for this problem have been observed even for N as large as fifteen [38]. The discrete ordinates solution in Fig. 9(f) is dominated by ray effects. Though the S_8 simulation uses same number of degrees of freedom as the FP_7 simulation, several features of the S_8 solution are qualitatively different from IMC and FP_{11} solutions.

6.3. A hohlraum problem

The other problem we solve is a modification to the simplified hohlraum problem that was originally suggested by Brunner [37]. The problem layout is given in Fig. 10; note temperature dependence of σ_t. On the left, there is a 1 keV blackbody source along the entire left boundary emitting radiation into the initially cold material. There is no analytic solution to this problem, though we expect the solution to have the following characteristics: (i) due to the fact that radiation cannot reach the back of the block directly, the central block should not be uniformly heated by the radiation source; (ii) due to shadowing, the region that lies behind the block with respect to the source should have less radiation energy than regions within the line of sight of the source.

For this problem, we will compare the FP_7 solution to an S_8 solution, an implicit Monte Carlo (IMC) solution [11], and a flux-limited diffusion [17] solution. The IMC and flux-limited diffusion solutions effectively bracket the angular resolution extremes: IMC has high angular fidelity because it samples the entire phase space of the intensity; flux-limited diffusion, on the other hand, tracks only the scalar flux, whose evolution is approximated by a diffusion equation. It should be noted

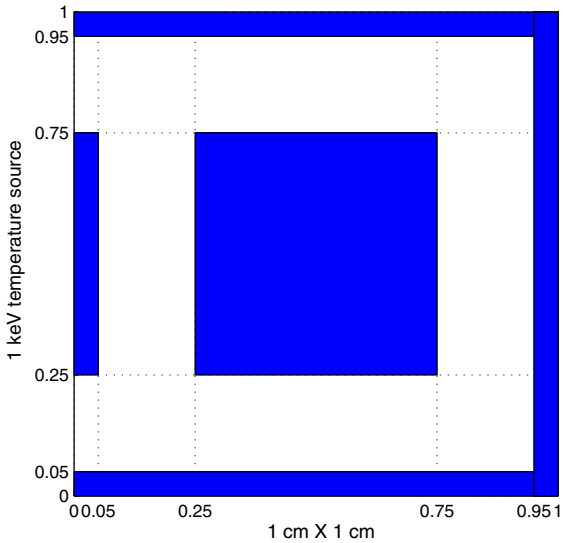

Fig. 10. Layout for the 2D hohlraum problem. The dark regions have $\sigma_a = \sigma_t = 100\,T^{-3}\ \text{cm}^{-1}$ for T in keV, $C_v = 0.3\ \text{GJ/cm}^3 - \text{keV}$; the white regions are vacuum, $\sigma_t = 0$.

