

Beyond moments: relativistic Lattice-Boltzmann methods for radiative transport in computational astrophysics

L. R. Weih¹★, A. Gabbana², D. Simeoni^{2,3,4}, L. Rezzolla^{1,5,6}, S. Succi^{7,8,9}, R. Tripiccione²

¹ *Institut für Theoretische Physik, Max-von-Laue-Str. 1, 60438 Frankfurt, Germany*

² *Università di Ferrara and INFN-Ferrara, I-44122 Ferrara, Italy*

³ *Bergische Universität Wuppertal, D-42119 Wuppertal, Germany*

⁴ *University of Cyprus, CY-1678 Nicosia, Cyprus*

⁵ *School of Mathematics, Trinity College, Dublin 2, Ireland*

⁶ *Helmholtz Research Academy Hesse for FAIR, Max-von-Laue-Str. 12, 60438 Frankfurt, Germany*

⁷ *Center for Life Nano Science @ La Sapienza, Italian Institute of Technology, Viale Regina Elena 295, I-00161 Roma, Italy*

⁸ *Istituto Applicazioni del Calcolo, National Research Council of Italy, Via dei Taurini 19, I-00185 Roma, Italy*

⁹ *Harvard-SEAS, Oxford Street 29, 02130 Cambridge, USA*

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

We present a new method for the numerical solution of the radiative-transfer equation (RTE) in multidimensional scenarios commonly encountered in computational astrophysics. The method is based on the direct solution of the Boltzmann equation via an extension of the Lattice Boltzmann (LB) methods and allows to model the evolution of the radiation field as it interacts with a background fluid, via absorption, emission, and scattering. As a first application of this method, we restrict our attention to a frequency independent (“grey”) formulation within a special-relativistic framework, which can be employed also for classical computational astrophysics. For a number of standard tests that consider the performance of the method in optically thin, optically thick and intermediate regimes with a static fluid, we show the ability of the LB method to produce accurate and convergent results matching the analytic solutions. We also contrast the LB method with commonly employed moment-based schemes for the solution of the RTE, such as the M1 scheme. In this way, we are able to highlight that the LB method provides the correct solution for both non-trivial free-streaming scenarios and the intermediate optical-depth regime, for which the M1 method either fails or provides inaccurate solutions. When coupling to a dynamical fluid, on the other hand, we present the first self-consistent solution of the RTE with LB methods within a relativistic-hydrodynamic scenario. Finally, we show that besides providing more accurate results in all regimes, the LB method features smaller or comparable computational costs compared to the M1 scheme. We conclude that LB methods represent a competitive and promising avenue to the solution of radiative transport, one of the most common and yet important problems in computational astrophysics.

Key words: radiative transfer – radiation: dynamics – neutrinos – scattering – methods: numerical

1 INTRODUCTION

The proper treatment of the dynamics of radiation, as it interacts with a matter fluid, is a fundamental problem in essentially all astrophysical phenomena and requires the solution of the radiative-transport equation (RTE). Given the complexity of the RTE and the nonlinear regimes normally encountered in astrophysical scenarios, the inclusion of ra-

diative effects inevitably commands the use of advanced numerical methods to solve the RTE.

An important and representative example is a binary system of merging neutron stars (see [Baiotti & Rezzolla \(2017\)](#); [Paschalidis \(2017\)](#) for an overview), where the radiative transport of neutrinos can alter significantly the chemical composition of the ejected matter, the efficiency of its ejection, as well as the stability of the post-merger object. Furthermore, radiative-transport effects are expected to play a fundamental role in the kilonova signal that is produced

★ weih@itp.uni-frankfurt.de

(Rosswog et al. 2014; Dietrich & Ujevic 2017; Siegel & Ciolfi 2016; Bovard et al. 2017; Perego et al. 2017; Fujibayashi et al. 2018; Siegel & Metzger 2017; Fernández et al. 2019) and might even be relevant for producing a short gamma-ray burst associated with such a merger. An equally important astrophysical scenario where radiation transport plays a fundamental role is the one explored in simulations of neutrino-driven core-collapse supernovae (Mezzacappa et al. 2001; O’Connor 2015; Just et al. 2015; Kuroda et al. 2016), where the radiation in form of neutrinos is essential for the explosion mechanism, which may require a fine balance, of the order of a few percent, between the energy deposited in the stalled accretion shock and the release of potential gravitational energy by the collapsing matter (see Janka et al. (2007) for an overview). Finally, one more classical astrophysical scenario where radiative-transfer effects cannot be ignored is the study of accretion flows around black holes (Zanotti et al. 2011; Fragile et al. 2012; Roedig et al. 2012; Sądowski et al. 2013; McKinney et al. 2014), for which a broad array of techniques has been developed over the years to compare with the observations (Event Horizon Telescope Collaboration et al. 2019).

Unfortunately, the computational cost associated with the solution of the RTE in numerical astrophysics also represents a significant obstacle to the inclusion of radiative effects in numerical simulations. This is due to the properties of the fundamental equation behind the RTE, i.e., the Boltzmann equation for massless particles, which lives in a seven-dimensional space of time (one dimension), configuration space (three dimensions) and momentum space (three more dimensions). As a result, the solution of the RTE for typical astrophysical scenarios as the ones mentioned above exceeds the current capacities of supercomputers and thus approximate methods have to be used¹.

A low-order approximation commonly employed in binary neutron-star simulations is the *leakage scheme* (Ruffert et al. 1996; Rosswog & Liebendörfer 2003; Galeazzi et al. 2013; Perego et al. 2014; Most et al. 2019), which only allows for cooling via the emission of neutrinos and is therefore not useful for core-collapse supernovae simulations, where heating is essential for reviving the shock. These effects can be included by an approximation of similar simplicity, the *flux-limited diffusion* approximation (Pomraning 1981; Levermore & Pomraning 1981); a recent implementation of this scheme has been presented by Rahman et al. (2019). In this method, the zeroth moment of the radiation distribution function, i.e., the radiation energy-density, is evolved together with the fluid quantities. Since the zeroth moment does not provide any information about the direction of the radiation fluxes, its implementation is useful only for systems with clear underlying symmetries. These symmetries are not present in the case of binary neutron-star simulations, so the flux-limited diffusion does not offer but a crude approximation of the radiative effects. This is also true for the M0 scheme developed by Radice et al. (2016), which also evolves the lowest moment of the distribution function, but in the free-streaming limit.

¹ In some cases the direct solution of the Boltzmann equation is indeed feasible by exploiting symmetries and reducing the spatial dimensionality of the problem.

A considerably better approximation is the so-called truncated moment-based scheme developed by Thorne (1981) and first implemented in general relativity by Rezzolla & Miller (1994) in one dimension and by Shibata et al. (2011); Cardall et al. (2013) in three dimensions. Within these schemes, the lowest moments up to order N of the distribution function are evolved, and the flux-limited diffusion method is then the limiting case for $N = 0$ of the general set of moment-based schemes. Increasing the order of the hierarchy to the case with $N = 1$ implies that the momentum-density vector is evolved together with the radiation energy density. Such a scheme is known in the literature as “M1 scheme” and is indeed one of the most commonly used methods for radiative transport throughout many different applications of computational relativistic astrophysics (Rezzolla & Miller 1994; Roedig et al. 2012; Sądowski et al. 2013; Fragile et al. 2014; McKinney et al. 2014; O’Connor 2015; Foucart et al. 2015; Skinner et al. 2019; Melon Fuksman & Mignone 2019; Weih et al. 2020b). With this method, it is possible to track the *average* direction of the radiation momentum, providing a significant improvement over the previously mentioned schemes, but can still lead to rather unphysical results such as those that emerge when radiation beams interact and cross (see e.g., Fragile et al. (2014); Foucart et al. (2015); Weih et al. (2020b)). Furthermore, as is typical in moment-based schemes, in the M1 approach the set of evolution equations for the zeroth and first moment depend on the second moment, which is not known within this hierarchy scheme and has to be approximated in the form of a closure relation. This results in the M1 scheme only being exact in the limits of high and/or low optical depths, but not in the intermediate regime. The situation does not improve when going to methods with $N > 1$, which suffer from increased computational costs and also need to specify a closure relation that expresses the $N + 1$ moment in terms of the lower-order ones.

A more accurate solution of the RTE is offered by a completely different class of methods employing Monte-Carlo techniques for the handling of the radiation field (Foucart 2018; Miller et al. 2019), which, however, suffer from low-statistics numerical noise and a comparatively high computational cost. To summarise, state-of-the-art numerical solutions of the RTE in computational relativistic astrophysics revolve around two main classes of methods: i) approximate methods based on the laws of hydrodynamics (such as leakage, flux-limited diffusion, M1) or ii) direct solutions for specific cases where symmetries can be exploited (such as Monte-Carlo approaches to the solution of the Boltzmann equation).

The scope of this paper is to introduce a new method for the solution of the RTE and hence for the treatment of radiative effects in computational astrophysics that promises to be more precise than the M1 scheme, providing a correct treatment of intersecting radiation beams and an accurate treatment of regimes of high, low, and intermediate optical depths. At the same time, it comes with an algorithmic complexity and an associated computational cost comparable to that of M1 schemes, thus making it well suited for multi-dimensional astrophysical simulations.

