HERO - A 3D general relativistic radiative post-processor for accretion discs around black holes

The Harvard community has made this article openly available. Please share how this access benefits you. Your story matters

HERO - A 3D General Relativistic Radiative Postprocessor for Accretion Discs around Black Holes

Yucong Zhu[1]⋆ Ramesh Narayan[1]⋆, Aleksander Sadowski[2]⋆, Dimitrios Psaltis[3]⋆,

[1]*Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA*

[3]*MIT Kavli Institute for Astrophysics and Space Research, 77 Massachusetts Ave, Cambridge, MA 02139, USA*

[3]*University of Arizona, 933 N. Cherry Ave, Tucson, AZ 85721, USA*

20 May 2015

ABSTRACT

HERO (Hybrid Evaluator for Radiative Objects) is a 3D general relativistic radiative transfer code which has been tailored to the problem of analyzing radiation from simulations of relativistic accretion discs around black holes. HERO is designed to be used as a postprocessor. Given some fixed fluid structure for the disc (i.e. density and velocity as a function of position from a hydrodynamics or magnetohydrodynamics simulation), the code obtains a self-consistent solution for the radiation field and for the gas temperatures using the condition of radiative equilibrium. The novel aspect of HERO is that it combines two techniques: 1) a short characteristics (SC) solver that quickly converges to a self consistent disc temperature and radiation field, with 2) a long characteristics (LC) solver that provides a more accurate solution for the radiation near the photosphere and in the optically thin regions. By combining these two techniques, we gain both the computational speed of SC and the high accuracy of LC. We present tests of HERO on a range of 1D, 2D and 3D problems in flat space and show that the results agree well with both analytical and benchmark solutions. We also test the ability of the code to handle relativistic problems in curved space. Finally, we discuss the important topic of ray-defects, a major limitation of the SC method, and describe our strategy for minimizing the induced error.

Key words: methods: numerical – radiative transfer – accretion, accretion discs – black hole physics

1 INTRODUCTION

Radiative transport (RT) plays a crucial role in astrophysics. It governs our primary source of information about the cosmos, viz., the luminosities and spectra of cosmic sources[1]. In addition, radiation acts as a channel for energy transport, shaping the dynamics and impacting the evolution of astrophysical systems on all scales: stellar evolution, stellar and galactic winds, super-Eddington accretion, epoch of reionization, etc.

Radiation is quite challenging to model since it behaves in highly nonlinear and nonlocal ways. The radiation field is determined by a six-dimensional (6D) system of coupled equations which link spatial positions (3D), ray directions (2D) and frequencies (1D), all as a function of a seventh coordinate, time (in the general time-dependent problem). The high dimensionality of the solution vector, combined with the complex integro-differential nonlocal nature of the problem, results in an extremely taxing computational challenge (both in terms of memory and computational speed).

Because of the intrinsic complexity of the problem, the earliest radiative solvers made use of fairly restrictive approximations. For example, enforcing spatial symmetries allows one to greatly reduce the dimensionality of the problem and led to the first generation of 1D/2D radiative codes (e.g. Feautrier method that exploits symmetric and antisymmetric radiation moments, see Mihalas et al. 1978). Several of the multidimensional RT ideas have also found application in the field of neutrino transport in core collapse supernovae (Burrows et al. 2000; Rampp & Janka 2002; Liebendorfer et al. 2004; Livne et al. 2004; Hanke et al. 2013). In this case, the infalling matter becomes optically thick to the neutrino flux, which propagates analagously to radiative transport in optically thick media.

3D radiative problems are extremely computationally

⋆ E-mail:
yzhu@cfa.harvard.edu (YZ);   rnarayan@cfa.harvard.edu (RN);
asadowsk@mit.edu (AS); dpsaltis@email.arizona.edu (DP);
[1] We do receive additional information from cosmic rays, neutrinos and, hopefully soon, gravitational waves, but these pale next to the enormous volume of information we receive via electromagnetic radiation.

taxing, which motivates the need for highly efficient techniques. Recent applications that make use of 3D RT include models of young stellar objects (Wolf, Fischer & Pau 1998), protostellar to protoplanetary discs (Indebetouw et al. 2006; Niccolini & Alcolea 2006), reflection nebulae (Witt & Gordon 1996), molecular clouds (Steinacker et al. 2005; Pelkonen, Juvela & Padoan 2009), spiral galaxies (Bianchi 2008; Schechtman-Rook, Bershady & Wood 2012), interacting and starburst galaxies (Chakrabarti et al. 2007; Hayward et al. 2011), AGNs (Schartmann et al. 2008; Stalevski et al. 2012), and cosmological reionization simulations (Ferland et al. 1998; Abel & Wandelt 2002; Finlator et al. 2009). These applications are primarily concerned with point sources illuminating optically thick media.

In more complex optically thick flows where extended diffuse emission is important, the typical approach is to simplify the directional structure of the radiation field. One idea is to decompose the radiation field into its moments (the first three moments being the radiation energy density, radiation flux, and radiation pressure) and to evolve the radiation field locally as a fluid using some closure relation on the moments. This approach is popular due to its local nature and fast speed, and has been implemented in many hydrodynamic codes (Turner & Stone 2001; Bruenn et al. 2006; Hayes et al. 2006; Gonzáles, Audit, & Huynh 2007; Krumholz et al. 2007; Gittings et al. 2008; Ohsuga et al. 2009; Swesty & Myra 2009; Commerçon et al. 2011; Zhang et al. 2011; Kolb et al. 2013). Flux limited diffusion (FLD, Levermore & Pomraning 1981), which is based on the first two moments, is the simplest and perhaps most popular moment-based radiative method. M1 closure is a generalization of the FLD method which uses the first three moments. It has recently gained traction as a fast solver for diffuse emission (Levermore 1984; Dubroca & Feugeas 1999; Stone et al. 1992; Gonzáles, Audit, & Huynh 2007; Sądowski et al. 2013).

The field of protoplanetary disc dust modelling is one area that has been active in developing RT techniques. In this domain, the monte-carlo approach has been the dominant paradigm owing to its ease of implementation and ability to handle anisotropic scattering kernels (Lopez, Mekarnia, & Lefevre 1995; Niccolini, Woitke, & Lopez 2003; Wolf, Henning & Stecklum 1999; Bjorkman & Wood 2001; Pinte, Duchene, & Bastien 2006). Monte-carlo RT has also been applied to the problem accretion discs, specifically to handle the problem of modeling Compton scattering in disc coronas (Dolence et al. 2009; Schnittman & Krolik 2013; Ghosh 2013) as well as dust scattering of emission from the Galactic centre (Odaka et al. 2011). The primary downside of this technique is the computational expense and noise associated with photon statistics. The method is also limited to modest optical depths.

The limitations of the various methods mentioned above led to the development of more deterministic discretized finite-differencing methods (Stenholm, Stoerzer, & Wehrse 1991; Steinacker, Bacmann, & Henning 2002) and discrete ordinates characteristic methods (Dullemond & Turolla 2000; Steinacker, Bacmann, & Henning 2002; Hayes & Norman 2003; Vögler et al. 1982; Heinemann et al. 2006; Woitke, Kamp & Thi 2009; Hayek et al. 2010; Davis et al. 2012; Jiang et al. 2014), which have a number of widely discussed advantages. However, all of these methods suffer from the phenomenon of "ray-defects" (see appendix A for

a detailed discussion). Another twist to the discretization strategy involves expanding and evolving the radiation field as a combination of spherical harmonics (Szu-cheng & Kuo-Nan 1982; Evans 1997; McClarren, Holloway, & Brunner 2008). Spherical harmonics respect rotational symmetry and therefore do not suffer from the linear ray defect patterns that plague other discrete ordinate methods. The tradeoff is that spherical methods typically suffer from artificial sidelobe patterns due to the finite order cutoff used in taking the series expansion of the radiation field.

Raytracing methods are also used for predicting the observed flux from astrophysical objects. These codes typically do not solve for the global radiation field within an object; instead the focus is on calculating the apparent intensities that reach a distant observer. In these codes, the typical assumption is to ignore scattering (i.e. to ignore nonlocal coupling of the radiation field) since this allows one to immediately compute the evolution of ray intensity by simply integrating the local emissivities along a photon geodesic (e.g. Cunningham & Bardeen 1973; Özel & Di Matteo 2001; Huang et al. 2007; Dexter & Agol 2009; Shcherbakov & Lei 2011; Vincent et al. 2011; Chan et al. 2013; Bohn et al. 2014).

Despite the multitude of radiative solvers available, none are currently able to solve the problem of optically thick emission from accretion discs around black holes (BHs). The main difficulty here is that a general relativistic 3D framework is needed to properly account for both light bending effects and doppler/gravitational redshifts. The work reported here is a first attempt at tackling the radiation problem in full glory around black holes. The code we describe here called HERO is a *postprocessor* – given some fixed gas structure (density, velocity, energy injection rate), as determined by a separate BH accretion disc simulation, our code *solves for the radiation field along with its self-consistent temperature solution*. We explain here how our code works and show verification tests to gauge its performance under various conditions.

The organization of the paper is as follows. In §2, we first describe the method by which HERO operates. We give a brief overview of the radiative transfer problem that HERO solves, and how it generalizes in curved space. We also detail how the two radiative solvers (short/long characteristics) operate, and the numerical methods involved (i.e. acceleration schemes, interpolation, discretization strategy). Next, in §3 we check the code by comparing to analytic and benchmark results for 1D, 2D, and 3D test configurations. Finally, in appendix §A we end with a discussion of ray-defects, a systematic problem that plagues short-characteristic radiative solvers.

2   RADIATIVE SOLVER

Our radiative code consists of three primary components:

(i) A **short characteristic solver** for quickly obtaining an approximate solution to the radiation field
(ii) A **long characteristic solver** for more accurate modeling of the radiation field
(iii) An optional **raytracer** for generating mock observations from some distant observing plane

We employ a *hybrid* approach to solve the radiation field, hence the name of the code: HERO (Hybrid Evaluator for Radiative Objects). The code is hybrid in the sense that it uses both (i) "short characteristics" (SC, Mihalas et al. 1978) and (ii) "long characteristics" (LC, Feautrier 1964). Here "short" and "long" refer to the length of the light rays that are considered in each iteration of the solver. The short characteristics method is only concerned with propagation of radiation from a given cell to its immediate neighbours, whereas the long characteristic method traces rays all the way to the edge of the computational grid.

The motivation for developing a hybrid approach is to allow for the accurate modelling of radiation in optically thin regions via long characteristics, while retaining the computational speed offered by short characteristics in optically thick regions (see §2.1 & §2.3 for a detailed comparison of the methods). For any given problem, we first apply short characteristics to quickly solve for the local radiation/temperature. We then feed this solution as the initial guess for a more detailed long characteristics calculation. Using such a hybrid approach allows us to combine the "best of both worlds."

In addition to solving for the radiation field, in both SC and LC, we also solve for a self-consistent gas temperature. The general idea is to loop back and forth between solving for the radiation field given a fixed temperature structure, and solving for the equilibrium temperature distribution given a fixed radiation field. This procedure is iterated until convergence, leaving us with a self-consistent solution for both the radiation and the temperature.

Our goal in developing HERO is to model/investigate the observed properties of relativistic accretion discs around black holes, and to compare the radiative properties of simulated discs with data collected by earthbound telescopes. Therefore, after completing the SC and LC steps described above, the final step is to input the radiative and temperature structure as computed from LC and to generate synthetic observations of the disc as seen by a distant observer. For this stage we use the standard raytracing approach (see §2.5 for details).

Due to the high temperatures present in accretion discs, Comptonization plays a crucial role in determining the observed radiation from relativistic disc photospheres. Because of the complex nature of the Compton scattering kernel, we defer discussion of Comptonization to a follow-up paper, which will focus exclusively on explaining and testing our relativistic scattering module.

2.1 Short Characteristics

Short characteristics is a popular algorithm for tackling multidimensional radiative transfer problems (Mihalas et al. 1978; Olson & Kunasz 1987; Kunasz & Auer 1988). It is a local method and hence very fast. The basic idea is to solve for the radiation field at a given point (reference cell) using only the information provided by neighbouring cells (see schematic in Figure 1).
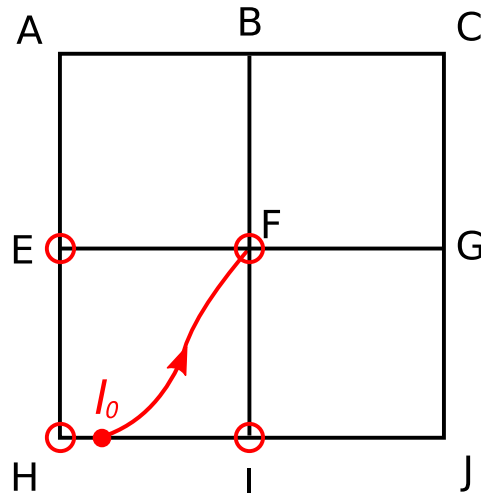
Figure 1. Schematic of the short characteristics method. A curved ray (null geodesic) is shot "upstream" to determine the intensity at F. This ray is terminated at the neighbouring cell boundaries, and the intensity I_0 at the boundary is computed by interpolation of neighbouring grid points (in this case points H and I). The source function $S(\tau')$ in Eq. 22 is also obtained by interpolation on neighbouring points (points E, F, H, I), and the intensity at F is thereby calculated. The procedure is repeated for a number of rays in different directions to obtain an estimate of the radiation field at F.

2.1.1 Ordinary Radiative Transfer Equation

To set the stage, we first discuss the standard case of flat space, where the radiative transfer equation takes the form:

$$\frac{dI_\nu}{ds} = -(\kappa_\nu + \sigma_\nu)I_\nu + j_\nu + \sigma_\nu \int \phi_\nu(\Omega, \Omega')I_\nu(\Omega')d\Omega'. \quad (1)$$

Here I_ν is the intensity of a ray travelling along path s directed towards Ω; κ_ν and σ_ν are the absorption and scattering coefficients; ϕ_ν is the scattering shape function, normalized such that $\int \phi_\nu(\Omega, \Omega')d\Omega' = 1$; and j_ν is the emission coefficient. Typically, Eq.1 is simplified to the form

$$\frac{dI_\nu}{d\tau_\nu} = -I_\nu + S_\nu, \quad (2)$$

where the the optical depth τ_ν is given by

$$d\tau_\nu = (\kappa_\nu + \sigma_\nu)ds, \quad (3)$$

and the source function S_ν is defined as

$$S_\nu = \frac{j_\nu + \sigma_\nu \int \phi_\nu(\Omega, \Omega')I_\nu(\Omega')d\Omega'}{\kappa_\nu + \sigma_\nu}. \quad (4)$$

In the case of isotropic scattering, $\phi_\nu = 1/4\pi$, and the scattering term simplifies to

$$\int \phi_\nu(\Omega, \Omega')I_\nu(\Omega')d\Omega' = \frac{1}{4\pi} \int I_\nu(\Omega')d\Omega' \equiv J_\nu, \quad (5)$$

where J_ν is the angle-averaged mean intensity. We assume isotropic scattering in all of the work reported here. We further assume a thermal source, for which the emissivity is given by $j_\nu = \kappa_\nu B_\nu(T)$, where $B_\nu(T)$ is the Planck function corresponding to the local temperature T of the medium. With these two simplifications, the source function can be compactly rewritten as:

$$S_\nu = \epsilon_\nu B_\nu + (1 - \epsilon_\nu)J_\nu, \quad (6)$$

where ϵ_ν is the photon interaction destruction probability, given by the ratio of the absorption and total opacity:

$$\epsilon_\nu \equiv \frac{\kappa_\nu}{\kappa_\nu + \sigma_\nu}. \tag{7}$$

To avoid confusion with summation indices, in the remainder of this paper we will no longer use a subscript ν to note the frequency dependence of radiative quantities. Unless otherwise stated, this will apply to all subsequent intensities, opacities, source functions, and emissivities.