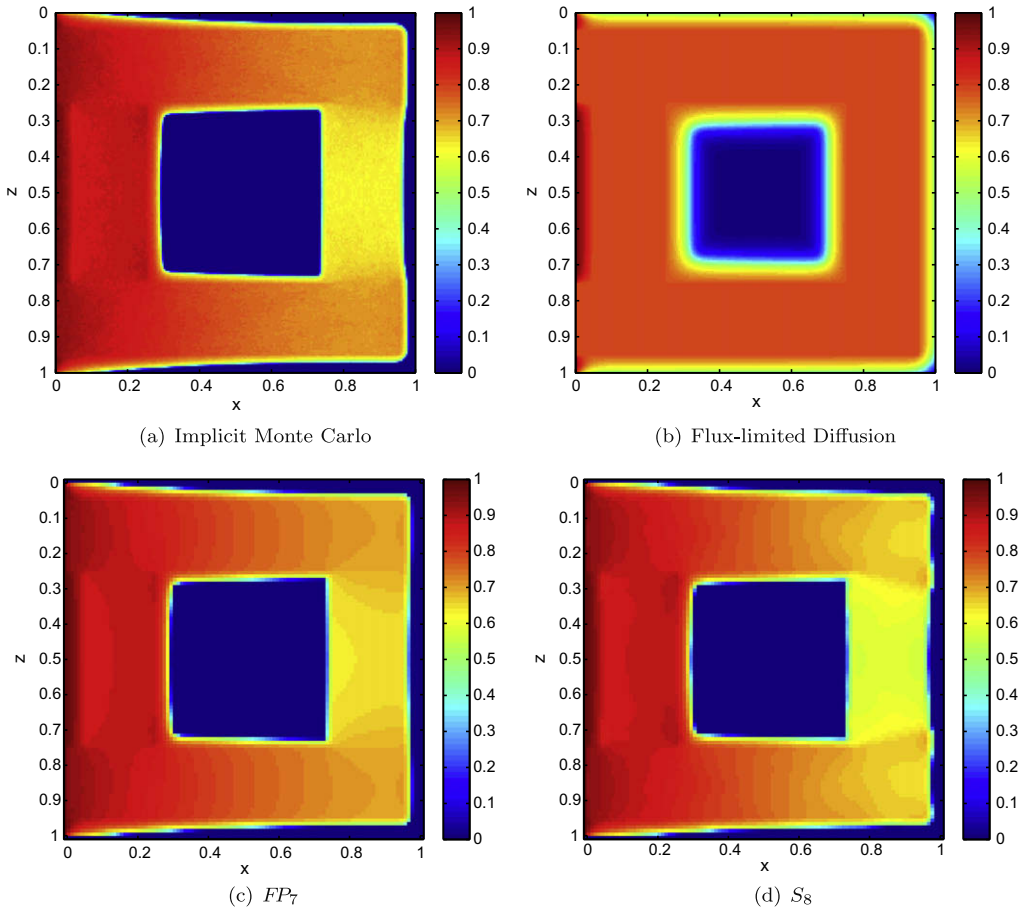

Fig. 11. Radiation temperature T_{rad} for the hohlraum problem at $t = 1$ ns.

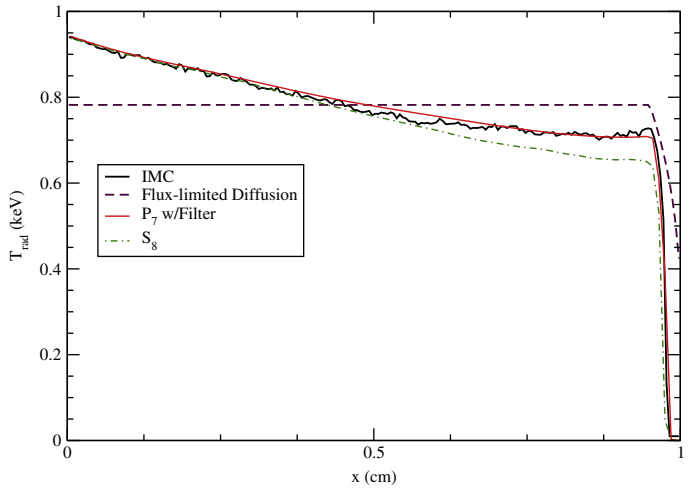

Fig. 12. Radiation temperature T_{rad} for the hohlraum problem at $t = 1$ ns along the line $z = 0.125$ cm.

that although IMC has high angular fidelity, it is not the exact solution even in the limit of an infinite number of sampled particles. This is due to the fact that IMC has spatial and temporal discretization errors that can lead to issues in diffusive regimes [39] and non-physical material heating [12,13].