In essence, this new radiative-transport scheme stems from the *Lattice Boltzmann method* (LB method) (Krüger et al. 2017; Succi 2018), which is commonly used in com-

computational fluid dynamics as an alternate scheme to direct hydrodynamic solvers.

The application of the LB method as a solver for radiation transfer problems is relatively new (Asinari et al. 2010), and most models proposed so far apply only to the analysis of steady-state radiation-transport problems in one and two dimensions (Bindra & Patil 2012; Mishra et al. 2014; McCulloch & Bindra 2016; Yi et al. 2016), with very few studies carried out in three dimensions (McHardy et al. 2016; Wang et al. 2019). More recent developments (Mink et al. 2020) have shown in detail that LB offers an accurate and efficient tool in the diffusive regime of radiation transport, though struggling in the transition towards ballistic conditions. Finally, we should stress that all these previous studies deal with radiative transfer in non-relativistic regimes.

We introduce a new LB solver for studying the time dependent evolution of radiation that interacts via emission, absorption and scattering with a (dynamic) background fluid. We make use of high-order spherical quadrature rules, which, thanks to their high-order isotropy, allow to significantly extend the applicability of the method to a wider range of kinetic regimes. We work in a special-relativistic framework and present, to the best of our knowledge, the first self-consistent coupled simulation of an LB solver for radiative transport with a dynamically evolving fluid background.

Our paper is structured as follows: in Sec. 2 we present a short summary and introduction of the classical LB method and discuss its advantages, which make it ideal for radiative-transport problems in computational astrophysics. In Sec. 3 we illustrate the details of the new LB method for radiative transport in special relativity and within a “grey” (i.e., energy averaged) approximation and show how to implement such a scheme suitably for simulations. We verify this implementation by a number of standard tests in Sec. 4. In Sec. 5 we couple our new LB-code to a hydrodynamics code, which is representative of the many relativistic-hydrodynamics codes used in the field of numerical astrophysics, and present a simulation of a relativistic jet; we show that its dynamics changes qualitatively due to the back-reaction of the produced radiation. Finally, we compare its accuracy and computational cost to the commonly used M1 scheme in Sec. 6. We conclude and discuss future prospects of our new method in Sec. 7.

Throughout this paper, we use units with $c = 1$ and only write the speed of light explicitly in equations where it is necessary for clarity. We write three-vectors in boldface while unit-vectors carry a hat.

2 A SHORT INTRODUCTION TO LATTICE BOLTZMANN

In this section we provide a brief overview of the LB method. The reader already familiar with the topic may safely jump directly to Sec. 3. We should stress that, for ease of presentation, in this section we will summarise the conceptual steps of the derivation and algorithmic structure of LB in a non-relativistic framework. For details on the derivation of the method in special-relativity, the interested reader is referred to a recent review (Gabbana et al. 2020).

The LB method has emerged in the past decades as a

computationally efficient numerical tool for the simulation of the dynamic of fluids in classical hydrodynamics. Its origin can be traced back to the pioneering work on discrete velocity models (Broadwell 1964) in the 1960s and later on to the work done on Lattice Gas Cellular Automata (Hardy et al. 1973; Frisch et al. 1986) in the late 1980s. Since then, the method has evolved as an independent and efficient alternative to direct Navier-Stokes solvers in the field of classical computational fluid dynamics (McNamara & Zanetti 1988; Higuera et al. 1989) and allied disciplines, primarily soft matter (Succi 2015; Dünweg & Ladd 2009).

Contrary to direct hydrodynamic solvers, the LB method relies on the underlying microscopic dynamics of the fluid constituents – be them molecules or photons – and therefore the natural theoretical framework to start approaching the method is kinetic theory and its mathematical cornerstone, the Boltzmann equation (here taken without source terms):

$$\left(\frac{\partial}{\partial t} + \mathbf{v} \cdot \nabla\right) f(\mathbf{r}, \mathbf{v}, t) = C(\mathbf{r}, \mathbf{v}, t). \quad (1)$$

The distribution function $f(\mathbf{r}, \mathbf{v}, t)$ refers to the number of particles with velocity \mathbf{v} at position \mathbf{r} at time t , while the collision operator $C(\mathbf{r}, \mathbf{v}, t)$ accounts for collisions between point-particles in the fluid and in Boltzmann’s theory takes the form of a non local integral in momentum space. It is customary to replace the full collisional operator with a simplified model, such as the well-known Bhatnagar-Gross-Krook (BGK) relaxation time approximation (Bhatnagar et al. 1954), encompassing the natural tendency of the system to relax towards an equilibrium, i.e., the tendency of f to reach a distribution function f^{eq} describing a local equilibrium state

$$C(\mathbf{r}, \mathbf{v}, t) = -\frac{1}{\tau}(f(\mathbf{r}, \mathbf{v}, t) - f^{\text{eq}}(\mathbf{r}, \mathbf{v}, t)). \quad (2)$$

In the above, τ represents the typical timescale needed to reach the equilibrium, a parameter which controls the hydrodynamic transport coefficients, hence dissipative phenomena within the fluid.

For a classical fluid with massive constituents, the equilibrium function is represented by the Maxwell-Boltzmann distribution (see, e.g., Rezzolla & Zanotti 2013)

$$f^{\text{eq}}(\mathbf{r}, \mathbf{v}, t) = \rho(\mathbf{r}, t) \left(\frac{m}{2\pi k_{\text{B}} T}\right)^{\frac{d}{2}} \exp\left[-\frac{m}{k_{\text{B}} T}(\mathbf{v} - \mathbf{u}(\mathbf{r}, t))^2\right], \quad (3)$$

where m is the mass of constituent particles and k_{B} is the Boltzmann constant. The rest-mass density $\rho(\mathbf{r}, t)$ and the velocity field $\mathbf{u}(\mathbf{r}, t)$ can be computed as the zeroth and first moment of the distribution function, respectively,

$$\rho(\mathbf{r}, t) := m \int f(\mathbf{r}, \mathbf{v}, t) d\mathbf{v}, \quad \mathbf{u}(\mathbf{r}, t) := \frac{m}{\rho(\mathbf{r}, t)} \int \mathbf{v} f(\mathbf{r}, \mathbf{v}, t) d\mathbf{v}. \quad (4)$$

Historically, the LB method was devised as a noise-free (pre-averaged) version of its lattice-gas cellular automaton (LGCA) ancestor (McNamara & Zanetti 1988). This represented a major conceptual leap, but left all other LGCA shortcomings untouched, primarily the exponential complexity barrier associated with Boolean collision operators and the ensuing low collisional rates which prevented LGCA from accessing high-Reynolds regimes (turbulent flows). All

of the above barriers were lifted just months later in a short sequence of papers (Higuera et al. 1989; Higuera & Jiménez 1989; Higuera & Succi 1989), which placed LB on the map of computational fluid dynamics.

For the sake of simplicity and continuity with continuum kinetic theory, it proves expedient to derive LB from the expansion of the equilibrium distribution in a series of orthogonal Hermite polynomials $\mathbf{H}^{(k)}(\mathbf{v})$; the advantage of using Hermite as the expansion basis is that the expansion coefficients $\mathbf{a}^{(k)}$ coincide with the moments of the distribution function (Grad 1949a,b). The expansion is then truncated to the desired order S , high enough to recover the macroscopic observables of interest (Shan & He 1998; Shan et al. 2006).

In d dimensions, this results in the following local equilibrium distribution function:

$$f^{\text{eq}} = \left(\frac{1}{2\pi}\right)^{\frac{d}{2}} \exp\left(-\frac{\mathbf{v}^2}{2}\right) \sum_{k=0}^S \mathbf{a}^{(k)}(\rho, \mathbf{u}) \mathbf{H}^{(k)}(\mathbf{v}), \quad (5)$$

where all physical quantities have been made dimensionless by appropriate scaling with respect to a characteristic velocity $\tilde{c} := \sqrt{k_{\text{B}}T_0/m_0}$, a characteristic temperature T_0 , a mass unit m_0 , and a length scale L_0 . Furthermore, the (microscopic) velocity vector of phase-space is discretized using a set of N_{pop} distinct populations \mathbf{v}_i with $i \in [1; N_{\text{pop}}]$. As a consequence, the distribution f itself becomes a set of N_{pop} functions $f_i(\mathbf{r}, t) = f(\mathbf{r}, \mathbf{v}_i, t)$, each accounting for the particles moving along the discrete direction \mathbf{v}_i .