2.1.2 Covariant Formulation of the Radiative Transfer Equation

Because HERO has been developed for relativistic problems, we use a generalization of the radiative transfer equation that accounts for relativistic effects such as light bending and doppler boosting. The generalization of the radiative transfer equation to curved space has been discussed in the literature (Mihalas & Mihalas 1984; Lindquist 1966), but we give a brief primer here for completeness. Starting from the time dependent version of Eq. 1:

$$\frac{1}{c}\frac{dI}{dt} + (\hat{n} \cdot \nabla)I = G, \tag{8}$$

where G on the righthand side contains all scattering and absorption source terms that interact with the radiation field, and $\hat{n}$ is a normalized 3-vector denoting the propagation direction of photons. We can rewrite Eq. 8 in a more revealing form by introducing the direction 4-vector k^α whose orthonormal form is (note that $k \cdot k = 0$):

$$k^\alpha = (1, \hat{n}). \tag{9}$$

This lets us rewrite the RT equation (setting $c = 1$) as:

$$k^\alpha \frac{d}{dx^\alpha}I = G \equiv -(\kappa + \sigma)I + j + \sigma J. \tag{10}$$

By considering the Lorentz transformation properties of the various quantities, Eq.10 can be recast in an invariant form (see Mihalas & Mihalas 1984, §7.1):

$$p^\alpha \frac{\partial}{\partial x^\alpha}\mathcal{I} = \mathcal{G}, \tag{11}$$

where the photon 4-momentum $p^\alpha = h\nu k^\alpha$ takes on the role of the propagation 4-vector. Instead of the usual definitions of intensity I and source term G, this equation considers their corresponding relativistically invariant versions,

$$\mathcal{I} = I/\nu^3, \qquad \mathcal{G} = hG/\nu^2, \tag{12}$$

where the scaling of I is determined by photon number conservation in phase space. The $(h\nu)^{-2}$ scaling for G arises after transforming Eq.10 to use $\mathcal{I}$ and allows us to construct the invariant scalings for the absorption and scattering coefficients as:

$$\mathfrak{k}_\nu = h\nu\kappa_\nu, \qquad \mathfrak{s}_\nu = h\nu\sigma_\nu, \tag{13}$$

and for the emissivity and source function,

$$\mathfrak{j}_\nu = hj_\nu/\nu^2, \qquad \mathcal{S} = S/\nu^3. \tag{14}$$

In terms of these new quantities, the invariant source $\mathcal{G}$ becomes

$$\mathcal{G} = -(\mathfrak{k} + \mathfrak{s})(\mathcal{I} - \mathcal{S}). \tag{15}$$

For convenience, since we usually work with comoving quantities, κ_c, σ_c, S_c, we can instead write $\mathcal{G}$ as:

$$\mathcal{G} = -h\nu_c(\kappa_c + \sigma_c)\left[-\mathcal{I} + \frac{S_c}{\nu_c^3}\right]$$
$$= p^\alpha u_\alpha(\kappa_c + \sigma_c)\left[-\mathcal{I} + \frac{S_c}{\nu_c^3}\right], \tag{16}$$

where u^μ is the four-velocity of the fluid. Eq. 11 is in manifestly covariant form and the generalization to curved space is simply accounted for by replacing the directional derivative $\partial/\partial x^\alpha$ with the covariant derivative:

$$\frac{D}{Dx^\alpha} = \frac{\partial}{\partial x^\alpha} - \Gamma^\gamma_{\alpha\beta}p^\beta\frac{\partial}{\partial p^\gamma}, \tag{17}$$

evaluated along a null geodesic (Lindquist 1966). The general covariant form of the RT equation then takes the form:

$$p^\alpha\frac{D}{Dx^\alpha}\mathcal{I} = \mathcal{G} \tag{18}$$

We can further simplify the above expression by using the photon geodesic equation,

$$\frac{dp^\gamma}{d\lambda} = -\Gamma^\gamma_{\alpha\beta}p^\alpha p^\beta, \tag{19}$$

where λ is an affine parameter defined by $p^\alpha = dx^\alpha/d\lambda$. Using Eq. 19 to eliminate the Christoffel symbols in Eq. 17, the RT equation becomes:

$$p^\alpha\frac{\partial\mathcal{I}}{\partial x^\alpha} + \frac{dp^\alpha}{d\lambda}\frac{\partial\mathcal{I}}{\partial p^\alpha} = \mathcal{G}. \tag{20}$$

Finally, using the chain rule for partial differentiation, we obtain the very simple equation

$$\frac{d\mathcal{I}}{d\lambda} = \mathcal{G}. \tag{21}$$

This is the form of the covariant RT equation that we employ in HERO. The only point that requires some care is choosing the correct frequency when evaluating the source term $\mathcal{G}$. For instance, if we wish to evaluate $\mathcal{I}$ at frequency ν in the lab frame, but we compute $\mathcal{G}$ in the fluid comoving frame (the natural frame since the opacity is defined there), then ν_c in Eq. 16 must be chosen such that it corresponds to the lab frame ν used to label $\mathcal{I}$.

2.2 Implementation of Short Characteristics

Given a ray in the reference cell F, we compute the upstream location where it intersects the boundary of that cell (see Fig. 1) and use this location to determine the relative weights contributed by each of the four nearest neighbours on that boundary. In flat space, it is possible to write down an analytical formula for the intersection point (e.g. Dullemond & Turolla 2000), but there is no generalization in terms of simple functions for curved space. HERO does the calculation numerically.

Although this is computationally somewhat expensive, it needs to be done only once as part of the initialization steps. The information from this initial computation is saved and used repeatedly during the SC iterations. Given the intersection of a ray with the nearest neighbour cell boundary, we compute the intensity of this ray in the reference cell F using the radiative transfer equation (Eq. 21):

$$I_F = I_0\exp(-\tau_0) + \int_0^{\tau_0}S(\tau')\exp(-\tau')d\tau', \tag{22}$$

where I_0 is the incoming intensity at the boundary, and τ_0 is the total optical depth of the photon path leading from the boundary to point F (computed using the average of the central and boundary opacities). The intensity I_0 is computed by interpolating from the four boundary points. Our current implementation of short characteristics uses linear interpolations at various stages (this is easily improved in the future). In this spirit, the integral evaluation in Eq. 22 is computed as

$$\int_0^\tau S(\tau')\exp(-\tau')d\tau'$$
$$= \frac{e^{-\tau}(S_0 - S_\tau(1+\tau)) + S_0(\tau-1) + S_\tau}{\tau}, \qquad (23)$$

which is the analytic solution when the source function S varies linearly with τ along the photon trajectory. Higher order interpolation schemes have been explored in other codes (e.g. Kunasz & Auer 1988); however, special care is needed to ensure non-negativity of the source function.

Our implementation of SC in HERO retains some memory of the radiative quantities from the past iteration. This helps with stability. Thus, in computing the intensity at iteration n, we use the linear combination

$$I^n = (1-m)I_{FS}^{n-1} + mI^{n-1}. \qquad (24)$$

where I_{FS}^{n-1} is the result from the RT formal solution (i.e. evolving the radiation field through SC). Typically, we set $m = 0.5$. HERO with SC converges within a reasonable number of iterations $\sim 10 - 1000s$, depending on the degree of scattering/coupling in the problem.

The above discussion assumed that the gas temperatures are given. More often, one does not know the temperature but has to solve for it based on boundary conditions and internal heat sources. In this case, between every two of the above iterations for the intensity, the code carries out a round of iterations (typically $\sim 10s$) to solve for the temperature using the condition of radiative equilibrium, i.e., the requirement that the heating and cooling rates of the gas should balance. Radiative equilibrium requires at each position x

$$Q_{\text{cool}}(x) = Q_{\text{heat}}(x) + Q_{\text{inject}}(x), \qquad (25)$$

where

$$Q_{\text{cool}}(x) = \int \kappa_\nu(x)B_\nu[T(x)]d\nu ,$$
$$Q_{\text{heat}}(x) = \int \kappa_\nu(x)J_\nu(x)d\nu, \qquad (26)$$

and $Q_{\text{inject}}(x)$ refers to any additional injection of energy into the fluid, e.g., through viscous dissipation in an accretion disc.

2.3 Long Characteristics

Similar to the short characteristics method, our long characteristics solver obtains the radiation field by evolving ray intensities according to the radiative transfer equation (c.f. Eq. 21). Figure 2 shows a schematic of our approach to LC – rays are shot upstream, and we evaluate the intensity at the observing cell according to

$$I_0 = I_b \exp(-\tau_b) + \int_0^{\tau_b} S(\tau')\exp(-\tau')d\tau'. \qquad (27)$$
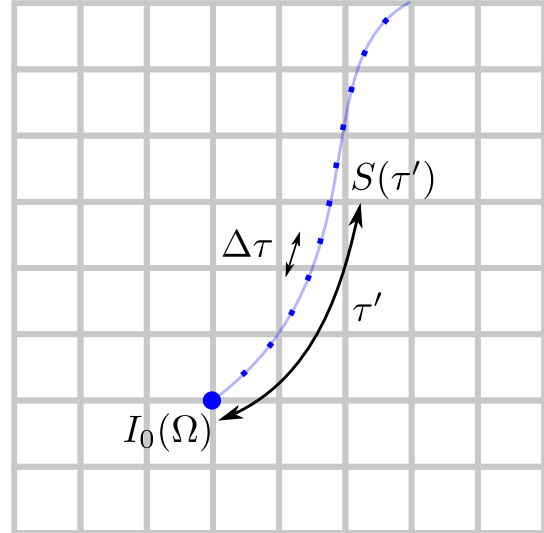
Figure 2. Schematic of long characteristics RT solver. We calculate the local mean intensity at an observing cell (large blue circle) by integrating the source function along photon geodesics (light blue path). The RT integral is evaluated by direct summation of contributions from linear piece-wise segments (small blue dots).

The above expression is identical to our SC calculation (Eq.22), except that I_b now represents the intensity distribution at the grid boundaries and the ray integral continues past the immediate neighbouring cells to traverse the entire spatial domain. The integral in Eq. 27 is evaluated in discrete steps ($\tau' \rightarrow \tau' + \Delta\tau$), where each source-function piece is evaluated using the same linear interpolation scheme as used in SC (Eq. 23).

One major difference between our LC and SC solvers is the choice of angular grid. In SC we used a fixed angular resolution ($N_A = 80$, see §2.7.3) for all cells. This fairly low resolution choice is adequate (though only barely) in SC since one only needs to resolve the neighbouring 27 grid cells. In LC, since the rays traverse the entire spatial domain, much higher angular resolution is needed to resolve the emission from distant source cells. This becomes particularly problematic with spherical coordinates where the grid spans a few decades in radius. Cells located at larger radii have a difficult time "seeing" the inner regions unless one takes care in choosing the ray directions.

Our solution is to dynamically set the angular resolution used by the LC solver such that there is sufficient coverage to resolve all the cells within the domain. Consider, as a typical example, a 3D grid in spherical coordinates which extends from an inner radius r_{in} to an outer radius $r_{\text{out}} \gg r_{\text{in}}$. The radial dynamic range may be up to three decades, as in some of the examples described later. Let us suppose that our spherical spatial grid uses uniform angular spacing in θ, with N_θ cells, and logarithmic spacing in radius with $N_r \lesssim N_\theta$ cells per decade. Consider a cell at some radius r. For this cell, the apparent density of other cells as viewed in its local "sky" is highly anisotropic. For rays traveling from the direction of the coordinate centre, this cell needs to sample the local radiation field with an angular resolution of order $\Delta\theta_{\text{res}} \sim r_{\text{in}}/rN_\theta$ in order to adequately sample the contribution of all the cells at the centre. In the opposite

direction, however, an angular resolution of order $\sim 1/N_\theta$ is sufficient.

The above requirement can be achieved by designing a suitable angular tiling of the sky. One simple prescription is to vary the angular resolution $\Delta\theta_{\rm res}$ as a function of direction χ relative to the local radius vector as follows:

$$\Delta\theta_{\rm res} = \frac{f}{N_\theta}\left(\frac{r_{\rm in}}{r} + \chi\right), \tag{28}$$

where f controls the factor by which we oversample to reduce noise. Usually $f = 1$ is adequate, however even with this relatively low resolution, the number of tiles varies from $\sim N_\theta^2$ for the innermost radial grid cells to several orders of magnitude larger for outer cells. The increase in resolution at the outer cells is needed in order to resolve the interior. The average number of rays per cell is thus a factor of 100 or more larger than the $N_A = 80$ rays that we consider with short characteristics. In addition, each ray in LC traverses the entire grid, compared to just one cell with SC. For both reasons, one iteration of LC often takes 10^4 times more computation time than one iteration of SC! On the other hand, one iteration of LC moves the solution much closer to convergence since information is propagated across the entire grid.

The main weakness of LC is its inability to handle optically thick regions in scattering-dominated problems. This is because of the huge number of iterations needed to model properly the diffusion process. Since LC cannot be feasibly run for more than a few dozen iterations, it cannot be used to update radiative information deep inside optically thick objects. Luckily, it is precisely in this optically thick regime where short characteristics is free from the ray-defect problem and can be trusted to produce unbiased results. This motivates our hybrid scheme: first we run many iterations of short characteristics to resolve the complete diffusive process and to correctly obtain the radiation field in the optically thick parts of the problem; second we run a few passes of LC to ensure that the optically thin regions of the radiation field are handled accurately and are devoid of ray-defects.

2.4 Acceleration Schemes

The radiative quantities S and J are inextricably linked. For future reference, we quantify the dependence of J on S by means of a Λ operator, defined such that

$$J_i = \Lambda_i^j S_j, \tag{29}$$

where the sub/superscripts denote grid cells. Each iteration of the radiative solver[2] acts as Λ, describing how J is related to S. However, since S itself has a scattering term, it needs to be updated along with J. Writing this explicitly (c.f. Eq. 6),

$$\begin{aligned} S^{n+1} &= (1-\epsilon)J^n + \epsilon B \\ &= (1-\epsilon)\Lambda[S^n] + \epsilon B. \end{aligned} \tag{30}$$

By switching between solving for J as a function of S (Eq. 29), and for S as a function of J (Eq. 30), we seek to find a

---

[2] In both SC and LC, the mean intensity term J is evaluated by averaging I via Eq.5

stationary self-consistent solution for all our radiative quantities. Unfortunately, simply applying these two transformations one after the other and iterating has well known convergence issues whenever scattering strongly dominates (say with $\epsilon < 10^{-2}$).