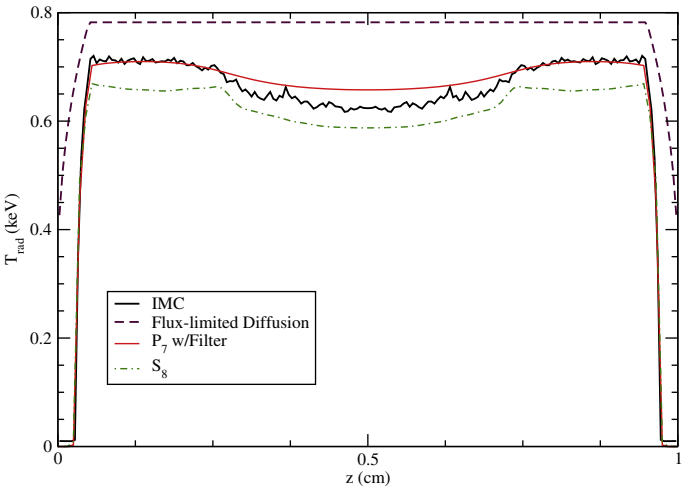

Fig. 13. Radiation temperature T_{rad} for the hohlraum problem at $t = 1$ ns along the line $x = 0.85$ cm.

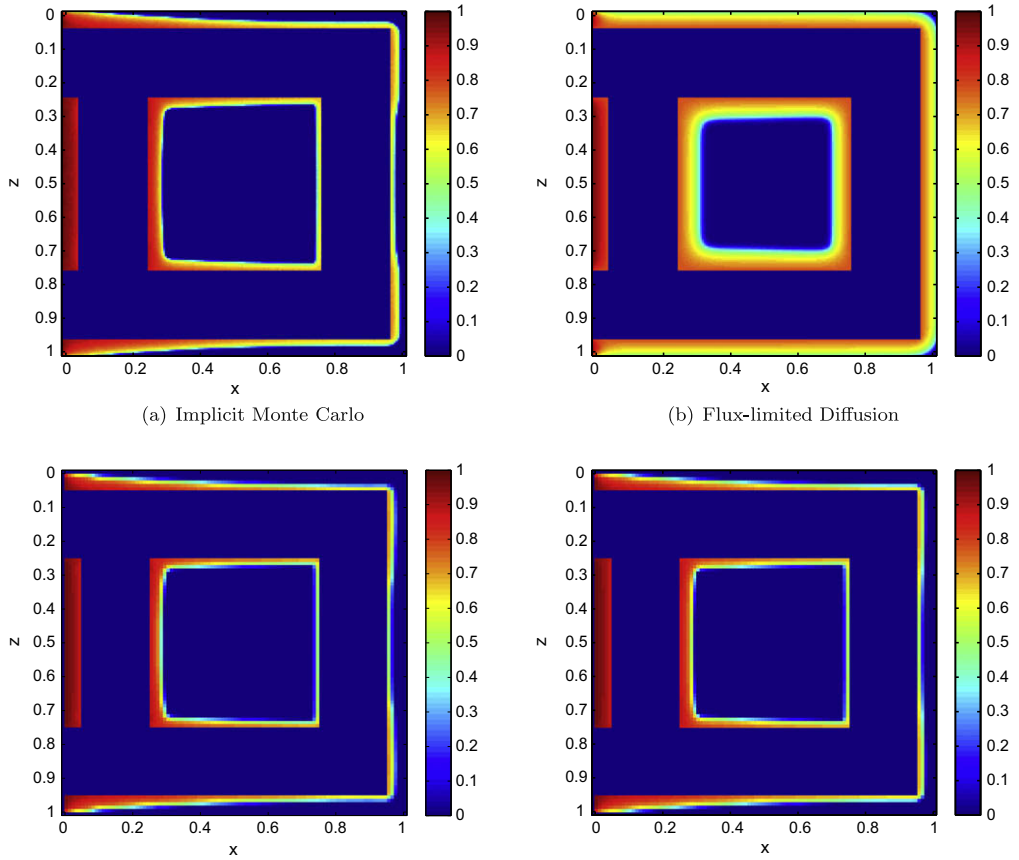

Fig. 14. Material temperature T_{rad} for the hohlraum problem at $t = 1$ ns.