The choice of the discrete velocities lies at the heart of the LB method, the informing criterion being of reproducing *exactly* the set of kinetic moments which describe the low-Knudsen hydrodynamic regime, namely the mass density (scalar), the flow current (vector) and the momentum flux tensor (second-order tensor). In the above, “exactly” means that no error altogether is incurred by replacing the integrals in continuum velocity space with the corresponding summations over the discrete velocities which characterise the LB representation. Formally, this can be linked to a Gauss-Hermite quadrature rule, where one defines \mathbf{v}_i , and the corresponding weights w_i , and requires *exact* preservation of the relevant hydrodynamic fields. In equations:

$$\rho(\mathbf{r}, t) = \sum_{i=1}^{N_{\text{pop}}} f_i(\mathbf{r}, t), \quad \mathbf{u}(\mathbf{r}, t) = \frac{1}{\rho(\mathbf{r}, t)} \sum_{i=1}^{N_{\text{pop}}} \mathbf{v}_i f_i(\mathbf{r}, t), \quad (6)$$

with the truncated equilibrium distribution Eq. (5) given by

$$f_i^{\text{eq}} = w_i \sum_{k=0}^S \mathbf{a}^{(k)}(\rho, \mathbf{u}) \mathbf{H}^{(k)}(\mathbf{v}_i). \quad (7)$$

The combination of the velocity discretization with explicit time-marching finally delivers the lattice Boltzmann equation:

$$f_i(\mathbf{r} + \mathbf{v}_i \Delta t, t + \Delta t) = f_i(\mathbf{r}, t) + \Delta t C_i(\mathbf{r}, t), \quad (8)$$

with Δt the time-step, $\Delta x = \mathbf{v}_i \Delta t$ the characteristic mesh spacing, and

$$C_i = \frac{1}{\tau} (f_i - f_i^{\text{eq}}). \quad (9)$$

The evolution of Eq. (8) follows the so-called “stream-and-collide” paradigm, where in the “collide step” each popula-

tion $f_i(\mathbf{r}, t)$ is updated by receiving a local collisional contribution:

$$f_i^*(\mathbf{r}, t) = f_i(\mathbf{r}, t) + \Delta t C_i(\mathbf{r}, t). \quad (10)$$

In the streaming step, instead, the post collision populations $f_i^*(\mathbf{r}, t)$ stream along their associated direction \mathbf{v}_i , landing on the corresponding neighbouring lattice site (no particle can fly off-grid):

$$f_i(\mathbf{r} + \mathbf{v}_i \Delta t, t + \Delta t) = f_i^*(\mathbf{r}, t). \quad (11)$$

Two major assets associated to the stream-collide paradigm are worth highlighting. First, the non-local operator (streaming) is linear and the nonlinear one (collision) is local, meaning that, at variance with the hydrodynamic representation, non-linearity and non-locality are disentangled. This is because information always travels along constant characteristics, the discrete velocities, regardless of the spacetime complexity of the emergent hydrodynamics. By contrast, in the fluid representation, information travels along space-time dependent material lines, defined by the local flow speed. This is a major advantage also for the handling of complex boundary conditions and parallel computing.

Second, since dissipation emerges from enslaving of the Boltzmann distribution to local equilibrium, there is no need for second order spatial derivatives. This is a significant advantage for the calculation of the stress tensor, especially near solid boundaries. In addition, since the collision operator is conservative to machine-accuracy, the LB method usually offers better accuracy than most grid-based discretisations of the Laplace operator.

The standard LB method described so far is suitable for the description of hydrodynamic systems, where the molecular mean free path is much shorter than the shortest hydrodynamic length scale, the ratio of the two being the Knudsen number of the fluid. On the other hand, for a fluid of radiation particles, e.g., photons or neutrinos, one should bear in mind that the radiation constituents do not interact among themselves, but only with the background fluid, which effectively produces, destroys, and scatters the radiation particles. Hence, when applying the LB method to radiation local conservation laws must be revisited and complemented with suitable source (sink) terms, accounting for the above processes. In particular, energy and momentum are conserved only over the combined system of radiation and fluid.

This still fits within the LB framework, which has been used for decades to study transport phenomena, such as advection-diffusion-reaction equations, (Massaioli et al. 1993; He et al. 1998; Peng et al. 2003; Karlin et al. 2013). Based on this idea, several LB models have been proposed to study radiative transport. Initially most of these efforts have focused on studying steady-state problems in one and two spatial dimensions (Asinari et al. 2010; Bindra & Patil 2012; Mishra et al. 2014; McCulloch & Bindra 2016; Yi et al. 2016), considering isotropic as well as anisotropic scattering (Vernekar & Mishra 2014). So far, however, only very few authors have considered the three-dimensional case (McHardy et al. 2016; Mink et al. 2020).

All the models mentioned above discretize the velocity space by means of standard space-filling lattices, typically used in LB methods. As a consequence, information travels in a single timestep to nodes located at different distances,

corresponding to a different magnitude of the discrete velocities, a typical example being the two-dimensional nine-velocity lattice comprising a rest particle with zero speed, four particles connecting to the nearest neighbours (speed 1) and four connecting the diagonals (speed $\sqrt{2}$). The latter are mandatory to the correct recovery of the two-dimensional Navier-Stokes equations, but not necessary (albeit recommended for matters of accuracy) for the case of advection-diffusion equations, a property which can be transferred to radiative LBEs (Gairola & Bindra 2017; Wang et al. 2019). The obvious drawback is that such models are appropriate only for collision-dominated, low-diffusive, regimes, while higher-order models based on extended velocity sets are required to describe the transition from low to high diffusivity and finally towards the free-streaming (ballistic) regimes.

In the following, we introduce a new LB method which is precisely meant to address this important issue. In particular, we employ high-order spherical quadratures to provide a unified numerical RTE solver, capable of handling both low and high diffusive regimes, up to the free-streaming scenario.

This major extension comes at a price: since the discrete velocities lie on a sphere, they no longer end on the nodes of a space-filling Cartesian grid. Hence interpolation is required to supplement the standard stream-and-collide algorithm, leading to the loss of exact streaming. Actual simulations show that in all cases inspected in this paper, the lack of exact streaming does not lead to any appreciable loss of accuracy of the method, thus identifying the radiative LB developed here as a viable and competitive numerical RTE solver.

3 THE LATTICE BOLTZMANN SCHEME FOR RADIATIVE TRANSPORT

3.1 Mathematical setup

The RTE is the master equation describing how radiation propagates through a medium that scatters, absorbs and emits radiation particles. As such, the RTE follows from the Boltzmann equation assuming massless particles (i.e., photons or neutrinos). All radiative fields are expressed in terms of the distribution function $f_\nu(\mathbf{x}, \hat{\mathbf{n}}, \nu)$ of neutrinos or photons within a given frequency band $d\nu$ at position \mathbf{x} and velocity within a solid angle $d\Omega$ in the direction $\hat{\mathbf{n}}$.

The subscript ν highlights the dependency on frequency, to distinguish from the corresponding frequency-independent quantities to be introduced later on. The RTE describes the evolution of the radiation distribution function in the direction $\hat{\mathbf{n}}$

$$\frac{1}{c} \frac{\partial f_\nu}{\partial t} + \hat{\mathbf{n}} \cdot \nabla f_\nu = -\kappa_{a,\nu} f_\nu + \eta_\nu + C_{\text{scat}} =: C_{\text{rad}}. \quad (12)$$

This expression is equivalent to Eq. (1), but for massless particles and with an explicit expression for the collisional operator C_{rad} on the right-hand side. This operator splits into three terms: the absorption, the emission, and the scattering term, respectively. The absorption term is proportional to the absorption coefficient $\kappa_{a,\nu}$, the emission term to the emissivity η_ν , and the scattering term C_{scat} is written in its

most general form as (Bruenn 1985; Rampp 1997)

$$C_{\text{scat}} = \int_0^\infty \nu'^2 d\nu' \int_{4\pi} f'_{\nu'}(1 - f_\nu) R^{\text{in}} - f_\nu(1 - f'_{\nu'}) R^{\text{out}} d\Omega', \quad (13)$$

where $R^{\text{in}}(\nu, \nu')$ and $R^{\text{out}}(\nu, \nu')$ are the incoming and outgoing scattering kernels, respectively.

These kernels depend on the underlying physical process and generally do not allow for an analytic solution of the scattering integral. In order to simplify the above integral, the incoming and outgoing scattering kernels are typically expanded as a Legendre series and truncated to the first two terms (Bruenn 1985; Rampp 1997; Shibata et al. 2011). This leads to

$$R^{\text{in/out}}(\nu, \nu') \approx \frac{1}{2} \Phi_0^{\text{in/out}}(\nu, \nu') + \frac{3}{2} \Phi_1^{\text{in/out}}(\nu, \nu') \cos \theta, \quad (14)$$

where $\Phi_{\ell=0,1}^{\text{in/out}}$ are the ℓ -th coefficients of the Legendre expansion and θ is the angle between the incoming and outgoing particle, $\cos \theta = \hat{\mathbf{n}} \cdot \hat{\mathbf{n}}'$.

Hereafter, we will consider only iso-energetic scatterings, i.e., we assume that the energy of the radiation particles is left unchanged by the scattering with the constituents of the underlying fluid.

It follows that: $\Phi_\ell^{\text{in}}(\nu, \nu') \equiv \Phi_\ell^{\text{out}}(\nu, \nu') =: \Phi_\ell(\nu) \delta(\nu - \nu')$.