To avoid slow convergence, acceleration schemes are often employed. Accelerated lambda iteration (ALI) is one such technique. To understand the idea behind ALI, rewrite Eq. 30 so as to isolate S:

$$S = [1 - (1-\epsilon)\Lambda]^{-1}[\epsilon B]. \tag{31}$$

Formally, one can obtain the complete solution by evaluating the inverse matrix on the right hand side of Eq.31 and calculating S from B. Then evaluating Eq. 29 gives J. However, in 3D, Λ is an enormous matrix and it is impractical to carry out full matrix inversion. Instead, ALI considers an *approximate* lambda operator, typically chosen to be the diagonal component Λ_{ii} (i.e. the contribution to the local J from solely the local S – usually the largest component). Denoting the shift in our old formal solution without acceleration as $\Delta S^{FS} = S_i^{n+1} - S_i^n$, solving the system taking into account Λ yields (see Hubeny 2003 for details):

$$\Delta S_i^{ALI} = \frac{\Delta S^{FS}}{1 - (1-\epsilon_i)\Lambda_{ii}}. \tag{32}$$

To obtain the diagonal of the Λ operator, we consider:

$$\Lambda_{ii} = \frac{J_{ii}}{S_i}, \tag{33}$$

where J_{ii} represents the local contribution to J from the source function within the reference grid cell. This self-illumination contribution to J is the following sum over the N_A ray angles considered:

$$J_{ii} = \frac{1}{N_A}\sum_{i=1}^{N_A} I_i^{local}, \tag{34}$$

where I is evaluated by summing up all local sources within the cell:

$$I_i^{local} = \int_0^{\tau_0} S(\tau')\exp(-\tau')d\tau'. \tag{35}$$

ALI is crucial for scattering-dominated problems whenever $\epsilon < 0.01$ – see convergence tests in §3.2. Even with ALI, we often need to run hundreds of iterations before the solution settles down and converges. This is feasible with SC, which is quite fast, but not practical with the much slower LC. The convergence criterion that we use is:

$$\max\left(\frac{|\Delta S_i|}{S_i}\right) \leqslant \delta_c, \tag{36}$$

where the convergence level δ_c is preset to a desired accuracy level (typically $\delta_c = 10^{-4}$). One caveat with this criterion is that it only indicates the level at which HERO is internally satisfied with the solution (i.e. that HERO converged onto a solution). HERO can still converge onto the wrong solution due to systematic errors caused by the numerical setup (i.e. the assumption of linearity for how S varies with τ, or the assumption that $I(\Omega)$ has linear dependence on Ω across neighboring cells). We find that in practice, these systematic effects lead to the HERO solution being biased at the $\sim 5\%$ level when compared with exact analytical results.

2.5 Raytracing

At the end of the day, we wish to calculate images and compute synthetic spectra from our radiating objects. This is accomplished by means of raytracing from some distant observation plane. We consider a large number of rays arriving at the observer and trace each backwards in time to determine its trajectory and point of origin. Then, using the radiative transfer equation and the source function as calculated via the LC method, we compute the observed intensity of each ray. Finally, we combine the rays to construct the observed image of the disc as well as the spectrum of the source. These synthetic results, which are easily calculated for different viewing angles, can be directly compared to observations of accretion discs.

Photon trajectories are handled by numerically evaluating the null geodesic equation Eq. 19 (a second order PDE in x and p):

$$\frac{d^2 x^\alpha}{d\lambda^2} = -\Gamma^\alpha_{\beta\gamma} p^\beta p^\gamma$$

$$\text{with } p^\alpha = \frac{dx^\alpha}{d\lambda}. \tag{37}$$

The emission seen at the observer plane is obtained by summing up the emissivity along the ray paths (same exercise as Eq. 27). This yields a grid of ray intensities ("image" of the system) as projected on the observation plane. Integrating all the rays generates the final observed spectrum via

$$F_\nu = \int I_\nu cos\theta d\Omega, \tag{38}$$

where $\theta \approx 0$ is the normal incident angle of the ray at the observer plane, I_ν is the intensity at the observer plane, and $d\Omega$ is the angular size of a single pixel at the observer plane.

2.6 Frequency Discretization

HERO is set up to handle frequency dependent opacities. Both the computational cost and memory requirement scale linearly with n_F the number of frequency bins. All radiative quantities (opacities, intensities, source functions) are calculated for the same discrete set of frequencies. The code could be set up to handle group mean opacities and emissivities with their appropriate quadrature weights to handle more complex line emission problems. In problems where there is frequency coupling (i.e. gravitational/doppler shifting) the redistribution of photons is handled by linear interpolation. Our typical choice for the frequency grid is a total range of 6 decades with a resolution of 10 points per decade (60 frequency bins).

2.7 Angular Discretization

The angular setup of our code differs based on the dimensionality of the problem that is being solved.

2.7.1 1D Angles

In 1D, we solve plane parallel slab problems which are actually 3D problems that can be collapsed down to 1D by invoking translational and azimuthal symmetry. We subdivide the 4π steradian sphere into equal solid angle slices in θ; thus, we set our angle spacing $d\theta$ such that $\sin\theta d\theta = 2/N_A$.

2.7.2 2D Angles

In all the 2D test problems described in this paper, the rays are considered to live in a flat 2D space, so we do not consider the 3D solid angle at all. We choose our angle grid so as to cover uniformly the full 2π of the 2D plane. We use simple linear interpolation on 2D angles to handle the mixing of angles when rays travel from one cell to another.

2.7.3 3D Angles

The goal is to subdivide the 4π sphere of solid angle into N_A solid angle wedges. A few common approaches include bisecting octants, following a "spiral" that winds around the sphere from the two poles, or a special grid of latitudes/longitudes. All these methods approximately achieve the required goal, but they suffer from undesirable patterns of symmetry lines or "seams" on the surface of the sphere. To sidestep this issue, we obtain our angular grid via a numerical approach.

The strategy that we employ is to deposit N_A fictitious charged particles constrained to move on the surface of a sphere, with initial positions set by simple heuristics (equal proper length spacings in θ, ϕ coordinates). Then, we allow the particles to evolve over time under their mutual inter-particle electrostatic forces until they settle on an equilibrium configuration. This naturally produces a set of positions that are nearly equidistant from one another. We use these positions to specify the angular grid (see Fig 3 for an example with $N_A = 80$).

For relativistic problems where light bending introduces mixing between different angle bins, or for curvilinear coordinate grid setups that invoke angle grids fixed along the locally rotated unit direction vectors, we handle angular interpolation via linear combinations. For a ray with a given unit direction vector $\hat{\psi}$, there will always by a set of 3 angular points (ψ_1, ψ_2, ψ_3) on the 4π sphere which defines a triangle enclosing $\hat{\psi}$. We find the angles (ψ_1, ψ_2, ψ_3) by simply locating the 3 angles with the largest dot product to $\hat{\psi}$ in the angle grid.

Once the three vertices of the triangle have been identified, we decompose $\hat{\psi}$ as a linear combination of ψ_1, ψ_2, ψ_3:

$$\hat{\psi} = c_1\psi_1 + c_2\psi_2 + c_3\psi_3. \tag{39}$$

The coefficients c_1, c_2, c_3 (after renormalizing to satisfy $c_1 + c_2 + c_3 = 1$) are then used as the interpolation weights. In rare cases (i.e. when ψ is slightly outside the boundary of the triangle formed by ψ_1, ψ_2, ψ_3), the interpolation weights may be slightly negative. For stability purposes, we clip any negative weights to zero.

In Fig. 3, we show how our linear interpolation scheme handles an example intensity pattern. Notice that in the slowly-varying equatorial regions, the interpolated result provides a much better match to the exact intensity distribution than simply using a discretized version of the intensity field. The linear reconstruction does poorly when the radiation field fluctuates on angular scales smaller than a single beamwidth of the $N_A = 80$ grid (e.g. see polar regions of Fig. 3).
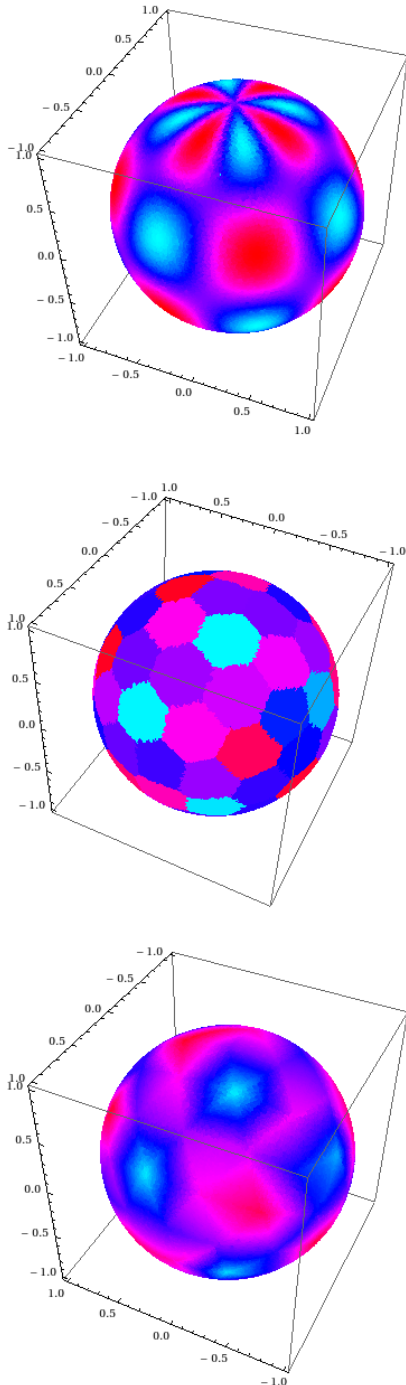
Figure 3. A plot of how our interpolation scheme performs compared to the exact solution. Top: Exact $I = \cos(4\theta)\cos(4\phi)$ beam pattern, Middle: discretized version of top panel using our $N_A = 80$ predefined grid, Bottom: I generated from linear interpolation of the $N_A = 80$ grid.

3 NUMERICAL TESTS

We wish to verify the following properties of HERO: 1) that it correctly solves for the radiation field using both short and long characteristics, 2) that it calculates a self consistent gas temperature taking into account the condition of radiative equilibrium, and 3) that it correctly treats the curved spacetime of GR.

In all the following examples, we run HERO decoupled from any hydrodynamics. HERO is a postprocessor to compute the radiation field I_ν and, if required, a self-consistent temperature T_{gas}, given the optical properties (opacities, emission, absorption) of the medium. Other properties of the background fluid, specifically density and velocity, are assumed to be given and fixed. It is also assumed that the system is time-independent, so we effectively solve for the steady state solution for the radiation.

We are particularly interested in problems where scattering dominates the opacity (as is often true with relativistic accretion flows), leading to diluted-blackbody radiation fields. Due to the complicated nature of scattering, many iterations are needed for the radiative solver to converge. This makes scattering problems ideally suited to test the convergence properties of the code. We describe below a series of 1D, 2D and 3D tests.

3.1 1D Plane-parallel Grey Atmosphere

Scattering problems are in general very difficult to solve due to the integro-differential nature of the scattering kernel in the radiative transfer equation (see righthand side of Eq. 1). As such, there are very few examples that have closed form analytic solutions.

The 1D plane-parallel isothermal slab with grey opacity turns out to be one particularly simple case with a well known closed form solution in the two-stream limit (c.f. §1 Rybicki & Lightman 1979). There is also a semi analytic solution in the limit of infinite angles. For this problem, all quantities, T, κ, σ, are uniform within the medium, leading to a solution that is only a function of the optical depth: $d\tau = (\kappa + \sigma)dz$. Under the two-stream approximation, and using the Eddington approximation to close the radiative moment equations, the analytic solution for the mean intensity is given by

$$J(\tau) = B\left[1 - \frac{\exp(-\sqrt{3\epsilon}\tau)}{1 + \sqrt{\epsilon}}\right], \qquad (40)$$

where B is the thermal blackbody source function, τ is the total optical depth from the surface, and ϵ is the photon interaction destruction probability given by Eq.7.

Note the exponential transition in τ that separates a thermal interior from a dilute-blackbody surface layer. In Figure 4, we show solutions computed by HERO compared to Eq. 40. The $\tau = 0$ surface outer boundary condition was set to have zero ingoing flux, and the innermost $\tau = 10^5$ boundary was set to $J = B$ for some fixed B. The agreement between the numerical and analytical solutions is quite good ($\sim 1\%$ errors) with a systematic trend of larger errors in low ϵ systems. These low ϵ atmospheres are highly scattering dominated and the induced systematic error is simply a consequence of the nonlocal nature of scattered radiation.

3.2 Convergence Tests

The previous plane-parallel atmosphere calculations for the high scattering cases ($\epsilon < 10^{-3}$) required ALI to converge. In Figure 5, We show the convergence properties of the code for a few example problems to emphasize the importance
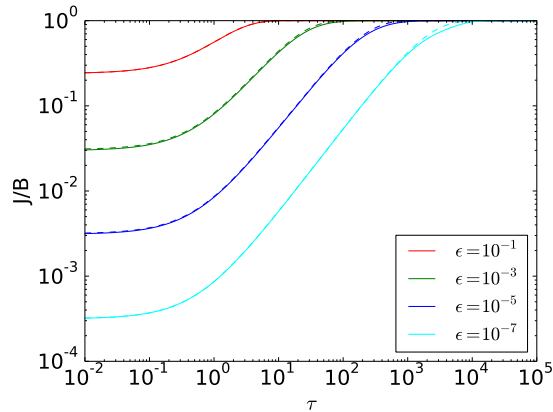
Figure 4. J computed from HERO (dashed lines) compared with the analytic solution (solid lines) for a plane parallel scattering atmosphere described under the 2-stream approximation. Four values of ϵ (see Eq. 7) are shown.
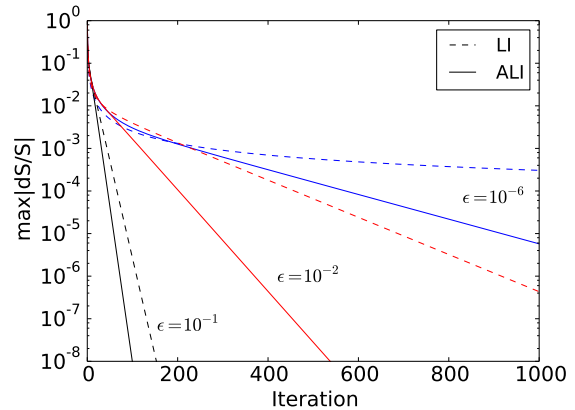

Figure 5. Convergence rates for highly scattering plane parallel atmospheres. Note that the standard Lambda Iteration (LI) procedure saturates at a fixed error threshold for small values of ϵ (when $\epsilon \ll 10^{-2}$). Accelerated Lambda Iteration (ALI) using Eq. 32 converges for all values of ϵ, even extremely small values.

of acceleration in obtaining the correct solution within a reasonable number of iterations.

In the absence of acceleration, radiation information can only propagate a distance $\delta\tau \sim 1$ per iteration. Thus, the number of iterations needed for the radiation to pass from one end to the other is $\sim \tau_{tot}^2$ iterations (corresponding to random walk diffusion with stepsize $\delta\tau = 1$).

The power of ALI rests in the fact that it boosts the single iteration range of influence from $\delta\tau = 1$ to $\delta\tau = \delta\tau_{cell}$. To see why this occurs, consider a tiny perturbation of the source function ΔS for some cell. The effect of ΔS on a neighbouring cell is attenuated by an exponential optical depth factor $\exp(-\Delta\tau)$, where $\Delta\tau$ represents the inter-cell optical depth. The ALI boost factor (c.f. Eq. 32 when $\epsilon \to 1$) in the scattering dominated limit is $\Delta S_{\mathrm{ALI}} = \Delta S \exp(\Delta\tau)$, which exactly compensates for the $\exp(-\Delta\tau)$ attenuation. This allows the ΔS information to propagate to the next cell in just a single iteration, hence requiring N_{cell} iterations for information to traverse the medium from end to end. For most scattering problems, this leads to a huge speedup (since $N_{cells} \ll \tau_{tot}^2$).

3.3 Multiray Temperature Solution

Here we describe a test of HERO's ability to calculate equilibrium temperatures. In Figure 6, we compute a constant flux nonscattering atmosphere and compare the numerical temperature profile with the analytic result (from a table of the Hopf function, see Chandrasekhar 1950). With sufficient angular resolution ($N_A > 16$), the temperature profile computed with HERO is correct to within 1%.