In Fig. 11, we display the radiation temperature

$$T_{\text{rad}} := \sqrt[4]{\frac{\phi}{ac}}$$

in the hohlraum at $t = 1$ ns, for each of the different methods. The P_N and S_N solutions were obtained using $N_x = N_z = 100$ and $\Delta t = 0.3\Delta x/c$. The IMC solution used 10^6 particles per time step with a time step size of 10^{-2} ns and 200 mesh cells per

direction. The flux-limited diffusion solution used Larsen's flux limiter [40] with $n = 2$ and 200 mesh cells in each direction. We do not show a standard P_7 solution, because in that simulation, negative values in the radiation energy caused the material energy to also become negative, which subsequently crashed the simulation. We also note that early in time, the radiation energy for the filtered P_7 solution also has some negative values in the vacuum regions at early times. However, the magnitude of these values was much smaller and quickly decayed and, as a result, did not cause the material to cool from its initial temperature.

From Fig. 11 we can see that there is a stark difference between the flux-limited diffusion solution and the other solutions. The diffusion solution has a radiation temperature that is uniformly distributed around the block, whereas the other solutions have noticeably less radiation to the right of the block. In addition, the radiation in the diffusion solution has penetrated much farther into the walls and block.

To better differentiate between the IMC, filtered P_7, and S_8 solutions, we plot the solution along the lines defined by $z = 0.125$ cm and $x = 0.85$ cm in Figs. 12 and 13. These lines are half-way between the wall and block on the bottom and right part of the problem respectively. From these figures, the discrepancy between the flux-limited diffusion solution and the other solutions is apparent; the noise in the IMC solution is also noticeable. In both figures, the IMC and filtered P_7 solutions are in rough agreement with the S_8 solution under-predicting the radiation temperature relative to the other two. This is an encouraging result for filtered spherical harmonics, because IMC is considered the benchmark for this problem.

The material temperature for each method is shown in Fig. 14. In these plots we observe that the block in the flux-limited diffusion case has been almost uniformly heated. Meanwhile, in the S_8 solution, ray effects have created hot spots on the back wall of the problem and at the back of the block. The IMC and filtered P_7 solution appear to be in good agreement, despite the fact the IMC solution uses a higher spatial resolution.

7. Conclusion

We have presented a filtered spherical harmonic expansion as a means of making the standard P_N approximation more robust and less oscillatory. The filtering procedure changes the standard spherical harmonic reconstruction of the intensity in order to suppress oscillations in the reconstruction. We also showed how filters can be implemented in existing spherical harmonics codes in a manner that is convergent and preserves the equilibrium diffusion limit.

Numerical results demonstrate that the application of filters to the standard P_N equations improved solutions considerably. For the line-source problem, the filtered P_N equations produce solutions that are comparable to S_N and P_N methods at long times but vastly superior at short times. In another problem of radiative transfer, a 2D Cartesian hohlraum problem, the filtered P_N solution have comparable accuracy to an implicit Monte Carlo solution, without the ray effects found in S_N simulations. The standard P_7 could not solve this problem because the negative radiation energies caused the material temperature to become negative, causing the calculation to crash.

We believe that filtered P_N expansions present a promising approach to solving radiative transfer problems. Filters take advantage of the benefits in the original expansion while reducing (or eliminating) the unwanted oscillations. While we feel that the properties of the filter we presented – namely preserving the equilibrium diffusion limit and a filter strength that vanishes as the expansion order goes to infinity – are necessary for any filter to have, we do not contend that our particular choice of filter is the best possible. We do, however, point to our numerical results that demonstrate that our filter does lead to accurate solutions without large oscillations. Other ways for choosing the filter strength are possible, and may prove to be more effective on a certain subset of problems. Besides the choice of filter strength, the theoretical underpinnings of our method should be investigated. We noted that discontinuities in the opacity introduces discontinuities in the filter and that the proper continuum limit of the FP_N equations is unknown. We expect that further investigation will uncover how these concerns can be mitigated/eliminated.