In this way, inserting Eq. (14) in Eq. (13), the scattering term in Eq. (12) reads as:

$$C_{\text{scat}} \approx -\kappa_{0,\nu} f_\nu + \kappa_{0,\nu} E_\nu + 3\kappa_{1,\nu} \hat{\mathbf{n}} \cdot \mathbf{F}_\nu, \quad (15)$$

where we have used the definition of the zeroth and first moment of the distribution function

$$E_\nu := \frac{1}{4\pi} \int_{4\pi} f_\nu d\Omega, \quad \mathbf{F}_\nu := \frac{1}{4\pi} \int_{4\pi} \hat{\mathbf{n}} f_\nu d\Omega, \quad (16)$$

and defined the energy-dependent opacities $\kappa_{\ell,\nu} = 2\pi\nu^2\Phi_\ell$. Note that the explicit form of these opacities depends on the type of radiation, namely, whether one is considering photons or neutrinos (see e.g., Eqs. (A.47) and (A.48) in Rampp (1997) for the case of neutrinos or Rybicki & Lightman (1986) for photons).

The frequency-integrated version of the RTE – often referred to as “grey” approximation – is obtained via multiplication by ν^3 and integration over ν , i.e.,

$$\frac{1}{c} \frac{\partial I}{\partial t} + \hat{\mathbf{n}} \cdot \nabla I = -\kappa_a I + \eta + \kappa_0(E - I) + 3\kappa_1 \hat{\mathbf{n}} \cdot \mathbf{F}, \quad (17)$$

where we used the definition of the frequency-integrated specific intensity

$$I := \int_0^\infty I_\nu d\nu = \int_0^\infty \nu^3 f_\nu d\nu, \quad (18)$$

and the frequency-integrated moments

$$E := \int_0^\infty \nu^3 E_\nu d\nu = \frac{1}{4\pi} \int_{4\pi} I d\Omega, \quad (19)$$

$$\mathbf{F} := \int_0^\infty \nu^3 \mathbf{F}_\nu d\nu = \frac{1}{4\pi} \int_{4\pi} \hat{\mathbf{n}} I d\Omega. \quad (20)$$

These moments can be interpreted as the radiation energy density and momentum density, respectively, and arguably represent the most important properties of the radiation field.

Finally, the energy-averaged opacities κ_a , κ_0 , κ_1 are given by

$$\kappa_* := \frac{\int_0^\infty \kappa_{*,\nu} I_\nu d\nu}{\int_0^\infty I_\nu d\nu}, \quad (21)$$

where $*$ = $a, 0, 1$ and the frequency-integrated emissivity η is given by

$$\eta := \int_0^\infty \nu^3 \eta_\nu d\nu. \quad (22)$$

3.2 LB discretization of the RTE

We next describe the steps needed to derive a LB-inspired discretization of the RTE within the grey approximation. We start by recasting Eq. (17) in a BGK-like form:

$$\frac{1}{c} \frac{\partial I}{\partial t} + \hat{\mathbf{n}} \cdot \nabla I = -\kappa_0 (I - I^{\text{eq}}) + S, \quad (23)$$

where the source term S , accounting for emission and absorption, is given by

$$S = -\kappa_a I + \eta, \quad (24)$$

while the scattering term has been rearranged by introducing the equilibrium radiation intensity

$$I^{\text{eq}} := E + \lambda \hat{\mathbf{n}} \cdot \mathbf{F}, \quad \lambda = 3 \frac{\kappa_1}{\kappa_0}. \quad (25)$$

We then perform a discretization of the velocity space, namely, the directions $\hat{\mathbf{n}}$ along which radiation propagates at the speed of light. In essence, we replace $\hat{\mathbf{n}}$ with a discrete set of N_{pop} directions $\hat{\mathbf{n}}_i$, which define the corresponding discrete intensities $I_i(\mathbf{x}, t) = I(\mathbf{x}, \hat{\mathbf{n}}_i, t)$ and their associated weights w_i (a more detailed description on how to choose the N_{pop} directions is given in Sec 3.2.1). Consequently, Eq. (23) splits into a set of N_{pop} equations, each describing the evolution of a specific intensity I_i via the i -th expression of the RTE

$$\frac{1}{c} \frac{\partial I_i}{\partial t} + \hat{\mathbf{n}}_i \cdot \nabla I_i = -\kappa_0 (I_i - I_i^{\text{eq}}) + S_i. \quad (26)$$

The discrete counterparts of the integrals in Eq. (19) and Eq. (20) then take the following form:

$$E \approx \sum_{i=1}^{N_{\text{pop}}} I_i, \quad \mathbf{F} \approx \sum_{i=1}^{N_{\text{pop}}} \hat{\mathbf{n}}_i I_i, \quad (27)$$

while higher moments can be computed accordingly as

$$M^{j_1 \dots j_m} \approx \sum_{i=1}^{N_{\text{pop}}} n_i^{j_1} \dots n_i^{j_m} I_i. \quad (28)$$

Finally, a first-order discretization in time with a time-step Δt leads to the radiative Lattice Boltzmann equation:

$$I_i(\mathbf{r} + c\hat{\mathbf{n}}_i \Delta t, t + \Delta t) = I_i(\mathbf{r}, t) - c\kappa_0 \Delta t \left(I_i(\mathbf{r}, t) - I_i^{\text{eq}}(\mathbf{r}, t) \right) + c\Delta t S_i(\mathbf{r}, t), \quad (29)$$

where

$$I_i^{\text{eq}} := w_i (E(\mathbf{r}, t) + \lambda \hat{\mathbf{n}}_i \cdot \mathbf{F}(\mathbf{r}, t)), \quad S_i := w_i \eta - \kappa_a I_i(\mathbf{r}, t). \quad (30)$$

3.2.1 Velocity Discretization

The key of the LB method lies within the velocity discretization, that is, the mapping of the continuum velocity space in terms of a finite set of discrete velocities $\{\mathbf{v}_i\}$, the so-called discrete velocity stencil.

Indeed, the definition of the discrete velocity set, and its associated weights, plays a crucial role in the LB method, and several approaches have been developed over the years. The one based on the Gauss-Hermite quadrature is arguably the most systematic (Shan et al. 2006; Philippi et al. 2006; Shan 2016). Another way to go, which has been used in the early days of LB, is to construct velocity sets as d -dimensional projections from known $(d + 1)$ velocity sets d'Humières et al. (1986).

Yet, another possible approach consists in defining general conditions that should be satisfied by the velocity set, in terms of symmetry and conservation, typically mass, momentum and momentum flux (isotropy).

As mentioned above, since we cannot rely on conservation laws that would allow us to derive quadrature rules in a systematic way, we need to define the velocity stencil on the basis of symmetry considerations and isotropy conditions, privileging those stencils that exhibit a sufficiently high order of isotropy, so as to handle the different kinematic regimes that are typically encountered in astrophysical scenarios.

Formally, we define k -rank tensors $T^{\alpha_1 \dots \alpha_k}$, that are constructed by combining all the products between the different directions n_i forming the velocity stencil (appropriately weighted with the weights w_i)

$$T^{\alpha_1 \dots \alpha_k} := \sum_i w_i n_i^{\alpha_1} \dots n_i^{\alpha_k}. \quad (31)$$

The microscopic relations that have to be satisfied by the lattice in order to ensure n -th order isotropy are given by (Rivet & Boon 2001)

$$T^{\alpha_1 \dots \alpha_k} \begin{cases} = 0 & k \text{ odd,} \\ \propto \sum_{\text{perm}} (\delta_{\alpha_1 \alpha_2} \dots \delta_{\alpha_{k-1} \alpha_k}) & k \text{ even,} \end{cases} \quad (32)$$

which need to be satisfied for all $k \leq n$.

One additional condition on the definition of the velocity stencil is that all the (pseudo)-particles travel at the same speed, i.e., the speed of light, so $\mathbf{v}_i = c \hat{\mathbf{n}}_i$. It follows that the discrete directions $\hat{\mathbf{n}}_i$ must have the same magnitude, hence span the surface of a sphere in three dimensions (3D), or a circle in two dimensions (2D).

The adoption of spherical velocity stencils is not new to LB: it has been used, for example, in LB models for the simulation of ultra-relativistic hydrodynamics. In such frameworks, since the interest is restricted to the hydrodynamic picture, it is still possible to define stencils which live on the intersection between a Cartesian grid and a sphere of fixed radius (Mendoza et al. 2013; Gabbana et al. 2018), thus preserving a very desirable property of LB: exact-streaming. On the other hand, going beyond hydrodynamics generally requires a much larger number of discrete directions, making the definition of on-lattice quadratures impractical. In these cases, it is therefore necessary to take into consideration off-lattice schemes (Coelho et al. 2018; Ambrus & Blaga 2018). Since we need to properly model free-streaming regimes, we

adopt this latter approach and work with off-lattice stencils spanning a unit sphere.

Moreover, the choice of the velocity set should be such to maximise the accuracy in the calculation of the moments of the specific intensity $I(\mathbf{x}, \hat{\mathbf{n}}, t)$, such as the energy density $E(\hat{\mathbf{x}}, t)$ and the momentum density $\mathbf{F}(\hat{\mathbf{x}}, t)$.