3.4 Test of Spectral Hardening

As a test of HERO's multifrequency capabilities, we now examine a frequency dependent 1D atmosphere. This problem is simple enough to admit an exact analytic solution for the spectrum at all depths within the atmosphere. Appendix B derives the depth dependent spectrum.

We are particularly interested in examining how a strongly scattering atmosphere modifies the thermal emission from deep within a plane parallel atmosphere. This phenomenon is known as spectral hardening, and acts to shift the apparent colour of the photosphere emission. The effect is particularly important in the context of modelling black hole accretion disc spectra (Shimura & Takahara 1995; Davis & Hubeny 2006), where the high degree of scattering leads to a shift in the thermal emission peak by a factor of $\sim 1.5 - 2$ towards higher energies.

Figure 7 demonstrates an extreme version of this spectral hardening effect for an $\epsilon = 10^{-6}$ plane-parallel medium. HERO produces spectra that agree to within 5% of the analytic solution for all energies and optical depths. Note that above the photosphere ($\tau < \tau_{\mathrm{eff}} = 1/\sqrt{3\epsilon} \approx 10^3$), the radiation spectrum differs quite substantially from the local thermal emissivity/temperature (dotted curves).

For simplicity, the example shown here is with the two-stream approximation for grey (i.e. frequency independent) homogeneous opacities. We impose a constant flux boundary condition at $\tau = 10^5$ and solve numerically for the equilibrium temperature and frequency-dependent radiation intensity as a function of depth. There is excellent agreement between the code output and the analytic temperature and radiation profiles (Eq. B13 in Appendix B).

3.5 Effect of a Heating Source

As a final accretion disc-like 1D test that utilizes most of the features of the code, we discuss a problem that includes a heating source. We consider a slab with a total optical depth of 2×10^5 between two free surfaces. We assume grey opacities with $\epsilon = 10^{-4}$ (a strongly scattering-dominated
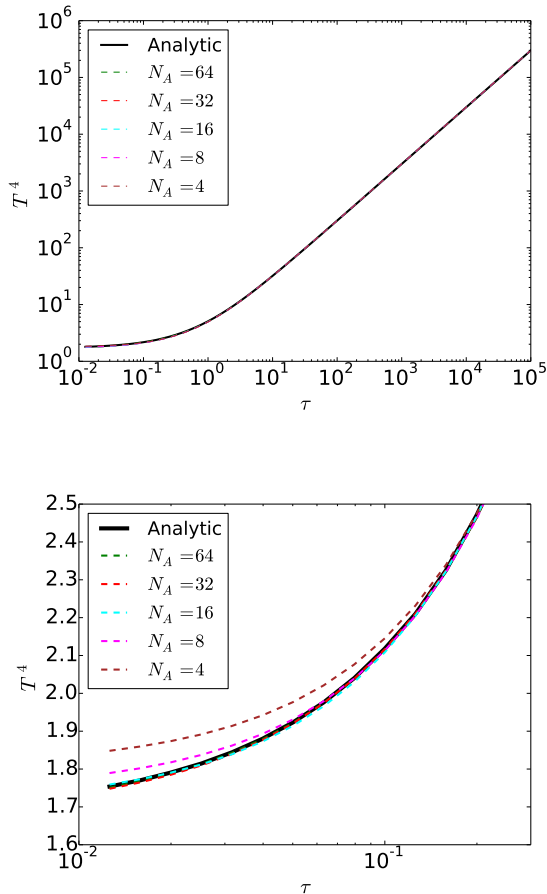
Figure 6. Comparison of analytic solution (solid line) to numerical results (dashed lines) for $\epsilon = 1$ constant flux atmospheres modeled with different numbers of angles. The bottom panel is a zoomed in version of the top panel. The temperature solution in the interior ($\tau \gg 1$) matches perfectly to the analytic solution. The agreement at the surface improves as the number of rays is increased.

disc). The material in the slab is heated at a steady rate of 10^{-5} (arbitrary units) per unit optical depth, so that the steady state flux from each surface is unity. We use 20 angles and 61 frequencies.

We start the calculation with some arbitrary temperature profile ($B = 10^4$ at all τ in this particular example) and we run HERO until it converges. After 3000 iterations (this is many more than needed, but we wished to converge as much as possible for this test), the radiative transfer equation has a maximum error of 10^{-6} and the radiative equilibrium equation a maximum error of 10^{-7}. The results are shown in Fig. 8. As expected, the temperature profile is quadratic in τ and symmetric with respect to the mid-plane (upper panel) and the radiation flux is linear in τ and antisymmetric (middle panel). The lower panel shows spectra at various depths. In the deep interior of the disc, the radiation is nearly blackbody, but as we approach the surface ($\tau_{\mathrm{eff}} = \tau\sqrt{3\epsilon} < 1$) there are signatures of spectral hardening. The radiation that escapes from the surface is distinctly hardened.
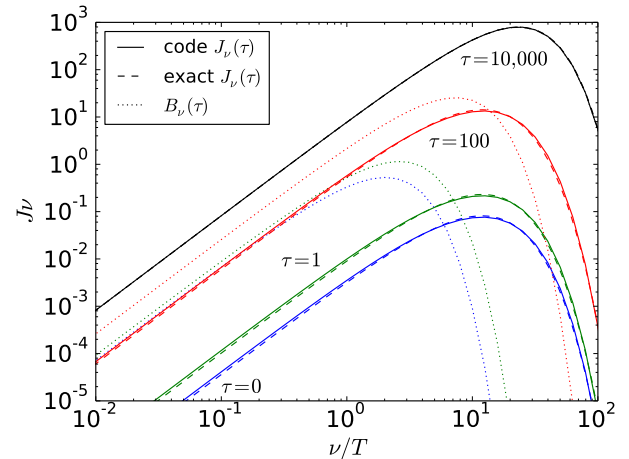
Figure 7. Comparison of depth dependent spectra calculated using HERO (solid lines) and the analytic solution (dashed lines) for a plane parallel thermal atmosphere. Note: this is a *highly* scattering test problem with $\epsilon = 10^{-6}$, which results in significant spectral hardening at the $\tau = 0$ surface (compare dotted blue blackbody curve with solid blue line). At all depths and frequencies, HERO is able to calculate the spectrum to within 5% of the true solution. The spatial resolution chosen for this calculation was 25 points per decade in τ, and 20 points per decade in ν. The two-stream approximation was used for handling angles ($N_A = 2$).

3.6   2D Solutions and Ray Defects

A well known limitation of radiative solvers with discretized angular zones is the appearance of "ray defects". These are sharp linear features that arise because interpolation on the angular grid is imperfect. The presence of ray defects is an important motivator for us to use a hybrid scheme. Our long characteristics method operates independently of an angular grid since it adaptively chooses an angular grid and hence does not suffer from ray-defects. In the following sections, we compare our two solvers, SC and LC, for a few test 2D problems. Appendix A provides a more extensive discussion of the ray defects.

3.6.1   *Opaque Wall Test*

Ray-defects are particularly severe when there is a compact source of radiation (see Fig. A1 in Appendix A). But there are noticeable effects even in the case of smooth extended sources. As an illustration, we analyze a simple 2D example where it is easy to compute the exact solution. In the top panel of Figure 9, we show the radiation field for an empty (zero opacity) box illuminated by a hot spot in the centre of the upper wall. The boundary condition at this wall is an opaque ($\tau = \infty$) isotropic source function with a gaussian profile:

$$\frac{S(x)}{S(x_c)} = \exp\left[-\frac{(x - x_c)^2}{w}\right], \qquad (41)$$

where x_c corresponds to the centre of the box, and the width w is equal to $1/5$ of the box size. Since the interior of the box has no opacity, the radiation field at any point can be eas-
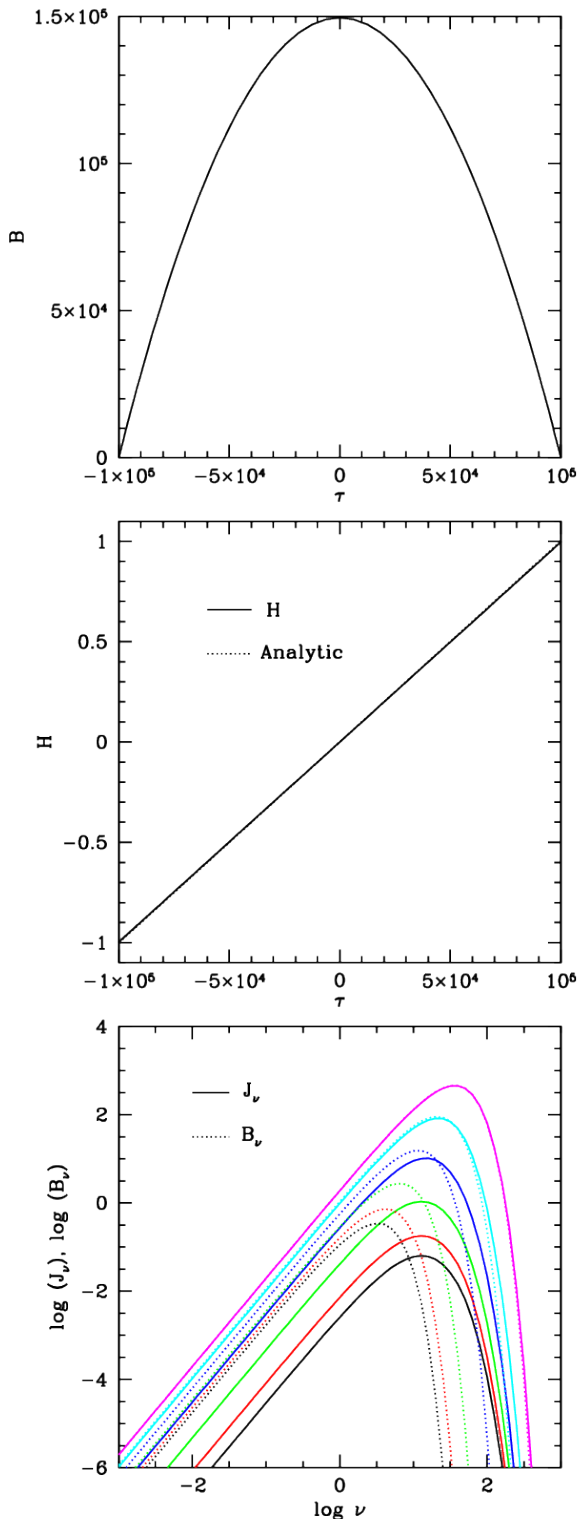
Figure 8. Top: Self-consistent solution with HERO for the temperature in the interior of a uniformly heated 1D slab. Middle: Numerical solution for the frequency integrated radiative flux H compared with the analytic solution. Bottom: Radiation spectrum J_ν compared with the blackbody spectrum at the local temperature B_ν for optical depths (from below) 0, 1, 10, 10^2, 10^3, 10^4 from the surface.

ily found by tracing rays backwards and finding the source function corresponding to the ray's intersection point with the upper wall. We can thus calculate the mean intensity J at every point inside the box.

In Figure 9, we compare the exact solution to the result from our two RT solvers. For the SC solver (middle panel), a clear wave pattern appears along each of the angle grid directions (we used $N_A = 20$ in this test). These are the ray defects mentioned earlier. The severity of the defects grows as we approach the upper wall, because of the sharp break in the intensity distribution there. On the other hand, LC perfectly reproduces the radiation pattern, even in the far field limit (compare top and bottom panels of 9).

3.6.2 Shadowing Test

Shadowing is another classic test that is useful for spotting systematic biases in a radiative solver. Simple moment closure schemes such as FLD (Levermore & Pomraning 1981) or M1 (Levermore 1984) typically have problems resolving the correct shadow structure. For instance, FLD has trouble maintaining the coherency of shadows over long distances due to its diffusive nature, and M1 has issues dealing with multiple light sources.

In Figure 10, we consider the shadow structure produced by an optically thick square box illuminated from above by an isotropically emitting wall. The top panel shows the true solution, with a clear umbra and penumbra. The middle panel shows the result using HERO SC with a crude angular grid $N_A = 16$. Note that the beam resolution of 20^o is our typical angular resolution in 3D problems (80 rays in 3D), so this is a realistic example of how SC would perform in 3D. Increasing the angular resolution makes the ray defects less pronounced – the top "exact" solution was produced with $N_A = 1000$ using our SC solver.

The bottom panel shows the result with the LC solver. Notice the high accuracy of LC in handling the radiation shadow pattern. This is a common theme in all of our tests. LC is always much superior to SC.

3.7 3D Solutions

We now shift our attention to 3D test problems, and begin by discussing ray-defects. All the ray defects discussed in 2D (previous subsection and also Appendix A) are present in 3D as well, especially in the case of cartesian grids. A spherical polar grid (the most natural choice for accretion problems) eliminates some problems by introducing ray mixing via the curvilinear nature of the coordinate system. However, this is at the expense of introducing a particularly serious defect for radial rays moving out from a central source. Specifically, if one applies the SC method blindly on a spherical grid, one will obtain a constant radiation energy density and flux at large radius instead of the inverse square law fall-off one expects.

Dullemond & Turolla (2000) discuss a way to "fix" the radial beam problem. They slightly modify the radiative transfer equation along the radial direction such that the expected inverse-square falloff is recovered. We build on their suggestion, except that, instead of modifying only one ray (the radial one), we treat all rays equally and apply an artificial diffusion that, when coupled with a logarithmic radial
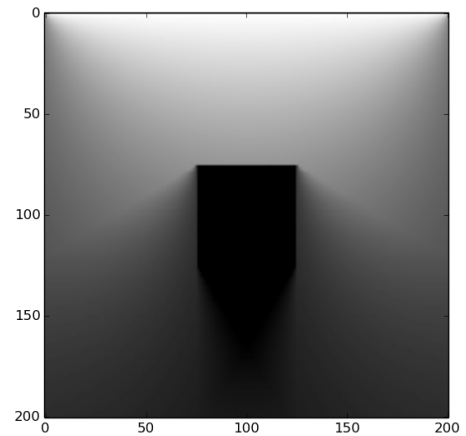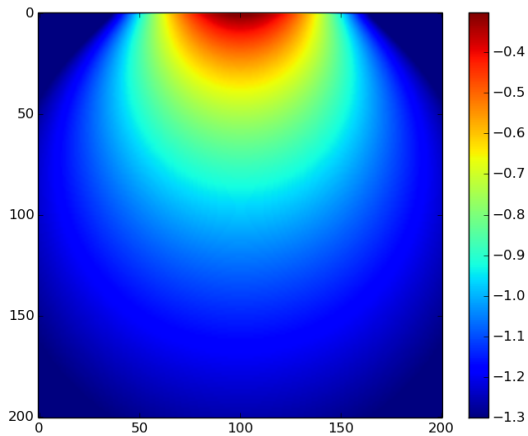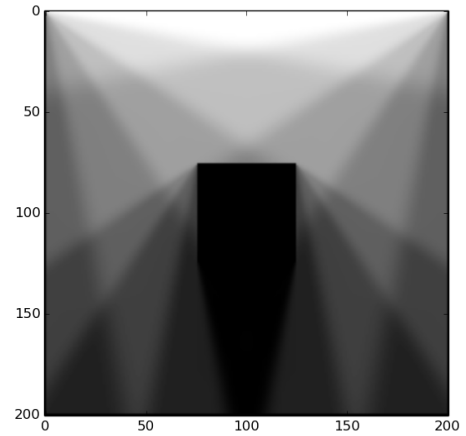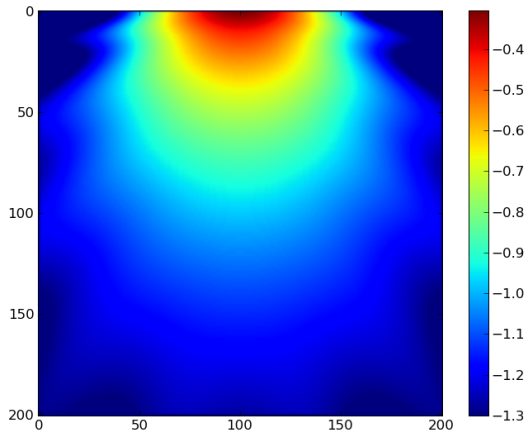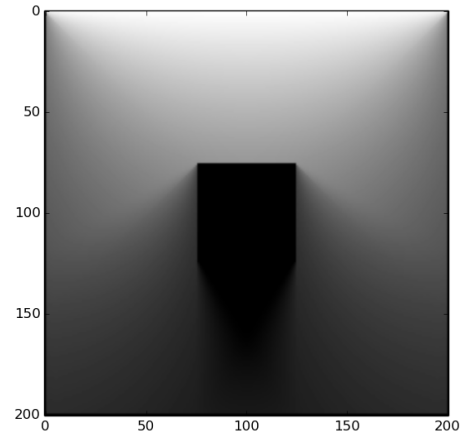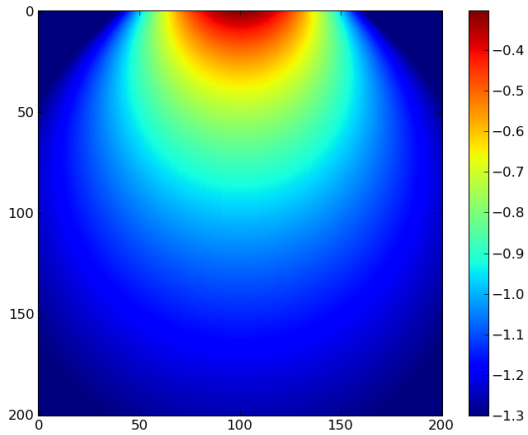
Figure 9. Comparison of radiation fields calculated using short and long characteristics with an exact analytic solution for an opaque gaussian wall emitting into vacuum (colour $\log J$). From top to bottom: a) exact solution, b) SC solution ($N_A = 20$), c) LC solution. Note the appearance of systematic banding/waves in the SC case. These are ray defects.