In the future we plan on applying our filtering approach to implicit numerical methods for the P_N equations in 2D axisymmetric and 3D coordinates, including the multigroup energy discretization. We are hopeful that results for those problems will be as encouraging as those presented above. Also, future work will need to address the computational cost of our filtered expansions. The results in this paper were obtained using a semi-implicit time integration scheme. Such a time integration method would see its largest benefits on massively parallel architectures; a topic we hope to explore in the future. It is our belief that with further research our method can compete with the current paradigms used for large-scale radiation transport problems in terms of computational cost.

Acknowledgments

The submitted manuscript has been authored, in part, by a contractor of the U.S. Government under Contract No. DE-AC05-00OR22725. Accordingly, the U.S. Government retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

References

[1] C.L. Fryer, G. Rockefeller, M.S. Warren, SNSPH: a parallel three-dimensional smoothed particle radiation hydrodynamics code, Astrophysical Journal 643 (2006) 292–305.

[2] F. Swesty, E. Myra, A numerical algorithm for modeling multigroup neutrino-radiation hydrodynamics in two spatial dimensions, The Astrophysical Journal Supplement Series 181 (2009) 1–51.
[3] M. Matzen, M. Sweeney, R. Adams, J. Asay, Pulsed-power-driven high energy density physics and inertial confinement fusion research, Physics of Plasmas 12 (5) (2005) 055503.
[4] M. Marinak, G. Kerbel, N. Gentile, O. Jones, Three-dimensional HYDRA simulations of national ignition facility targets, Physics of Plasmas 8 (5) (2001) 2275.
[5] A. Reighard, R. Drake, K. Dannenberg, Observation of collapsing radiative shocks in laboratory experiments, Physics of Plasmas 13 (8) (2006) 082901.
[6] R.P. Drake, Theory of radiative shocks in optically thick media, Physics of Plasmas 14 (4) (2007) 043301.
[7] R.P. Drake, High Energy Density Physics, Springer, 2006.
[8] J.E. Morel, T.A. Wareing, R.B. Lowrie, D.K. Parsons, Analysis of ray-effect mitigation techniques, Nuclear Science and Engineering 144 (1) (2003) 1–22.
[9] K.A. Mathews, On the propagation of rays in discrete ordinates, Nuclear Science and Engineering 132 (1999) 155.
[10] I.K. Abu-Shumays, Angular quadratures for improved transport computations, Transport Theory and Statistical Physics 30 (2–3) (2001) 169–204.
[11] J.A. Fleck Jr., J.D. Cummings, An implicit Monte Carlo scheme for calculating time and frequency dependent nonlinear radiation transport, Journal of Computational Physics 8 (1971) 313–342.
[12] E.W. Larsen, B. Mercier, Analysis of Monte Carlo method for nonlinear radiative transfer, Journal of Computational Physics 71 (1987) 50–64.
[13] R.G. McClarren, T.J. Urbatsch, A modified implicit Monte Carlo method for time-dependent radiative transfer with adaptive material coupling, Journal of Computational Physics 228 (16) (2009) 5669–5686.
[14] R.G. McClarren, J.P. Holloway, T.A. Brunner, On solutions to the P_n equations for thermal radiative transfer, Journal of Computational Physics 227 (5) (2008) 2864–2885.
[15] G.L. Olson, Second-order time evolution of P_N equations for radiation transport, Journal of Computational Physics 228 (8) (2009) 3072–3083.
[16] D.S. Kershaw, Flux limiting nature's own way – a new method for numerical solution of the transport equation, Technical Report, UCRL-78378, Lawrence Livermore National Laboratory, July 1976.
[17] G.L. Olson, L.H. Auer, M.L. Hall, Diffusion, P_1, and other approximate forms of radiation transport, Journal of Quantitative Spectroscopy and Radiative Transfer 64 (6) (2000) 619–634.