In practice, we request the discrete sums in Eq. (27) to correctly reproduce their continuous counterparts, i.e., Eq. (19) and Eq. (20). These are spherical integrals of the form

$$Q(h) = \frac{1}{4\pi} \int_{4\pi} h(\hat{\mathbf{n}}) d\Omega \approx \sum_{i=1}^{N_{\text{pop}}} w_i h(\hat{\mathbf{n}}_i), \quad (33)$$

where $h(\hat{\mathbf{n}})$ represents a generic function of the direction $\hat{\mathbf{n}}$. To begin with, we recall that any square integrable function can be expanded on the unit sphere as a series of orthogonal spherical harmonics (Atkinson & Han 2012)

$$h(\hat{\mathbf{n}}) = h(\theta, \phi) = \sum_{\ell=0}^{+\infty} \sum_{m=-\ell}^{\ell} c_{\ell m} Y_{\ell}^m(\theta, \phi), \quad (34)$$

where the convergence rate depends on the coefficients $c_{\ell m}$. The spherical quadrature rule determining the weights w_i and discrete directions $\hat{\mathbf{n}}_i = (\theta_i, \phi_i)$ is then said to be of order p if it integrates exactly all the spherical harmonics $Y_{\ell}^m(\theta, \phi)$ up to the degree $\ell = p$

$$\frac{1}{4\pi} \int_{4\pi} Y_{\ell}^m(\theta, \phi) d\Omega = \sum_{k=0}^{N_{\text{pop}}-1} w_k Y_{\ell}^m(\theta_k, \phi_k) \quad \forall \ell \leq p. \quad (35)$$

It follows that a spherical quadrature of order p integrates exactly all integrals $Q(h)$ of functions $h(\theta, \phi)$ described by linear combinations of the first p harmonics. On the other hand, functions $h(\theta, \phi)$ that contain in their series harmonics of higher order are only approximated by the quadrature, with errors that depend on the smoothness of the function itself. Note that it is possible to prove that quadratures of order p satisfy Eq. (32) up to the level p , enabling in this way to evaluate the quality of the stencils with the use of a single parameter, i.e., p that we call the quadrature order.

Determining the quadrature satisfying the conditions discussed above is trivial in the case of two dimensions: the integral in Eq. (33) simply has to be computed on the unit circle. Furthermore, the expansion in Eq. (34) reduces to a Fourier series since the three-dimensional spherical harmonics reduces to circular functions (sine and cosines of the only angular coordinate ϕ)

$$h(\hat{\mathbf{n}}) = h(\phi) = \sum_{m=-\infty}^{\infty} c_m e^{im\phi}, \quad (36)$$

and one has to satisfy exactly the relations

$$\frac{1}{2\pi} \int_0^{2\pi} e^{im\phi} d\phi = \sum_{k=0}^{N_{\text{pop}}-1} w_k e^{im\phi_k} \quad \forall |m| \leq p. \quad (37)$$

As a result, to obtain a 2D quadrature of order p , i.e., in order to satisfy Eq. (37), it is sufficient to consider $N_{\text{pop}} = p + 1$ uniformly spaced points on the unit circle, displaced by the angles $\phi_i = 2\pi k/N_{\text{pop}}$, thus $\hat{\mathbf{n}}_i = [\cos(\phi_i), \sin(\phi_i)]^T$, all having equal weights $w_k = 1/N_{\text{pop}}$. A representation of such a velocity stencil with its discrete directions $\hat{\mathbf{n}}_i$ is presented in Fig. 1.

The problem becomes considerably more involved for

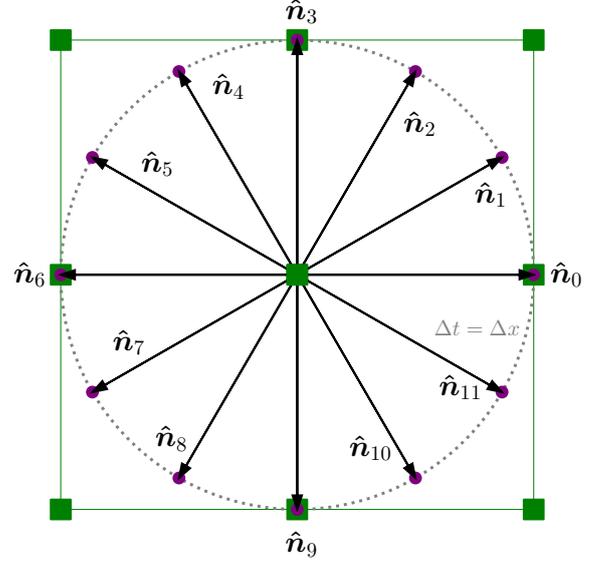


Figure 1. Example of two-dimensional velocity stencils showing $N_{\text{pop}} = 12$ discrete velocity directions (purple dots), in which the radiation is allowed to propagate. Green squares indicate the grid's cell-centres. The arrows end on a circle with radius $c\Delta t$, where Δt is assumed to be in units of the grid-spacing.

the three-dimensional case, which is clearly the most relevant one in terms of astrophysical applications. Indeed, the definition of quadrature rules on the surface of a sphere (i.e., a 2-sphere) is still an active area of research, with several different approaches coming with corresponding advantages and drawbacks, depending on the target application (Beentjes 2015; Gamba et al. 2017; Gross & Atzberger 2018; Lutsko & Lam 2018; Stepán, Jirí et al. 2020).

Here, we consider three different types of spherical quadrature schemes:

- (i) Gauss-Legendre quadrature
- (ii) Lebedev quadrature
- (iii) Spherical design

The application of the Gauss-Legendre quadrature to a 2-sphere can be obtained by making use of the product quadrature rule. Exploiting the separability of spherical harmonics, the integrals in Eq. (35) can be expressed as the product of a circular function $e^{im\phi}$ and Legendre polynomials:

$$\int_{4\pi} Y_{\ell}^m(\theta, \phi) d\Omega \propto \left(\int_0^{\pi} P_{\ell}^m(\cos \theta) \sin \theta d\theta \right) \left(\int_0^{2\pi} e^{im\phi} d\phi \right), \quad (38)$$

where $P_{\ell}^m(\cos(\theta))$ are the associated Legendre polynomials. At this point, the two integrals can be evaluated using two one-dimensional quadrature rules, where the integral in the direction θ is performed with a one-dimensional Gauss-Legendre quadrature (Hildebrand 1956), while the integral in ϕ can be evaluated, for example, with the trapezoidal rule. As it is apparent from the left panel of Fig. 2, the Gaussian-Legendre approach generates an accumulation of points near the north and south poles of the sphere, and it is less efficient than the other two quadratures, in the sense that it

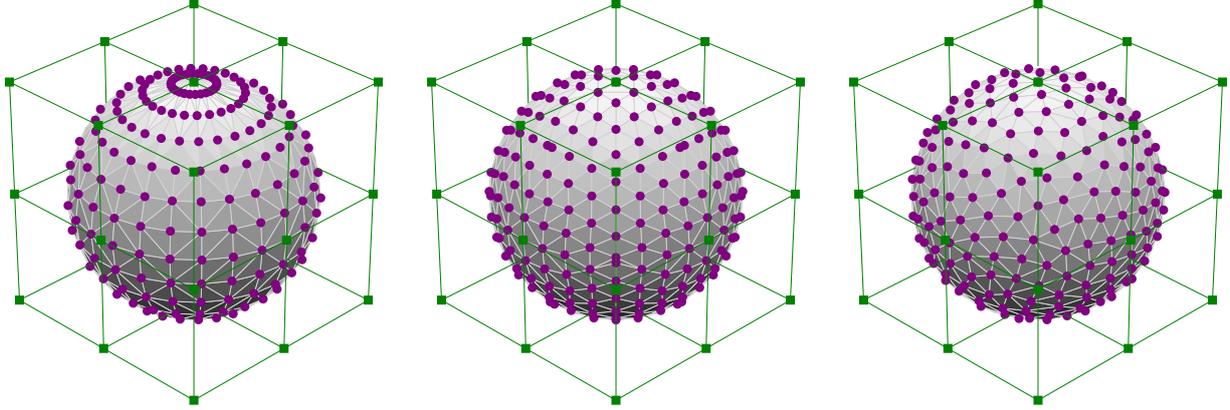


Figure 2. Examples of three-dimensional velocity stencils, comparing the distribution of the velocity directions (purple dots) for three different quadrature typologies; from left to right we respectively show an example for Gauss-Legendre, Lebedev, and spherical design. Analogous to Fig. 1 the purple dots are located on a sphere of radius $c\Delta t$ and the green squares denote the cell centres.

requires a larger number of points to achieve the same order of precision.

The Lebedev quadrature (central panel of Fig. 2), on the other hand, follows a different approach: instead of considering the product of two single quadratures, the integrals in Eq. (35) are used to build a system of nonlinear equations in the variables $\{w_k, \theta_k, \phi_k\}$. The central intuition (due to Sobolev (1962)) is that the number of nonlinear equations can be greatly reduced by considering only the spherical harmonics of degree $\leq p$ that exhibit invariance under all transformations that belong to a pre-determined group G . This procedure generates quadratures that are invariant under G , that is, these stencils evaluate exactly both $Q(h)$ and $Q(g(h))$ for all elements $g \in G$, and still retain the same order of integration p . In its original definition Lebedev's quadrature is by construction invariant under the octahedral group (Lebedev 1975, 1976, 1977), although quadratures based on different symmetry groups are also present in the literature (see, e.g., (Ahrens & Beylkin 2009) for a quadrature based on the icosahedral group). From the central panel of Fig. 2 one appreciates that points in Lebedev's quadratures offer a more homogeneous distribution over the sphere with respect to the Gauss product rule. Nevertheless, a few points can be seen to be almost overlapping.