Figure 10. Comparison of shadow patterns calculated using short and long characteristics for an isotropically emitting top wall shining on a central opaque box. From top to bottom: a) "exact" solution ($N_A = 1000$), b) SC solution ($N_A = 16$), c) LC solution. Note the discrete levels (ray defects) in the SC shadow pattern.

grid, naturally produces an inverse square falloff in the flux. Details of our diffusion method are explained in Appendix §A2. The following tests as well as those in §3.8 employ this ray diffusion scheme.

3.7.1 Ring Benchmark

One particularly simple test problem is an axisymmetric opaque emitting ring in vacuum, for which it is straightforward to calculate the radiation field at any point in space analytically. For an infinitesimally thin ring emitting at the equatorial plane, the total light reaching any position $\vec{r}$ is given by a 1-dimensional integral.

$$J(\vec{r}) = \int \frac{C}{|\vec{r} - \vec{r}_{\text{ring}}(\phi')|^2} d\phi', \qquad (42)$$

where C is a constant specifying the emission per unit length of the ring. In spherical coordinates, we have

$$J(r, \theta, \phi) = \int \frac{C}{r^2 + r'^2 - 2rr' \sin\theta \sin\theta' \cos(\phi - \phi')} d\phi', \qquad (43)$$

where $r' = 0.5$, $\theta' = \pi/2$, $\phi' \in [0, 2\pi]$ are the coordinates of the ring in the setup described here.

In Figure 11, we compare the results of SC and LC to the analytic result; the constant C has been appropriately normalized to account for the finite emitting area used in the LC/SC calculations. The LC calculation captures the radiation field perfectly ($\delta J/J \approx 10^{-3}$), whereas the SC result shows strong systematic ray-defect patterns. Note that to enforce a $1/r^2$ falloff of the radiation field, the SC calculation employs our ray diffusion scheme (otherwise, the result would be significantly worse).

3.7.2 Dusty Torus Benchmark

Previously, in §3.1, we demonstrated that HERO correctly computes both the radiation field and the gas temperature in 1D. Unfortunately, there are no analytic solutions available for nontrivial 3D problems. Therefore, we turn to a standard benchmark problem that has been widely discussed in the literature and use this problem to numerically compare HERO with other radiative codes. The model in question consists of an axisymmetric dusty torus Pascucci et al. (2004) with density structure given by:

$$\rho(r, z) = \rho_0 \cdot f_1(r) \cdot f_2(z) \qquad (44)$$
$$f_1(r) = (r/r_d)^{-1}$$
$$f_2(z) = \exp\{-\pi/4[z/h(r)]^2\}$$
$$h(r) = z_d (r/r_d)^{9/8}$$

The opacities are tabulated and correspond to $0.12\mu m$ silicate grains (Draine & Lee 1984). Scattering is assumed to be isotropic, dominating in the wavelength range $0.2 - 1.0\mu m$. A stellar point source is located at the centre of the disc. This point source shines on the disc and dictates its energetics and temperature structure. In HERO, the radiation emanating from the central point source is set to the exact stellar solution at the innermost radial cells of the grid. The propagation of this radiation outwards is then handled by the short (or long) characteristics solver. To avoid being killed by severe ray-defects, we include the diffusive term in

the short characteristics solver as described earlier (see §A2). We apply free outflowing radiation boundary conditions at the outer radius of the grid.

Given the above boundary conditions, HERO solves for the radiation field everywhere, both inside and outside the disc, as well as the self consistent disc temperature, i.e. the temperature that satisfies Eq. 26. To benchmark the code, we consider the most difficult example presented in Pascucci et al. 2004, the case corresponding to $\tau = 100$. The upper panels of Figure 12 compare the disc midplane temperatures computed by HERO with the benchmark models. Overall, the agreement is reasonable – the slight differences are likely due to ray defects. We recover the temperature structure to within 10%, with the worst cells being located in the optically thick disc midplane.

The equilibrium temperature profile is highly sensitive to the amount of scattered light (which dominates over the direct stellar illumination near the disc surface by an order of magnitude). The good agreement between HERO and Pascucci et al. (2004) indicates that : 1) HERO correctly handles/redistributes the scattered light; 2) the temperature solver is robust.

As part of this test, we also show in Fig. 12 how the short and long characteristics version of HERO perform in determining the self-consistent disc temperature. Short characteristics does a reasonable job everywhere within the optically thick parts of the disc, but systematically underestimates the radiation field and temperature in the optically thin regions. Generally, SC has difficulties propagating radiation towards the coordinate poles. Panel 2 of Figure 12 shows the resultant underestimated radiation temperatures ($\sim 10\%$ error) and that the bias increases with increasing θ resolution.

The long characteristics solver does not suffer from this systematic error and recovers the correct temperature profile within tens of iterations throughout the optically thin region. This confirms that it is generally a good idea to run a few iterations of long characteristics after the short characteristics solver has been run to convergence. We make a special point to emphasize the importance of the LC pass. One might be tempted to instead run a high resolution SC pass to pin down a more accurate radiative solution. However we find that high resolution does not reduce the systematic biases that plague the SC method. Despite the computational expense incurred by the LC method, it is the only way to obtain an accurate solution to the radiation field.

Finally, we also show in the lower two panels of Fig. 12 integrated disc spectra. Agreement between HERO and the other benchmark codes is within 20%, which is comparable to the spread amongst the four independent codes discussed in Pascucci et al. (2004). The HERO spectra were computed by raytracing from a distant observing plane, making use of the complete radiative solution (i.e. S_ν) obtained from our SC+LC hybrid solver.

3.8 GR Solutions

3.8.1 Light Bending

HERO is designed to solve for the radiation field in general relativistic curved spacetimes. One important effect is light bending which is demonstrated in the test problem shown
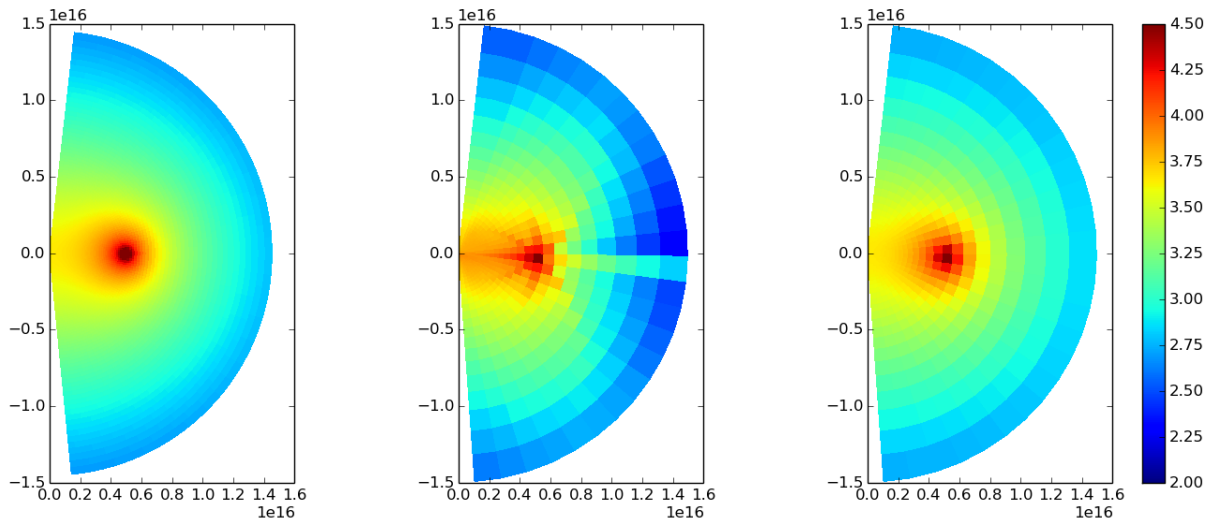
Figure 11. Comparison of true solution (left) with the results obtained with SC (middle) and LC (right) for an axisymmetric emitting ring (colours represent $\log J$). Note the effect of ray defects producing a "spider" pattern for SC. The LC calculation matches the exact analytic answer to within 0.1%.

in Figure 13. J is computed with both the SC and LC in the 3 panels for a beamed light source located just outside the photon orbit ($r = 3M$) propagating in Schwarzschild space-time. We see that the regions with the highest intensity follow the expected curved trajectory. However, the beam is broadened substantially. This is a consequence of the finite angular grid (number of angles $N_A = 80$) used in the computation.

The middle and right panels of Figure 13 show LC solutions to the same problem. The LC beam remains narrow and coherent, agreeing very well with the expected behaviour for free-streaming radiation at the photon orbit. The middle panel shows the LC result using an emitting source whose beam has been artificially broadened to match the angular resolution of the SC $N_A = 80$ angular grid. The right panel shows the true resolving power of LC, where the emitting source corresponds to a δ-function in angle.

3.8.2   Disc Spectra

To check that our handling of Doppler and gravitational redshifting is correct, we consider the problem of black hole accretion disc spectra. This problem has been tackled numerous times over the years using many independent codes (e.g. Li et al. 2005; Davis & Hubeny 2006; Kulkarni et al. 2011; Zhu et al. 2012) and is a simple but useful benchmark test.

We place an optically thick, geometrically thin disc around a Schwarzschild ($a = 0$) black hole radiating as per the idealized Novikov & Thorne (1973) disc model. In this problem, the disk emission is treated as isotropic thermal radiation with flux given by the NT model. For simplicity, we ignore any spectral hardening effects since we treat the emission in HERO as emanating from a single equatorial grid cell (i.e. for this test, we do not resolve the photon diffusion process that gives rise to spectral hardening). The

inner edge of the disk is fixed at the ISCO $r_{\mathrm{ISCO}} = 6$ and the outer edge is located at $r = 1000$.

We solve for the radiation field in the disc exterior using HERO in full GR. In HERO, the calculation of the disc spectrum is typically handled in two stages – 1) we first solve for the 3D radiation field above the disc using the short/long characteristics solver (this is not needed in the present example because we specify the disc emission profile and assume vacuum outside the disc), and 2) we trace rays backwards (via Eq. 22) from a distant observer plane located to create a synthetic image of the disc.

The HERO code solves for the full three-dimensional source function and radiation field within the $2M < r < 1000M$ spatial domain of the grid. This information is then fed into a separate raytracing subroutine. In the raytracing stage, parallel rays are shot towards the disk distributed according to a squeezed logarithmic polar grid (the same setup as Kulkarni et al. 2011) from an observer plane located at $r = 100,000M$. The final spectrum is generated by integrating the flux across the observer plane.

The top panel of Figure 14) shows the computed image for the particular example problem. Integrating over this image yields the observed disc spectrum, which we show in the lower panel. The disc spectrum computed with HERO agrees very well with a previous calculation by Kulkarni et al. (2011), lending confidence that HERO correctly handles GR effects.