[18] T.A. Brunner, One-dimensional Riemann solvers and the maximum entropy closure, Journal of Quantitative Spectroscopy and Radiative Transfer 69 (5) (2001) 543–566.
[19] B. Su, Variable eddington factors and flux limiters in radiative transfer, Nuclaer Science and Engineering 137 (3) (2001) 281–297.
[20] G.C. Pomraning, A generalized P_N approximation for neutron transport problems, Nukleonik 6.
[21] K.S. Oh, J.P. Holloway, A quasi-static closure for 3rd order spherical harmonics time-dependent radiation transport in 2D, Proceedings of the 2009 International Conference on Mathematics and Computational Methods and Reactor Physics, American Nuclear Society, 2008.
[22] C.D. Hauck, High-order entropy-based closures for linear transport in slab geometries, Communications in Mathematical Sciences, in press.
[23] P. Monreal, M. Frank, Higher order minimum entropy approximations in radiative transfer. Arxiv preprint arXiv:0812.3063, 2008.
[24] D. Wright, M. Frank, A. Klar, The minimum entropy approximation to the radiative transfer equation (arXiv:0812.3063 [comp-ph]).
[25] C.D. Hauck, R.G. McClarren, Positive P_N closures, SIAM Journal on Scientific Computing, submitted for publication.
[26] R.G. McClarren, C.D. Hauck, R.B. Lowrie, Filtered spherical harmonics methods for transport problems, Proceedings of the 2009 International Conference on Mathematics and Computational Methods and Reactor Physics, American Nuclear Society, 2008.
[27] G.C. Pomraning, The Equations of Radiation Hydrodynamics, Dover Publications, 2005.
[28] T.A. Brunner, Riemann Solvers for Time-Dependent Transport Based on the Maximum Entropy and Spherical Harmonics Closures, Ph.D. Thesis, University of Michigan, 2000.
[29] J.P. Boyd, Chebyshev and Fourier Spectral Methods, Dover Publications, Mineola, New York, 2001.
[30] K.M. Case, P.F. Zweifel, Linear Transport Theory, Addison-Wesley, Reading, Massachusetts, 1967.
[31] E. Larsen, G. Pomraning, V. Badham, Asymptotic analysis of radiative transfer problems, Journal of Quantitative Spectroscopy and Radiative Transfer, 29 (4).
[32] R.G. McClarren, T.M. Evans, R.B. Lowrie, J.D. Densmore, Semi-implicit time integration for P_N thermal radiative transfer, Journal of Computational Physics 227 (16) (2008) 7561–7586.
[33] C.D. Hauck, R.B. Lowrie, R.G. McClarren, Methods for diffusive relaxation in the P_N equations, in: G. Russo, G. Puppo (Eds.), Numerical Methods for Relaxation Systems and Balance Equations, Quaderni di Matematica, Dipartimento di Matematica, Seconda Universita di Napoli, Italy, in press.
[34] R.G. McClarren, R.B. Lowrie, The effects of slope limiting on asymptotic-preserving numerical methods for hyperbolic conservation laws, Journal of Computational Physics 227 (23) (2008) 9711–9726.
[35] B.D. Ganapol, Homogeneous infinite media time-dependent analytic benchmarks for X-TM transport methods development, Technical Report, Los Alamos National Laboratory, March 1999.
[36] R.G. McClarren, J.P. Holloway, T.A. Brunner, Analytic P_1 solutions for time-dependent, thermal radiative transfer in several geometries, Journal of Quantitative Spectroscopy and Radiative Transfer 109 (3) (2008) 389–403.
[37] T.A. Brunner, Forms of approximate radiation transport, Technical Report, SAND2002-1778, Sandia National Laboratories, 2002.
[38] T.A. Brunner, J.P. Holloway, Two-dimensional time dependent Riemann solvers for neutron transport, Journal of Computational Physics 210 (2005) 386–399.
[39] J.D. Densmore, E.W. Larsen, Asymptotic equilibrium diffusion analysis of time-dependent Monte Carlo methods for gray radiative transfer, Journal of Computational Physics 199 (2004) 175–204.
[40] J.E. Morel, Diffusion-limit asymptotics of the transport equation, the $P_{1/3}$ equations, and two flux-limited diffusion theories, Journal of Quantitative Spectroscopy and Radiative Transfer 65 (2000) 769–778.