Lastly, we consider the spherical-design quadrature rules, first introduced by Delsarte et al. (1977). This type of quadrature requires that the weights associated to all nodes be equal. The task is then to define a minimum set of points which integrates correctly all the spherical harmonics up to order p . There is no known rule for the definition of a generic order p spherical-design quadrature, and the topic is indeed still object of ongoing research. Nevertheless, numerical results leading to the definition of quadrature rules up to very high order are available online. In this work we refer to the set of stencils presented by Womersley (2018), for which we provide an example in the right panel of Fig. 2.

3.3 Numerical procedure

Having presented the equations to be solved and their discretization, we proceed to summarise the steps required to evolve the LB-discretized RTE.

More specifically, Eq. (29) can be solved following the standard stream-and-collide approach, where the streaming step is performed first, yielding provisional values of the N_{pop} intensities, i.e.,

$$I_i^*(\mathbf{r}, t) = I_i(\mathbf{r} - \hat{\mathbf{n}}_i \Delta t, t). \quad (39)$$

However, since the discrete velocities fall off-grid for the above discussed stencils, an interpolation is required. For simplicity we consider a trilinear interpolation scheme, from which follows

$$I_i^*(\mathbf{r} - \hat{\mathbf{n}}_i \Delta t, t) = \frac{1}{\Delta x \Delta y \Delta z} \times \left\{ \begin{aligned} &I_i(\mathbf{r} - \hat{\mathbf{x}} - \hat{\mathbf{y}} - \hat{\mathbf{z}}, t) \left(\Delta t |\hat{n}_i^x| \right) \left(\Delta t |\hat{n}_i^y| \right) \left(\Delta t |\hat{n}_i^z| \right) \\ &I_i(\mathbf{r} - \hat{\mathbf{y}} - \hat{\mathbf{z}}, t) \left(\Delta x - \Delta t |\hat{n}_i^x| \right) \left(\Delta t |\hat{n}_i^y| \right) \left(\Delta t |\hat{n}_i^z| \right) \\ &I_i(\mathbf{r} - \hat{\mathbf{x}} - \hat{\mathbf{z}}, t) \left(\Delta t |\hat{n}_i^x| \right) \left(\Delta y - \Delta t |\hat{n}_i^y| \right) \left(\Delta t |\hat{n}_i^z| \right) \\ &I_i(\mathbf{r} - \hat{\mathbf{x}} - \hat{\mathbf{y}}, t) \left(\Delta t |\hat{n}_i^x| \right) \left(\Delta t |\hat{n}_i^y| \right) \left(\Delta z - \Delta t |\hat{n}_i^z| \right) \\ &I_i(\mathbf{r} - \hat{\mathbf{z}}, t) \left(\Delta x - \Delta t |\hat{n}_i^x| \right) \left(\Delta y - \Delta t |\hat{n}_i^y| \right) \left(\Delta t |\hat{n}_i^z| \right) \\ &I_i(\mathbf{r} - \hat{\mathbf{y}}, t) \left(\Delta x - \Delta t |\hat{n}_i^x| \right) \left(\Delta t |\hat{n}_i^y| \right) \left(\Delta z - \Delta t |\hat{n}_i^z| \right) \\ &I_i(\mathbf{r} - \hat{\mathbf{x}}, t) \left(\Delta t |\hat{n}_i^x| \right) \left(\Delta y - \Delta t |\hat{n}_i^y| \right) \left(\Delta z - \Delta t |\hat{n}_i^z| \right) \\ &I_i(\mathbf{r} - \hat{\mathbf{z}}, t) \left(\Delta x - \Delta t |\hat{n}_i^x| \right) \left(\Delta y - \Delta t |\hat{n}_i^y| \right) \left(\Delta z - \Delta t |\hat{n}_i^z| \right) \end{aligned} \right\}, \quad (40)$$

with

$$\hat{\mathbf{x}} = [\text{sgn}(\hat{n}_i^x) \Delta x, 0, 0]^T \quad (41)$$

$$\hat{\mathbf{y}} = [0, \text{sgn}(\hat{n}_i^y) \Delta y, 0]^T \quad (42)$$

$$\hat{\mathbf{z}} = [0, 0, \text{sgn}(\hat{n}_i^z) \Delta z]^T. \quad (43)$$

For the two dimensional case a bilinear interpolation is used equivalently.

Next, the macroscopic moments are computed according to Eq. (27), from which the source term and the collision operator are computed using Eq. (30).

The collision step is then performed as follows:

$$I_i(\mathbf{r}, t + \Delta t) = I_i^*(\mathbf{r}, t) - \kappa_0 \Delta t [I_i^*(\mathbf{r}, t) - I_i^{\text{eq}}(\mathbf{r}, t)] + \Delta t S_i(\mathbf{r}, t). \quad (44)$$

Note that since in most astrophysical applications the emissivities and opacities vary over several orders of magnitude, these source terms may take values much larger than the evolved variable I_i , thus making the RTE a stiff equation that requires special treatment to be solved efficiently (see, e.g., Weih et al. 2020b).

Thus, rather than solving the explicit Eq. (44), we solve the implicit form, i.e.,

$$I_i(\mathbf{r} + \hat{\mathbf{n}}_i \Delta t, t + \Delta t) = I_i^*(\mathbf{r}, t) - \kappa_0 \Delta t [I_i(\mathbf{r}, t + \Delta t) - I_i^{\text{eq}}(\mathbf{r}, t + \Delta t)] + \Delta t S_i(\mathbf{r}, t + \Delta t), \quad (45)$$

Because of the presence of the first two moments in the definition of the source term, we obtain a system of N_{pop} linear equations with N_{pop} unknowns. This could be solved via the inversion of an $N_{\text{pop}} \times N_{\text{pop}}$ -matrix at every grid-point, which is, however, computationally unfeasible, especially in 3D, when a large number N_{pop} of discrete velocity directions is required. An alternate and computationally much more viable method is known as *Lambda iteration* (Rampp 1997), which works as follows:

1. Take as initial guess for $E(\mathbf{r}, t + \Delta t)$ and $\mathbf{F}(\mathbf{r}, t + \Delta t)$ the moments computed from I_i at time t .
2. Eq. (45) constitutes a system of N_{pop} linear decoupled equations that can be solved analytically and independently, to obtain a first estimate of I_i at time $t + \Delta t$.
3. Use this estimate of $I_i(\mathbf{r}, t + \Delta t)$ to formulate an improved guess for $E(\mathbf{r}, t + \Delta t)$ and $\mathbf{F}(\mathbf{r}, t + \Delta t)$.
4. Cycle back to step 2. and repeat until all $I_i(\mathbf{r}, t + \Delta t)$ converge, with an error below a prescribed threshold.

To summarise, our scheme consists of two steps; first the streaming-step according to Eq. (40) and then the collision-step according to Eq. (44), which is solved following the iterative procedure described above. If the radiation is coupled to a fluid, the radiative four-force that enters the standard equations of relativistic (magneto-) hydrodynamics (RMHD) can be computed after every timestep from the moments E and \mathbf{F} , which themselves are computed approximately from I_i according to Eq. (27). At the beginning of the next timestep, the coefficients η , κ_a , κ_0 and κ_1 can then be computed from the updated fluid variables.

As a result, the coupling to a standard RMHD code is exactly² the same as for the commonly used M1 scheme (see Weih et al. 2020b, for a detailed description of this coupling). Finally, we note that while we restrict ourselves for simplicity to a first-order time-stepper, it is conceptually straightforward to extend the evolution to higher orders using for example implicit-explicit (IMEX) schemes (Pareschi & Russo 2005), where the streaming-step is treated explicitly and the collision-step implicitly.

² Special attention has to be paid to the case of a moving background fluid (see Sec. 5 and appendix B).

4 NUMERICAL TESTS: STATIC FLUID

In astrophysical simulations, the ordinary fluid interacting with the radiation features optical depths that vary considerably, ranging from the optically thin regime – where radiation is in free streaming – to the optically thick regime – where radiation is coupled with the fluid and propagates by diffusion. While in several studies only one of these regimes is considered, (see e.g., Fragile et al. 2012; Roedig et al. 2012), it is our aim to provide a numerical method that can handle both limits, as well as the intermediate regime. The latter is particularly difficult to be described accurately by moment-based schemes, mostly because the closure relation – which is essential and inevitable in this schemes – is normally defined in either the optically thin or the optically thick regime and is then interpolated between these two limits (Weih et al. 2020b).

In the series of tests presented below, we discuss the performance of the LB method in these different limits, starting in Sec. 4.1 with the optically thin one – which represents the most difficult challenge. The optically thick regime is tested in Sec. 4.2 – where the LB method performs extremely well. Finally, in Sec. 4.2.2, we present an example of the intermediate regime, where LB is shown to provide very accurate results, at variance with the commonly used M1 scheme.

4.1 Optically thin limit

Since the LB method is designed for collisional fluids, we begin the verification of the LB method and its implementation with a number of code tests in the most difficult regime, namely, the one in which the radiation is actually freely streaming, as is the case when $\eta = \kappa_a = \kappa_0 = \kappa_1 = 0$.