3.8.3   Iron Line Spectra

In a similar vein as for the previous test, which was based on the NT continuum spectrum of the disk, we also benchmark our code via a calculation of the emission profile due to a monoenergetic line (e.g. Fe-Kα as seen in many Seyfert galaxies and some microquasars). For a system that emits only monoenergetic δ-function lines, the final integrated line
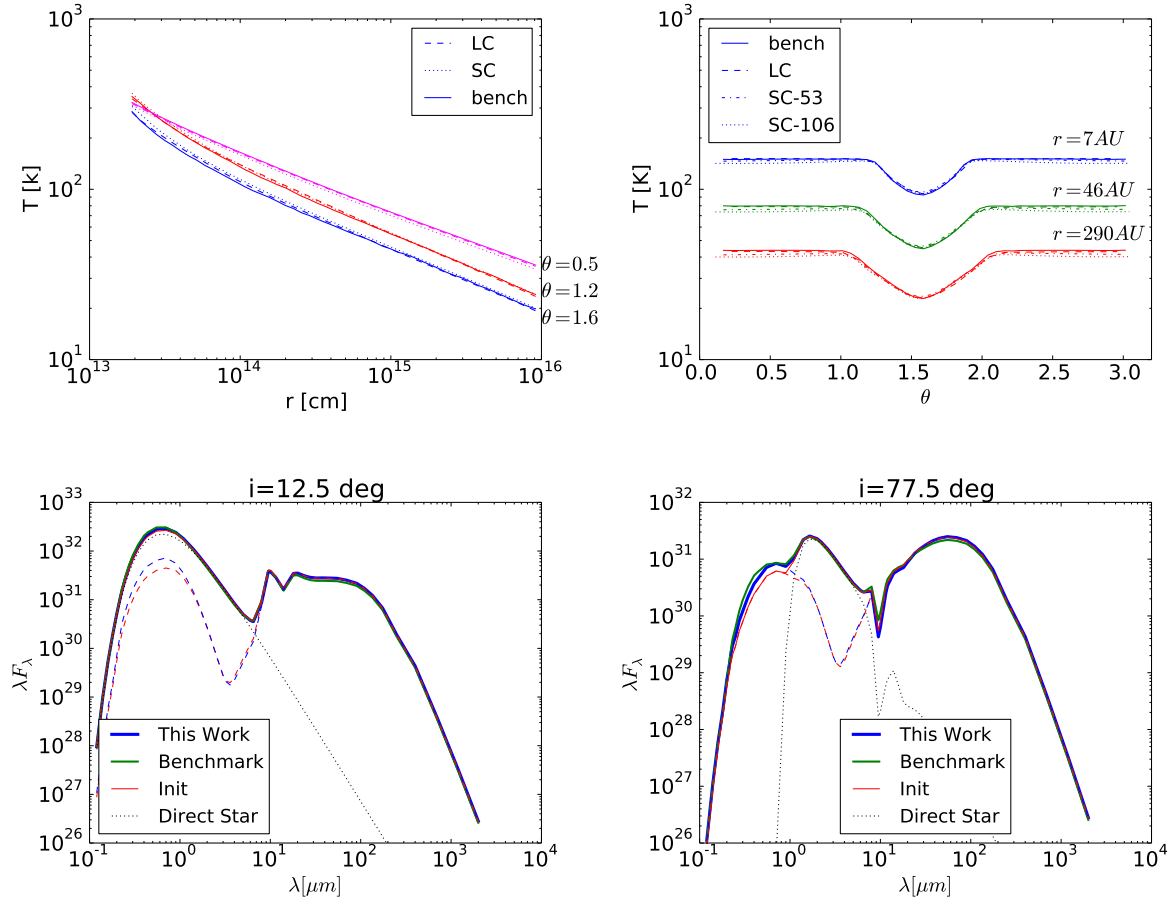
Figure 12. Comparison of HERO with the most difficult $\tau = 100$ dusty torus benchmark test of Pascucci et al. 2004. The top two panels show slices of the self-consistent temperature solution (top left: radial slices, top right: poloidal slices). We also show the SC solution evaluated for two different choices of grid resolution: $N_\theta = (53, 106)$. We find that the SC solution systematically underestimates the radiation/temperature field near the polar regions, with the effect becoming enhanced at high resolutions. This is a consequence of the angular diffusion scheme used in SC (see Appendix A2 for more discussion). The bottom two panels compare spectra at different inclination angles as computed from LC raytracing.
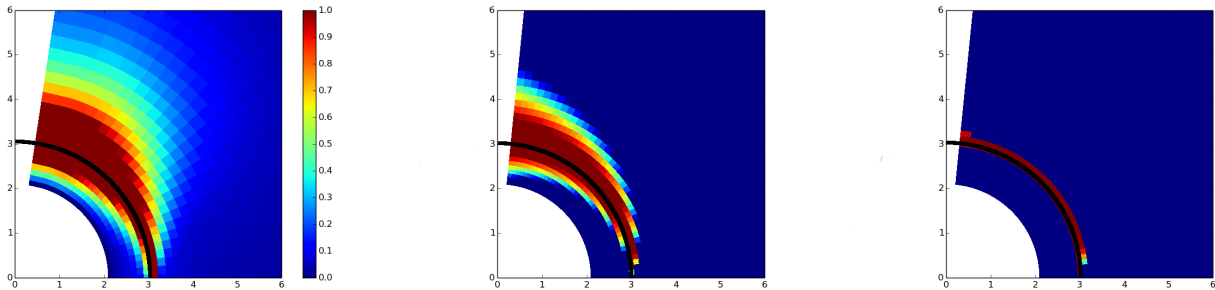
Figure 13. Light bending test for our two radiative solvers using a narrow laser beam injected tangent to the $r = 3$ photon orbit. From left to right: a) SC, b) LC with an artificially broadened beam to match the SC angle grid size, c) LC pure. The solid black line shows the analytic result corresponding to the null geodesic.
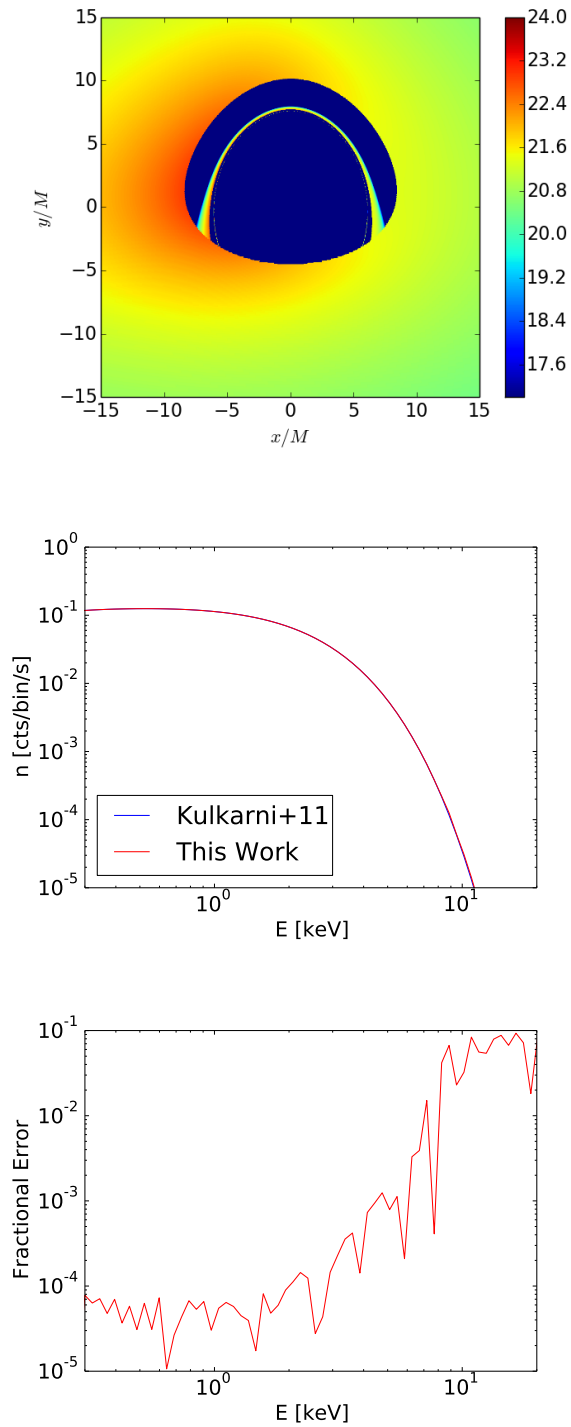
Figure 14. Upper panel: Raytraced image of a razor thin disc as viewed by an observer located with inclination angle $i = 60$. Colors indicate log of the frequency integrated intensity. The asymmetry is due to the doppler effect combined with gravitational redshifting. The central arc feature corresponds to a secondary image of the accretion disk that arises from strong gravitational lensing about the black hole. Middle panel: Integrated disc spectrum as computed by HERO (red points), compared to the result with the Kulkarni et al. 2011 code (blue line). Bottom panel: fractional errors between HERO and Kulkarni et al. 2011.
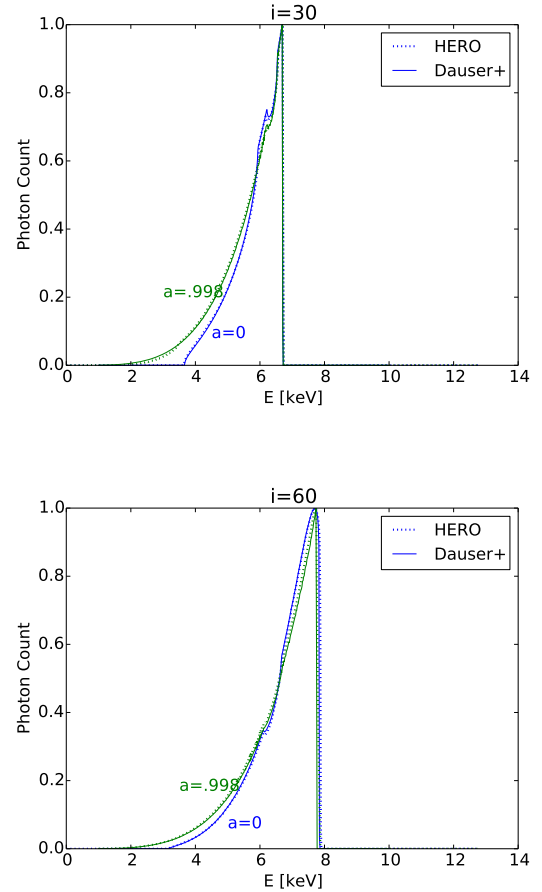

Figure 15. A comparison of the Fe-Kα line profile as computed by HERO and RELLINE (Dauser et al. 2013) for nonspinning (blue, $a/M = 0$) and nearly maximally spinning (green, $a/M = 0.998$) black holes. Line profiles are computed for different viewing inclinations: top panel shows $i = 30$, whereas bottom panel shows $i = 60$.

profile depends only on the geometrical properties of the system (i.e. redshifting from doppler/gravitational curvature and lensing effects from the Kerr metric).

We consider a Keplerian accretion disk in the Kerr metric whose line emission is modulated by a power law emissivity profile $F(r) \propto r^{-3}$. We also set the domain of the disk to span from $r_{in} = r_{ISCO}$ out to $r_{out} = 400$. Figure 15 shows the resultant spectra computed for two different choices of BH spin and two different choices of viewing angle. We find good agreement between HERO and the benchmark code RELLINE (Dauser et al. 2010, 2013), with the largest discrepancies occurring in the low energy red tail of the line profiles.

The setup in HERO is identical to that of the previous NT disc tests, except that a monoenergetic 6.4 keV line was used as the local emissivity instead of a thermal continuum source. In addition, we use much higher resolution for the raytracing grid since the low energy red tail of the line profile depends sensitively on how well the inner edge of the disc is resolved.
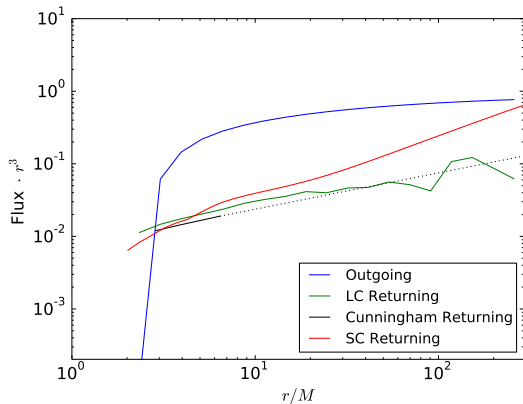
Figure 16. A calculation of the returning radiation for a razor thin accretion disc around a moderately spinning $a_* = 0.9$ black hole. The standard Page & Thorne (1974) luminosity profile is used to set the outgoing flux/rays from the disk plane (blue). We compare the returning flux as calculated by the two radiative solvers in HERO (red, green) with the result from Cunningham (1976) obtained via relativistic transfer functions (black). The black dashed line is an extrapolation of the values from Cunningham's data table.

3.8.4 Returning Radiation

Another classic accretion disk problem is that of computing the returning radiation due to relativistic light bending around the black hole. For this test, we setup a razor thin accretion disk that emits according to the standard thin disc luminosity profile (Page & Thorne 1974) and measure the amount of returning radiation incident on the disk. The goal is to benchmark both our SC and LC radiative solvers against the solutions of Cunningham (1976), who tackled the same problem by means of relativistic transfer functions.

In HERO, we model the returning radiation problem with a grey calculation (1 bolometric frequency bin) on an axisymmetric polar grid with $(n_r, n_\theta) = (60,30)$ restricted to the upper half plane. The spatial grid was set as uniformly spaced in angle and $\log(r)$. Boundary conditions invoked are: reflecting for the polar axis (to account for light that passes through the pole), constant flux injection at the equatorial disk plane according to (Page & Thorne 1974), and zero incident radiation at the inner and outer radial boundaries.

In figure 16, we plot the incoming radiation flux at various locations above the disk and compare to the solution of Cunningham (1976) for a moderately spinning $a_* = 0.9$ black hole. We find that SC systematically overestimates the amount of returning flux at large radii, presumably caused by the angular interpolation bleeding some of the outbound radiation into inbound rays. LC on the other does not experience any strong systematic biases, but suffers from a lack of resolving power at large radii since at these large distances, the LC ray grid can miss the inner photon ring that is responsible for most of the returning radiation.

3.8.5 Vacuum Test

As a final test of the general relativistic capabilities of HERO, we turn to the problem of light propagation in vac-

uum from an opaque spherical shell. In Schwarzchild geometry, it is a simple exercise to solve for the apparent angular size of a constant radius shell as viewed by an observer at some other (larger) radius. Furthermore, if the surface of the shell radiates isotropically like a blackbody with a constant surface temperature, one can calculate the radiation quantities (i.e. J, F, L) by simply integrating the constant intensity across the apparent solid angle subtended by the shell. For instance, the radial profile of luminosity has a particularly simple form, scaling directly with gravitational redshift as

$$L_{\text{local}} = L_\infty (1 + z)^2 \qquad (45)$$

where $L_{\text{local}} = F_{\text{local}} 4\pi r^2$ is the total luminosity as measured by an observer at radius r and L_∞ is the luminosity at infinity. In Figure 17, we show the results for the luminosity as calculated by our general-relativistic short and long characteristics solvers. We normalize the luminosity by the redshift factors so the analytic solution simply corresponds to a flat horizontal line (i.e., constant luminosity as measured at infinity).

In general, we find that our short characteristics solver systematically underestimates the luminosity profile at large distances. These tests were carried out with the diffusion prescription described earlier, therefore if the shell radiated purely in the radial direction, SC would by construction give the correct answer. Here the surface radiates isotropically, and there is a deviation from the true answer because of angle interpolation and diffusion. The various SC curves in Figure 17 correspond to different choices for the radius of the inner emitting shell. The high degree of similarity for all choices of inner radius (i.e. ranging from from highly relativistic $r_0 = 3$ in units of GM/c^2 to nonrelativistic $r_0 = 10^5$) implies that the SC bias is independent of relativistic effects. It arises purely from the discretization of the spatial and angular grid. While the bias is not negligible – the luminosity is reduced by a factor of 2.5 at large radius – note that the luminosity would be a factor $\sim 10^6$ too large if we did not include diffusion. As with the other test problems, LC gives an essentially perfect answer.

Panel two of Figure 17 shows another interesting phenomenon, viz., there is a strong dependence between the SC bias on the choice of spatial grid resolution. As we increase θ-resolution (keeping the r-resolution fixed), the luminosity profile of our isotropically emitting shell exhibits a stronger bias to lower values. This effect is ultimately caused by the angle diffusion scheme that we employ (c.f. Appendix A2). The diffusion is tuned such that the radial cell-to-cell attenuation of light recovers the inverse-square law. If the cell aspect ratio in $r-\theta$ is too rectangular, then the diffusion has a strong directional preference. Light rays traveling along the short dimension of the cell are overattenuated since they hit cell boundaries more often than rays traveling in the long-dimension (the degree of diffusion is directly proportional to how often the light ray traverses cells within a given spatial distance). Based on our tests, grid cells should ideally have a square aspect ratio to minimize the error, and any ratio in excess of 2:1 should be avoided.
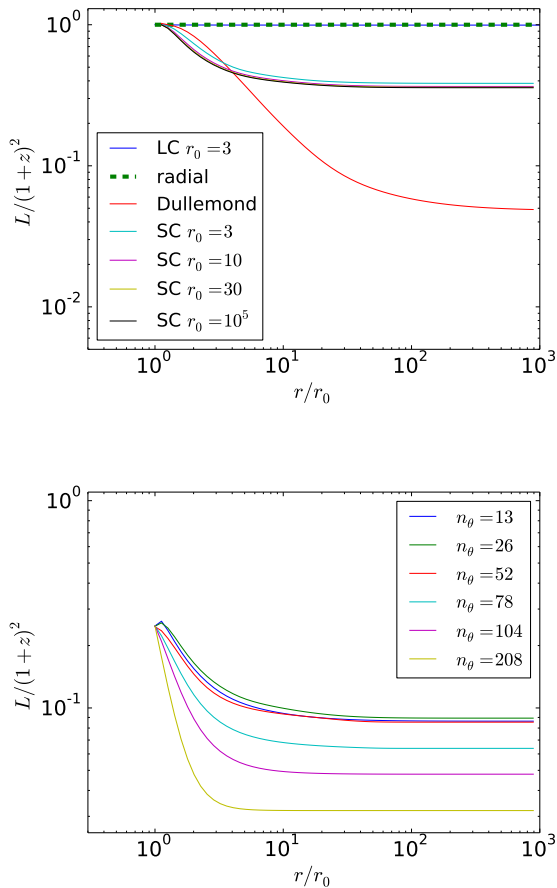
Figure 17. Variation in radial luminosity profiles as a function of a) relativistic effects (lensing, redshifting) and b) grid-θ resolution. Note that the luminosity plotted is the corrected luminosity as measured by an observer at infinity. The analytic solution is therefore a constant $L/(1+z)^2$ for all radii.