In other words, in a radiative-transport application, the LB method works best when the emission, absorption and scattering of the radiation with an underlying fluid is actually taking place, which obviously is not the case in the free-streaming regime. Nevertheless, since free-streaming is ultimately taking place in any astrophysical scenario of interest, such as supernova explosions and neutron-star mergers – where the radiation is composed of neutrinos – it is clear that testing the LB method in this regime is most important.

4.1.1 Beam tests

We start with a “classical” beam test, namely, the propagation of a well defined beam of radiation injected from the boundary of the computational domain. Having set $\eta = \kappa_a = \kappa_0 = \kappa_1 = 0$, Eq. (45) clearly states that free-streaming is trivially achieved when each of the intensities I_i is propagated from one grid-cell to the next following the underlying stencil. The top-left panel of Fig. 3 shows an example of a freely streaming beam on a grid of size $-0.5 < x < 0.5$, $-0.5 < y < 0.5$, which we cover with 100^2 equal-size grid cells. This test is performed in 2D using the stencil shown in Fig. 1 with $N_{\text{pop}} = 8$ discrete velocity directions. Throughout the simulation we constantly inject a radiation beam from the grid’s left boundary. To do so, at all times we enforce the following condition:

$$I_i = \begin{cases} 1 & i = 0 \\ 0 & i \neq 0, \end{cases} \quad (46)$$

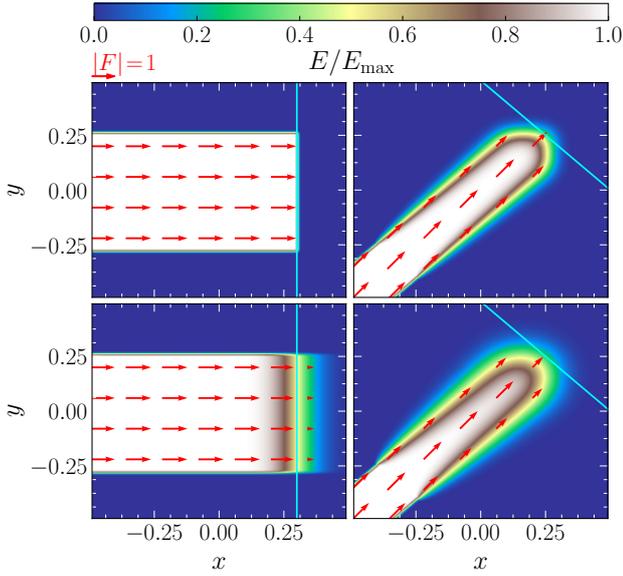


Figure 3. *Left:* Straight beam of radiation with a CFL-number of 1.0 (top) and a CFL-number of 0.2 (bottom). Colour coded is the energy-density, while the momentum density is shown by red arrows. The cyan line indicates, how far the beam should have propagated until $t = 0.7$. *Right:* The same as the left panel, but for a beam propagating diagonally. The region with $x, y < -0.25$ is frozen via a boundary condition in order to continuously shoot the beam into the grid from the bottom left.

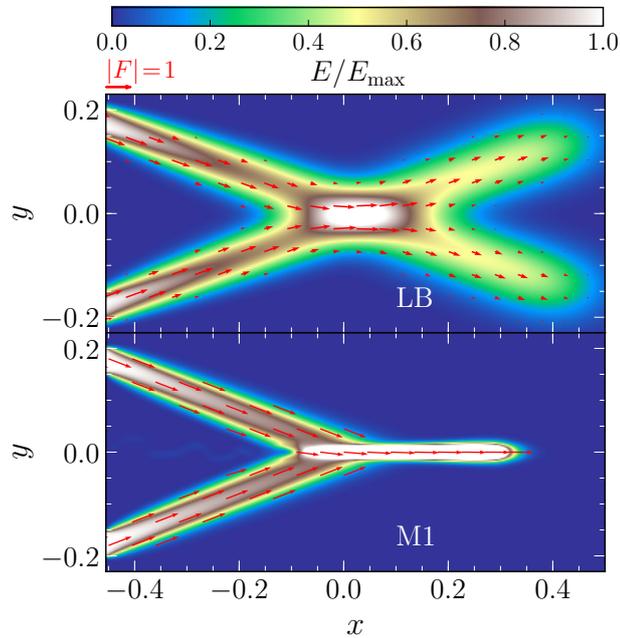


Figure 4. Comparison of the performance for the beam-crossing problem between LB (top) and M1 (bottom).

at $x = -0.5$ and $|y| < 0.25$.

As expected, the beam propagates parallel to the x -axis from left to right at the speed of light. In Fig. 3 we show with a colorcode the radiation energy density and with arrows the momentum density, as computed according to Eq. (27) at time $t = 0.7$, and where we have used a timestep $\Delta t = \Delta x = \Delta y$. This timestep is a standard choice in classical LB methods and leads to perfect streaming, i.e., no diffusion along the beam’s direction of propagation. This is due to the fact that the streaming step, i.e., Eqs. (39) and (40), reduces to

$$I_i^*(\mathbf{r}, t + \Delta t) = I_i(\mathbf{r} - \hat{\mathbf{n}}_i \Delta t, t) = \frac{|\hat{n}_i^x|}{\Delta x} \times I_i(\mathbf{r} - \hat{\mathbf{x}}, t). \quad (47)$$

Considering that $|\hat{n}_i^x|/\Delta x = 1$, where $|\hat{n}_i^x|$ is the x -th component of the i -th velocity vector, we find that the intensity is simply propagated from one cell to its right neighbour during each iteration.

Most of the astrophysical codes solving the equations of RMHD employ either finite-volume or finite-differencing schemes and require a timestep $\Delta t = \text{CFL} \Delta x$ for numerical stability, where $\text{CFL} \leq 1$ is the Courant-Friedrichs-Lewy coefficient (Rezzolla & Zanotti 2013). Keeping in mind that for simulations of astrophysical systems, the LB method needs to be coupled to such a RMHD code, we also perform the above beam tests with $\Delta t = 0.2\Delta x$, where $\text{CFL} = 0.2$ is a typical value chosen in multidimensional relativistic hydrodynamics. This amounts to squeezing the stencil and requires the complete evaluation of Eq. (40) during the streaming step.

The result of this simulation is reported in the bottom left panel of Fig. 3. In contrast to the case of perfect streaming reported in the top panel, the bottom panel shows that there is some diffusion of radiation ahead of the beam (see, for comparison, Fig. 1 of Weih et al. (2020b) for the same behaviour in an M1 code). Note that even though the previous test is a physically trivial one, it is nonetheless very useful to verify the correct implementation of the streaming step and the interpolation, according to Eq. (40).

A more challenging setup is that of a beam propagating along a direction not parallel to any coordinate axis. To model this case, we choose the same setup as above, but enforcing I_1 , which corresponds to $\hat{\mathbf{n}}_1 = [\cos(\pi/4), \sin(\pi/4)]$, rather than I_0 to a nonzero value; of course $I_{i \neq 1} = 0$.

Results are presented in the right panels of Fig. 3, where we show again the cases $\Delta t = \Delta x$ (top panel) and $\Delta t = 0.2\Delta x$ (bottom panel). In both cases, the beam now diffuses much more, an effect that can also be observed for the commonly used moment-schemes (Weih et al. 2020b). It is worth remarking that in all of the above tests, the beams propagate along one of the discrete velocity directions. However, radiation beams can propagate in any direction in the continuum, not necessarily along a discretized velocity direction. In this case, the LB scheme would inevitably incur increasing errors as we approach the optically thin regime. These can be tamed by developing high-order phase-space interpolators, possibly involving non-local neighbours both in configuration and velocity space. Clearly, this exposes a tension between accuracy and efficiency which still needs to be explored and resolved in full. In this paper, we rely upon trilinear and nearest-neighbour interpolation in configuration and velocity space, respectively.

Next, we consider the performance of our code for another “classical” and yet fundamental free-streaming test: two crossing beams. Assuming the radiation to consist of photons or neutrinos of the same flavour, one would expect the two beams to cross each other without interacting. The M1 scheme is known to perform very poorly under these conditions (Fragile et al. 2014; McKinney et al. 2014; Foucart et al. 2015; Rivera-Paleo & Guzmán 2019; Weih et al. 2020b), due to the fact that it retains only the two lowest moments, hence only the average direction of propagation.

This information could in principle be obtained by employing higher moments, at the cost of increased computational costs. On the other hand, in an LB method the various directions of propagation are evolved separately and thus crossing beams can be correctly evolved in time. This is shown in Fig. 4, where the results from an LB simulation (top panel) are compared to the ones from an M1 code (bottom panel). In both cases, we perform the simulation on a grid of size $-0.5 < x < 0.5$ and $-0.25 < y < 0.25$ with a resolution of 200×100 and choose again a CFL coefficient of 0.2. For the M1 solution, we initialise the momentum densities of the two beams to point towards the center, while for LB we simply initialise the intensities for the corresponding directions.