4 SUMMARY

We have described in this paper HERO, a new general relativistic radiative transfer code. The primary aim of this code is to model the radiation field in accretion flows around black holes. The unique features of HERO are: 1) a hybrid short/long characteristics radiative solver that enables accurate and fast modelling of complex anisotropic radiation fields; 2) implementation in a general relativistic framework taking into account the effects of light bending, doppler beaming, and gravitational redshifting.

HERO is written as a post-processing code decoupled from the hydrodynamic evolution of the fluid. It computes the time-independent radiation field assuming a given fixed background fluid structure. Strictly speaking, this approach is valid only for problems where the fluid velocities are small compared to the speed of light (i.e. nonrelativistic flows). Alternatively, and this is the primary application we have in mind, it could also be applied to time-steady relativistic problems.

We provide a detailed explanation of the long/short characteristics method used to solve for the radiation field and our approach for solving the self-consistent gas tem-

perature. To verify that HERO produces physically correct answers, we have performed a comprehensive set of tests designed to examine the code's convergence properties, accuracy, and capability to handle multidimensional relativistic problems. We confirm the well known result that 2D and 3D problems with compact sources suffer from significant ray defects in the far field when analyzed with the short characteristics method. We present an approximate fix which mitigates the effects in the case of a 3D spherical grid. However, for accurate results, it is necessary to switch to a long characteristics solver which is unaffected by ray-defects.

As the subject of a follow-up paper, we intend to apply HERO to radiative MHD simulations of accretion discs – particularly simulations undergoing super-Eddington accretion, where radiation feedback strongly dominates the dynamics of the flow. Using HERO, we will investigate the integrated spectra to see the role that self-shadowing and irradiation plays in these systems. This application requires a Comptonization module which will be described in our next follow-up paper.

REFERENCES

Abel, T., & Wandelt, B. D. 2002, MNRAS, 330, L53

Bianchi S., 2008, A&A, 490, 461

Bjorkman J. E., Wood K., 2001, ApJ, 554, 615

Bohn A., Throwe W., Hebert F., Henriksson K., Bunandar D., Taylor N. W., Scheel M. A., 2014, preprint (arXiv:1410.7775)

Bruenn S. W., Dirk C. J., Mezzacappa A., Hayes J. C., Blondin J. M., Hix W. R., Messer O. E. B., 2006, JPhCS, 46, 393

Burrows A., Young T., Pinto P., Eastman R., Thompson T. A., 2000, ApJ, 539, 865

Chakrabarti S., Cox T. J., Hernquist L., Hopkins P. F., Robertson B., Di Matteo T., 2007, ApJ, 658, 840

Chan C.-K., Psaltis D., & Ozel, F., 2013, ApJ, 777, 13

Chandrasekhar S., 1950, Radiative Transfer, Oxford, Clarendon Press

Commerçon B., Teyssier R., Audit E., Hennebelle P., Chabrier G., 2011, A&A, 529A, 35

Cunningham C. T., Bardeen J. M., 1973, ApJ, 183, 237

Cunningham C., 1976, ApJ, 208, 534

Dauser T., Wilms J., Reynolds C.S., Brenneman L.W., 2010, MNRAS, 409, 1534

Dauser T., Garcia J., Wilms J., Böck M., Brenneman L.W., Falanga M., Fukumura K., Reynolds C.S., 2013, MNRAS, 430, 1694

Davis S. W., & Hubeny I. 2006, ApJS, 164, 530

Davis S. W., Stone J. M., & Jiang Y.-F., 2012, ApJS, 199, 9

Dexter J., Agol E., 2009, ApJ, 696, 1616

Dolence J. C., Gammie C. F., Mościbrodzka M., Leung P. K., 2009, ApJS, 184, 387

Draine B. T., & Lee H. M., 1984, ApJ, 285, 89

Dubroca B., & Feugeas J. L. 1999, CRAS, 329, 915

Dullemond C. P., Turolla R., 2000, A&A, 360, 1187

Evans K. F., 1996, Journ. Atmosph. Sci., 55, 429

Feautrier P., 1964, Comptes Rendus Academie des Sciences (serie non speciee), 258, 3189

Ferland G. J., Korista K. T., Verner D. A., Ferguson J. W., Kingdon J. B., Verner E. M., 1998, PASP, 110, 761

Finlator K., Özel F., Davé R., Oppenheimer B. D., 2009, MNRAS, 400, 1049

Ghosh H, 2013, PhD Thesis, Jadavpur University (arXiv.org:0909.0708)

Gittings M., Weaver R., Clover M., Betlach T., et al., 2008, CS&D, 1a, 5005

Gonález, M., Audit E., Huynh P., 2007, A&A, 464, 429

Hanke F., Müller B., Wongwathanarat A., Marek A., Janka H.-T., 2013, ApJ, 770, 66

Hayek, W., Asplund, M., Carlsson, M., Trampedach, R., Collet, R., Gudiksen, B. V., Hansteen, V. H., & Leenaarts, J. 2010, A&A, 517, A49

Hayes, J. C., & Norman, M. L. 2003, ApJS, 147, 197

Hayes J. C., Norman M. L., Fiedler R. A., Bordner J. O., Li P. S., Clark S. E., ud-Doula A., Mac Low M.-M., 2006, ApJS, 165, 188

Hayward C. C., Keres D., Jonsson P., Narayanan D. Cox, T. J., Hernquist L., 2011, ApJ, 743, 159

Heinemann, T., Dobler, W., Nordlund, A., & Brandenburg, A. 2006, A&A, 448, 731

Huang L., Cai M., Shen Z-Q., Yuan F., 2007, MNRAS, 379, 833

Hubeny, I., 2003, ASPC, 288, 17

Indebetouw, R., Whitney, B. A., Johnson, K. E., Wood, K., 2006, ApJ, 636, 362

Jiang Y.-F., Stone J. M., Davis S. W., 2014, ApJS, 213, 7

Kolb S. M., Stute M., Kley W., Mignone A., 2013, A&A, 559A, 80

Krumholz, M. R., Klein, R. I., McKee, C. F., & Bolstad, J. 2007a, ApJ, 667, 626

Kulkarni A. K. et al., 2011, MNRAS, 414, 1183

Kunasz, P., & Auer, L. H. 1988, J. Quant. Spec. Radiat. Transf., 39, 67

Levermore, C. D., & Pomraning, G. C. 1981, ApJ, 248, 321

Levermore, C. D. 1984, Journal of Quantitative Spectroscopy and Radiative Transfer, 31, 149

Li, L.-X., Zimmerman, E. R., Narayan, R., & McClintock, J. E., 2005, ApJS, 157, 335

Liebendörfer M., Messer B., Mezzacappa A., Bruenn S. W., Cardall C. Y., Thielemann F.-K., 2004, ApJS, 150, 263

Lindquist, R. W., 1966, Ann. Phys., 37, 487

Livne E., Burrows A., Walder R., Lichtenstadt I., Thompson T. A., 2004, ApJ, 609, 277

Lopez B., Mekarnia D., Lefevre J., 1995, A&A, 296, 752

McClarren R. G., Holloway J. P., Brunner T. A., 2008, J. Comput. Phys., 227(3), 2864

Mihalas, D., Auer, L. H., & Mihalas, B. R. 1978, ApJ, 220, 1001

Mihalas, D., & Mihalas, B. W. 1984, Foundations of radiation hydrodynamics, ed. Mihalas, D. & Mihalas, B. W.

Niccolini G., Woitke P., Lopez, B., 2003, A&A, 399, 703

Niccolini, G., Alcolea, J., 2006, A&A, 456, 1

Novikov I. D., Thorne K. S., 1973, in Dewitt C., Dewitt B. S., eds, Black Holes (Les Astres Occulus), Gordon and Breach, Paris, p.343

Odaka H., Aharonian F., Watanabe S., Tanaka Y., Khangulyan D., Takahashi T., 2011, ApJ, 740, 103

Ohsuga K., Mineshige S., Mori M., Kato Y., 2009, PASJ, 61, 7O

Olson, G. L., & Kunasz, P. B. 1987, J. Quant. Spec. Radiat. Transf., 38, 325

Özel F., Di Matteo T., 2001, ApJ, 548, 213

Page D. N., Thorne I. D. 1974, ApJ, 191, 499

Pascucci I., Wolf S., Steinacker J., Dullemond C.P., Henning T., Noccolini G., Woitke P., & Lopez B., 2004, A&A, 417, 793

Pelkonen V.-M., Juvela M., Padoan P., 2009, A&A, 502, 833

Pinte C., Ménard F., Duchêne G., Bastien P., 2006, A&A, 459, 797

Rampp M., Janka H. T., 2002, A&A, 396, 361

Rybicki G. B., Lightman A. P., 1979, Radiative Processes in Astrophysics, John Wiley & Sons, New York, p39

Sądowski, A., Narayan, R., Tchekhovskoy, A., & Zhu, Y., 2013, MNRAS, 429, 3533

Schartmann M., Meisenheimer K., Camenzind M., Wolf S., Tristram K. R. W., Henning, T., 2008, A&A, 482, 67

Schechtman-Rook A., Bershady M. A., Wood K., 2012, ApJ, 746, 70

Schnittman J. D., Krolik J. H., 2013, ApJ, 777, 11

Shcherbakov R. V., & Lei H., 2011, MNRAS, 410, 1052

Shimura T., & Takahara F., 1995, ApJ, 445, 780

Stalevski M., Fritz J., Baes M., Nakos T., Popovic L. C., 2012, MNRAS, 420, 2756

Steinacker J., Bacmann A., Henning T., 2002, JQSRT, 75, 765

Steinacker J., Bacmann A., Henning T., Klessen R., Stickel M., 2005, A&A, 434, 167

Stenholm L. G., Stoerzer H., Wehrse R., 1991, JQSRT, 45, 47

Stone, J. M., Mihalas, D., & Norman, M. L. 1992, ApJS, 80, 819

Swesty F. D., Myra E. S., 2009, ApJS, 181, 1

Szu-Cheng S., Kuo-Nan L., 1982, JQSRT, 28, 271

Turner, N. J., & Stone, J. M. 2001, ApJS, 135, 95

Vincent F. H., Paumard T., Gourgoulhon E., Perrin G., 2011, CQGra, 28, 5011

Vögler, A., Shelyag, S., Schussler, M., Cattaneo, F., Emonet, T., & Linde, T. 2005, A&A, 429, 335

Witt, A. N., Gordon, K. D., 1996, ApJ, 463, 681
Woitke P., Kamp I., Thi W.-F., 2009, A&A, 501, 383
Wolf, S., Fischer, O., Pfau, W., 1998, A&A, 340, 103
Wolf S., Henning Th., Stecklum B., 1999, A&A, 349, 839
Zhang, W., Howell, L., Almgren, A., Burrows, A., & Bell, J. 2011, ApJS, 196, 2
Zhu Y., Davis S. W., Narayan R., Kulkarni A. K., Penna R. F., McClintock J. E., 2012, MNRAS, 424, 2504

APPENDIX A: RAY DEFECTS

The method of short characteristics constitutes the primary workhorse of our radiative solver and is used to generate a good first approximation to the radiation field. However, "ray-defects" are a well known limitation of SC, which is why we need to follow up SC with the more accurate LC method. Here we discuss some of the properties and explore the cause of ray defects. Figure A1 shows a few examples of ray defects in 2D arising from point source emitters. The defects manifest as unphysical beam-like patterns far from the emitting source.

Point sources generate the strongest defect pattern so we use them in the following discussion to illustrate the main issues. Point-source-like emission does appear in the problem of black-hole accretion discs, e.g. the hottest innermost region of the disc shines extremely brightly and due to its compact spatial scale, acts like a point source at large distances from the centre.

For a fixed angular grid, the point source ray defects form a series of radial beams that reflect the underlying structure of the angular grid (see Figures A1 for a few examples in different coordinate systems). Beam collimation is enhanced for rays travelling in directions where neighbouring grid cells cover a smaller angular size. This "grid-lattice" effect is best seen in panel c) (sheared box) in Figure A1, where the thinnest beams are those travelling to the upper-right/lower-left sectors (i.e. the directions where the neighbour points as defined in Figure 1 span the smallest angular extent). The same, but less pronounced result is seen in panel d) of Figure A1 (polar coordinate system). Here, the rays pointed radially inward suffer less dispersion than their radially outward counterparts, again for the same reason as in the sheared box case (tighter angular packing occurs for the neighbour cells at smaller radius).

Finally, ray defects are particularly enhanced for rays that directly intersect neighbouring cell centres (as an example, note that the 45^o rays exhibit overwhelmingly strong defect patterns in our cartesian setup in the top two panels of Fig A1).

Ideally, we would like to represent the discretized radiation pattern as a smooth field instead of a superposition of laser beams. After much trial and error, we have arrived at an angular diffusion based solution for spherical log polar grids. To motivate our final solution, we first examine the root cause of the defects.
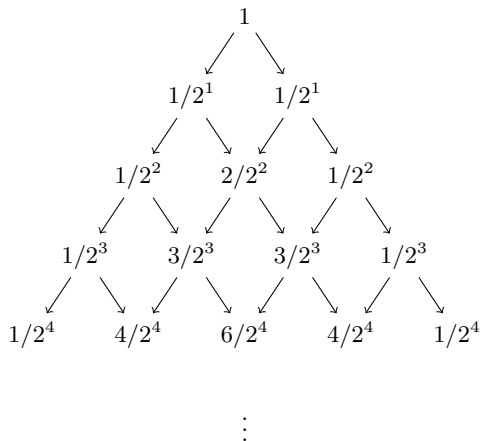
A1   Mathematical Origin of Ray Defects

Ray defects ultimately arise due to the local nature of the short characteristics solver. Methods that operate using just
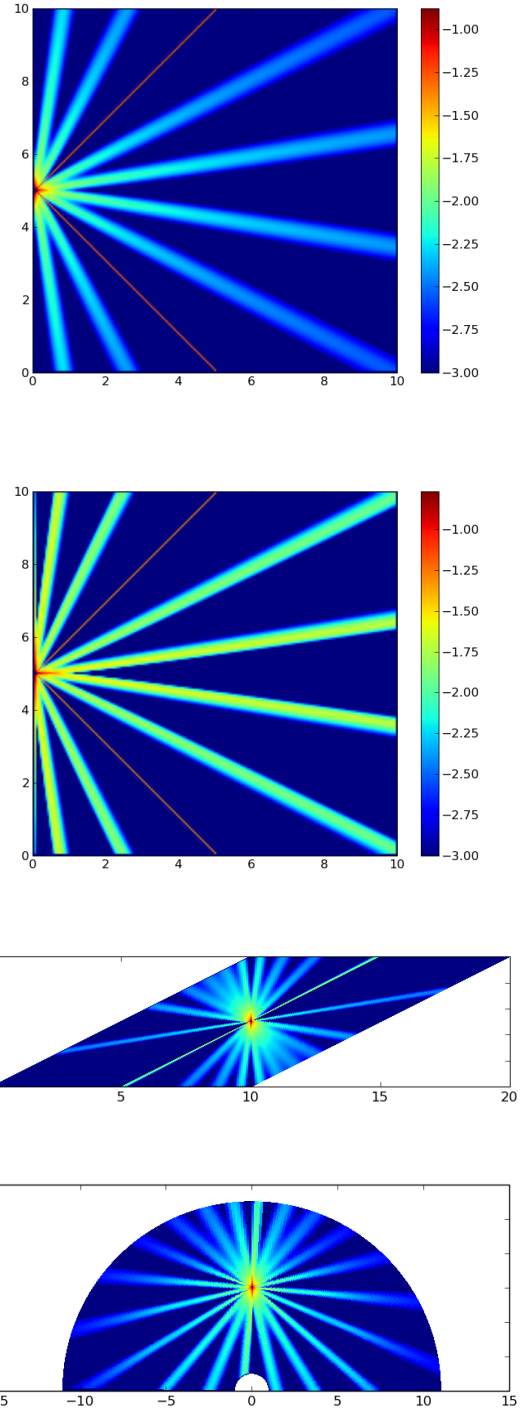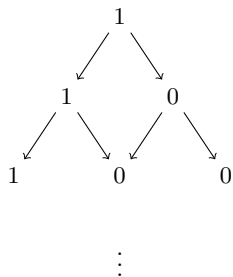

Figure A1. Comparison of the radiation solution for a point source computed by short characteristics for different grid choices ($N_A = 20$). Colour indicates $\log(J)$. From top to bottom: a) cartesian grid (100x100); b) same as a), but using quadratic interpolation; c) sheared cartesian (100x100); d) Polar grid (60x100). In all cases, we place a delta function source function that emits isotropically. The pencil beams that appear are the ray-defects – physically, the far-field radiation pattern should be spherical and smooth but the limitations of treating radiation only locally in the SC method results in the thin beams (see discussion in §A1).

local propagation of light develop a "Pascal's Triangle" characteristic beam pattern with increasing distance. If the radiation from a source cell is split with some linear combination to its neighbouring cells and the same method is applied uniformly for all cells, then the propagation of light has the following characteristic shape – consider the case where the mixing coefficients are $[1/2, 1/2]$:

$$
\begin{array}{ccccccccc}
& & & & 1 & & & & \\
& & & \swarrow & & \searrow & & & \\
& & 1/2^1 & & & & 1/2^1 & & \\
& \swarrow & & \searrow & & \swarrow & & \searrow & \\
1/2^2 & & & & 2/2^2 & & & & 1/2^2 \\
\end{array}
$$

$$
\begin{array}{ccccccc}
1/2^3 & & 3/2^3 & & 3/2^3 & & 1/2^3 \\
\end{array}
$$

$$
\begin{array}{ccccccccc}
1/2^4 & & 4/2^4 & & 6/2^4 & & 4/2^4 & & 1/2^4 \\
\end{array}
$$

$$\vdots$$

The final far field pattern simply corresponds to the weights of a discrete random walk. This shape is a Gaussian, with a characteristic width of $w \propto \sqrt{d}$ where d is the total propagation distance from the source. The beam far from the source will therefore have angular size $w/d \propto 1/\sqrt{d} \to 0$ as $d \to \infty$. Thus, at a sufficiently large distance, the propagation of light via short characteristics will inevitably result in a series of laser beams.

The defect problem can be further exacerbated if a ray happens to pass exactly through a neighbouring cell centre. In this case, the mixing/interpolation coefficients are $[1, 0]$ and the propagation diagram reduces to:

$$
\begin{array}{ccccc}
& & 1 & & \\
& \swarrow & & \searrow & \\
1 & & & & 0 \\
\end{array}
$$

$$
\begin{array}{ccccc}
1 & & 0 & & 0 \\
\end{array}
$$

$$\vdots$$

Notice that for this limiting case the beam pattern does not spread at all, resulting in a zero width beam of constant intensity at all distances. This kind of defect is particularly disastrous in the case of spherical coordinates since it affects all radial rays. Thus, any compact light source near the centre of the spherical grid will develop serious ray-defects at larger radii. The effect is quite devastating if there is a strong point source at the centre. The radial ray defect acts to force a constant intensity on all radial rays, resulting in constant radiation energy density independent of distance rather than the expected inverse square falloff. The natural way to correct non-spreading beams is to introduce some degree of artificial broadening (i.e. force the mixing coefficients to have some minimum floor value). We describe our approach in detail in the next section.

A2 Ray Defect Correction Schemes

The most obvious approach to mitigate the impact of ray defects is simply increasing the angular resolution. However this approach is not feasible for complex problems with large spatial domains due to the unfavorable computational cost scaling with resolution.

One idea to combat ray defects is to use higher order interpolation schemes in treating the propagation of radiation. Davis et al. (2012) found that quadratic monotonic interpolation works well and suggest that it be used over standard linear interpolation. Unfortunately, this does not address the fundamental problem of beam spreading – higher order schemes are actually counterproductive in that they produce even thinner beam patterns and amplify the significance of ray defects (Compare the top two panels of Figure A1 for linear vs. quadratic interpolation) .

Ultimately, the problem (as illustrated in §A1) has to do with the lack of spreading for a single beam. This motivates us to implement a diffusive scheme to bring back the necessary amount of spreading and to set the intensity values in such a way that an inverse square falloff of radiation density and flux is recovered.

Our approach is the following: every time we read out intensity values from our angular grid (i.e. the interpolation scheme described in §2.7), we apply a floor to our interpolation coefficients so that there is always some degree of mixing between nearby angles. That is, we modify the coefficients c_1, c_2, c_3 in Eq. 39 to new values c_1', c_2', c_3' given by:

$$c_i' = \frac{c_i + w}{\left(\sum_i c_i\right) + 3w} \tag{A1}$$

where w is the diffusion coefficient that controls the amount of angle mixing. A more physical way to understand this diffusion coefficient is to convert it into an attenuation factor for nearby cells. For a monodirectional beam (i.e. $I = 0$ for all but a single angle bin), the beam will attenuate by a factor a after propagating across a single cell, i.e.,

$$I^{n+1} = aI^n, \tag{A2}$$

where from Eq. A1, we find:

$$a = (1 + w)/(1 + 3w). \tag{A3}$$

This motivates a method for choosing the appropriate value of w for the case of a log-polar grid (where the radial spacing is uniform in $\log r$). Since we desire an inverse square law falloff of the radiation field, we simply solve for a in the equation $a^{N_{dec}} = 0.01$ where N_{dec} is the number of cells per radial decade. For example if $N_{dec} = 20$ points per decade (our canonical choice), we have have $a = 0.794$ and $w = 0.149$.

In Figure A2, we show radial profiles of F calculated with the SC for a point source central star. Note that, in the absence of diffusion, the effect of radial ray defects is to produce an unphysical solution where F is constant with radius.

In Figure A3, we show the results from this diffusion scheme compared to the exact solution for an off-centre equatorial ring source (e.g. the same setup as described in §3.7.1) for different choices of the spatial grid resolution. We find that our diffusive scheme suffers from systematic over-attentuation in the polar regions for spatial grids that too
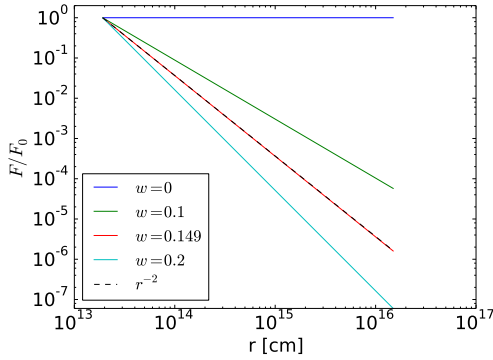
Figure A2. Radial profiles of F for various choices of diffusion coefficient as calculated by short characteristics code on a spherical polar grid $(n_r, n_\theta) = (60, 100)$. Given these grid dimensions, the diffusion coefficient must be set to $w = 0.149$ in order to reproduce an inverse square falloff for flux.

rectangular. We recommend keeping the grid cell aspect ratio as close to square (1:1) as possible in order to minimize this effect.

The scheme described here is a generalization of that proposed by Dullemond & Turolla 2000. They applied the same factor a derived above except that they did it solely for the radial ray. For a central point source, which produces only radial rays, their method and ours produce the same correct inverse-law behaviour of flux. However, when the source is extended, e.g., the spherical shell test problem described in §3.8.5, their method causes the intensity to fall too steeply. Our method also suffers a similar systematic bias, but it is less severe (see Fig. 17). In the limit of an isotropic radiation field, e.g., the interior of a blackbody enclosure, their method would still cause intensity to decline with increasing radius. Our diffusive method, on the other hand, would give the correct result, viz., no change in intensity. These differences are relatively minor. The main feature of both approaches is that they recover the inverse square law at least approximately. However, a major limitation of our diffusive approach is that it only works for the case of uniform log-radial grids. For more complex grids (especially those with nonregular arbitrary/adaptive meshes), the only generalizable solution is to invoke a few iterations of LC to resolve the inverse square falloff from each emitting source. In particular, if the illumination comes primarily from a very limited set of point sources, then the explicit LC handling of just these point sources shouldn't significantly impact the overall runtime.

APPENDIX B: ANALYTIC 1D ATMOSPHERE SPECTRUM

In this problem, we set up a constant flux 1D atmosphere subject to fixed grey opacities ($\kappa_\nu = $ const) with an absorption fraction set to $\epsilon = 10^{-6}$ (i.e. a highly scattering-dominated atmosphere). The local source function is a combination of thermal emission and reflected light (c.f. Eq 6).

The analytic solution to this problem can be easily obtained by considering the moments of the radiative transfer

equation. We define the first few moments of the intensity field as

$$J_\nu = \frac{1}{2} \int_{-1}^{1} I_\nu d\mu,$$

$$H_\nu = \frac{1}{2} \int_{-1}^{1} \mu I_\nu d\mu,$$

$$K_\nu = \frac{1}{2} \int_{-1}^{1} \mu^2 I_\nu d\mu, \tag{B1}$$

where $\mu \equiv \cos(\theta)$ is the angle cosine with respect to the plane normal. Using these quantities allows us to write the moments of the radiative transfer equation as:

$$\frac{dH_\nu}{d\tau_\nu} = \epsilon(J_\nu - B_\nu) \tag{B2}$$

$$\frac{dK_\nu}{d\tau_\nu} = H_\nu \tag{B3}$$

Combining Eqs. B2 + B3 and invoking the Eddington approximation ($K_\nu = J_\nu/3$) yields

$$\frac{d^2 J_\nu}{d\tau_\nu^2} = 3\epsilon(J_\nu - B_\nu), \tag{B4}$$

which is simply a constant coefficient second-order inhomogeneous differential equation in J_ν. This allows us to construct an exact solution using standard methods. The solutions to the homogeneous system are simply

$$\phi_1(\tau) = \exp(\sqrt{3\epsilon}\tau),$$
$$\phi_2(\tau) = \exp(-\sqrt{3\epsilon}\tau). \tag{B5}$$

The particular solution J_p that satisfies the inhomogeneous system is given by[3]

$$J_p(\tau) = \phi_1(\tau) \int_{\tau}^{\infty} \frac{\phi_2 g}{W} d\tau' + \phi_2(\tau) \int_{0}^{\tau} \frac{\phi_1 g}{W} d\tau', \tag{B6}$$

where $g = 3\epsilon B$ is the inhomogeneous function, and W denotes the Wronskian, defined by

$$W \equiv \frac{d\phi_1}{d\tau}\phi_2 - \frac{d\phi_2}{d\tau}\phi_1 \tag{B7}$$

$$= 2\sqrt{3\epsilon} \tag{B8}$$

Putting everything together, the solution takes the form

$$J = J_p + c_1\phi_1 + c_2\phi_2, \tag{B9}$$

where the undetermined constants are set by the boundary conditions of the problem. At the $\tau \to \infty$ inner boundary we expect $J \to B$, so we must eliminate the exponentially growing mode by setting $c_1 = 0$. To set the surface boundary condition, we make use of the two-stream approximation and evaluate

$$H(\tau) = \frac{1}{\sqrt{3}} \frac{dJ}{d\tau} \tag{B10}$$

---

[3] The particular solution is constructed via the "variation of parameters" method
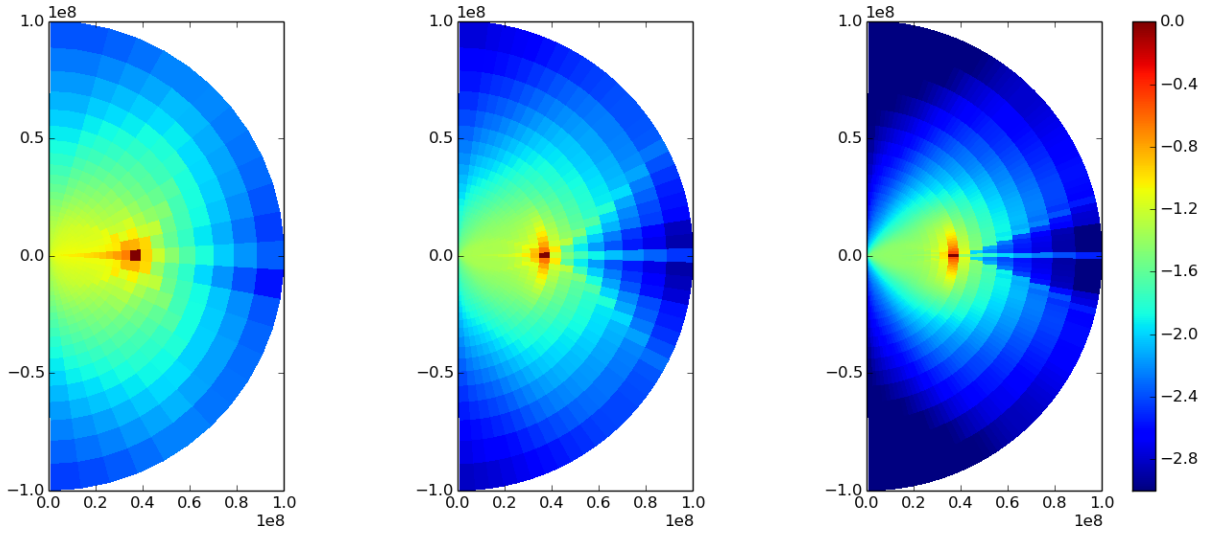
Figure A3. Ray defect pattern from an isotropically emitting ring for 3 different choices of grid resolution. Here we use a spherical spatial grid with locally defined ray angles. Panels: a) ($n_\theta = 26$) square aspect ratio; b) ($n_\theta = 52$) 1:2 aspect ratio; c) ($n_\theta = 104$) 1:4 aspect ratio. Notice the systematic reduction of radiation near the polar regions when n_θ increases and also the characteristic "spider" pattern arising from our choice of spatial grid interpolation.

and enforce a surface boundary condition that is consistent with the Eddington approximation:

$$\frac{H(0)}{J(0)} = \frac{1}{\sqrt{3}}, \tag{B11}$$

which sets

$$c_2 = -\left(\frac{1-\sqrt{3}}{1+\sqrt{3}}\right)\phi_2 \int_0^\infty \frac{\phi_2 g}{W} d\tau'. \tag{B12}$$

The final step is to determine the thermal source ($B(\tau)$) that is consistent with our radiation solution from Eq. B9. This can be calculated using our two radiative transfer moment equations. Since we have a constant flux atmosphere, $dH/d\tau = 0$, which implies $J(\tau) = B(\tau)$ from Eq. B2. We combine this with integrating the pressure equation (Eq. B3) to yield the full solution

$$J(\tau) = 3H\tau + J(0)$$
$$\rightarrow B(\tau) = 3H\left(\tau + \frac{1}{\sqrt{3}}\right) \tag{B13}$$