It can be seen clearly from Fig. 4 that in the LB simulation the two radiation beams cross as expected, while they merge to an averaged beam when using M1. The proper description of this behaviour is of crucial importance in astrophysical simulations, where beams of radiation – such as those emitted from the torus of a binary neutron-star merger remnant or in a supernova explosion – are expected to meet and interact. Hence, the successful outcome of this test provides encouraging evidence that the present LB method can be used in the place of moment-based schemes.

4.1.2 Radiation wave in free-streaming regime

In order to evaluate if radiation is propagated isotropically by the numerical scheme, we show here the results obtained for the case of a spherically symmetric propagation of a radiation wave. The wave is expected to expand at the speed of light in all directions, with its energy density decreasing over time following an inverse-square law for the distance travelled by the wavefront. We run all simulations on a 3D grid with 200^3 uniform sized grid cells with $\Delta t = 0.2$, and initialise the intensities by setting $I_i = 1$ within a sphere of radius $R = 16\Delta x$ grid-cells and $I_i = 0$ everywhere else. Stencils based on the three quadrature methods introduced in Sec. 3.2 are employed, with different quadrature orders, so as to compare both the methods and the various orders.

Figure 5 shows the results of these simulations in terms of the radiation energy-density, after 200 iterations, in the (x, y) and (y, z) planes. It is clear that a low number of discrete velocities N_{pop} in the stencils does not allow the radiation to stream isotropically. However, upon increasing N_{pop} , the isotropy of the system rapidly improves up to very satisfactory levels. Also to be noted, none of the three methods emerges as a neat winner, all yielding results of similar quality. This is true even for the Lebedev quadrature, which features a consistently higher quadrature order as the other two types for the same N_{pop} .

Finally, we show that the wave’s maximum energy de-

cays proportionally to r^{-2} , which is expected for this 3D test, where the energy is conserved and spreads over a spherical shell of area πr^2 in time. Figure 6 shows this behaviour for the reference case of the Lebedev quadrature of order $p = 23$ (middle panel in Fig. 5). The figure reports the profile of the radiation energy density along a diagonal cut at different times. Note that the maxima of these cuts align well with an inverse square law (red-dashed line).

4.2 Optically thick limit

Having tested the LB method for the solution of the RTE in the optically thin regime, we now focus on the optically thick regime. We recall that an optically thick medium either absorbs (in the case of a high value of the absorption opacity κ_a) or scatters radiation (in the case of a high scattering opacities κ_0 and κ_1). As a result, we present tests of the absorption scenario in Secs. 4.2.1 and 4.2.2 and the scattering scenario, i.e., the diffusion limit, in Sec. 4.2.3.

4.2.1 Shadow test

We start again with a beam using a similar setup as for the straight beam in Sec. 4.1.1, but we place an optically thick obstacle on the beam path. “Thick obstacle” means that we set the absorption opacity to a large value in the region occupied by the obstacle. From a physical point of view, the obstacle can then be viewed as a black-body absorbing the incoming radiation without re-emitting it ($\eta = 0$).

More specifically, on a 3D grid of size $-0.75 < x < 0.75$, $-0.25 < y, z < 0.25$ and covered with $150 \times 50 \times 50$ uniform sized grid cells, we set $\kappa_a = 10^7$ within a sphere of radius $R = 0.1$, centred at the origin. The beam is again injected into the grid from the left boundary. To this end, the 3D stencil must be chosen so as to provide a discrete velocity direction parallel to the x -axis. All other directions do not matter, since the radiation only propagates parallel to the x -axis in this test. Figure 7 shows the beam in the (x, y) plane after it has propagated across the whole grid and passed the obstacle marked as a cyan-coloured circle.

It can be seen that, as expected, the beam is blocked by the obstacle and only a negligible amount of radiation diffuses inside the optically thick region. This is due to the finite resolution of the grid and is progressively suppressed as the grid is refined.

We should note that, in contrast with all the tests performed so far, the set of equations that we solve here is very stiff because of the high numerical value of κ_a (as compared to the evolved variables). Obtaining a numerically stable solution with a reasonable timestep (we here use again $\Delta t = 0.2\Delta x$) is then only possible when using an implicit solver in time. As described in Sec. 3.3, we have implemented the *Lambda iteration*-method. It is found that, at least for this simple test, it converges in at most three cycles at every timestep, within a tolerance of $\Delta E/E = 10^{-14}$.

4.2.2 Radiating sphere

We next consider a test involving the emission of radiation as it occurs for a background static fluid in thermodynamic equilibrium with the radiation, that is, when the emissivity

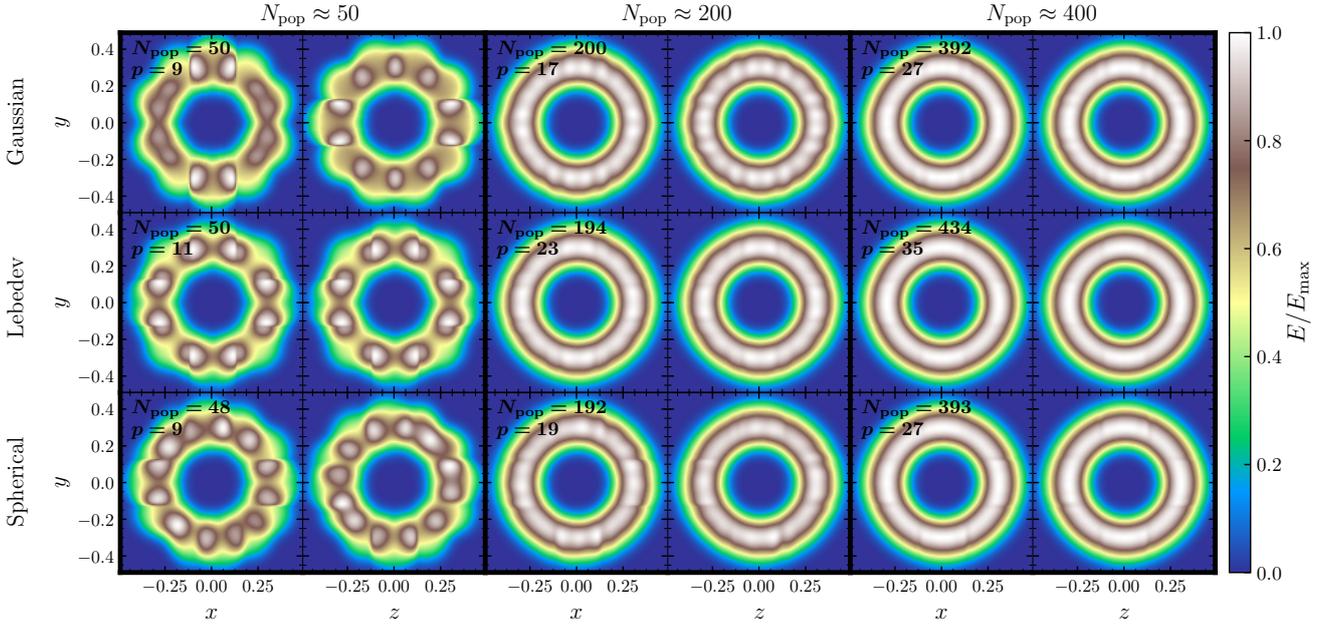


Figure 5. Spherical freely-streaming wave for different types of stencils. From top to bottom we show stencils derived from (i) Gaussian product quadrature, (ii) Lebedev quadrature and (iii) spherical design. Each column compares these three stencils for a comparable number of discrete velocities. Snapshots show the radiation energy-density in the (x, y) (left) and (y, z) (right) plane after 200 iterations. In every panel we also report N_{pop} and the quadrature order p .

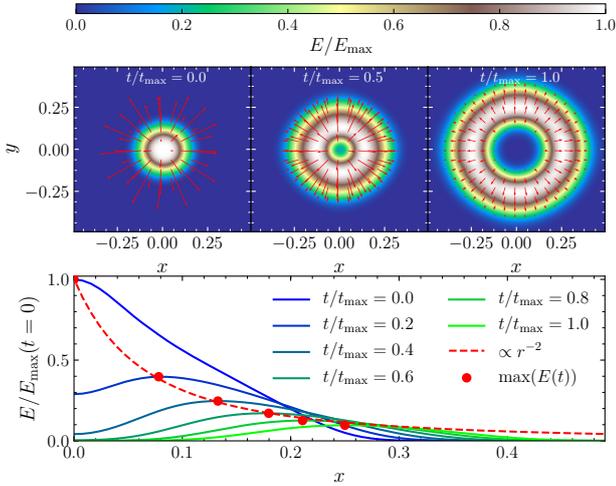


Figure 6. *Top:* Radiation energy density (colour coded) and momentum density (red arrows) at three representative times in the (x, y) plane for the radiation wave test with the Lebedev quadrature of order $p = 23$. *Bottom:* Diagonal profiles of the energy density at different times (blue to green) and inverse square law (red-dashed) fitted to the maxima of each time (red dots).

η is equal to the absorption opacity κ_a . More specifically, we simulate a system in which $\eta = \kappa_a = \text{const.}$ within a sphere of radius R and zero outside the sphere. This is known as the homogeneous-/radiative-/emitting-sphere test, as first proposed by [Smit et al. \(1997\)](#). From a physical point of view, this system can be thought of as a dense sphere with a sharp boundary to vacuum, as is the case of a neutron

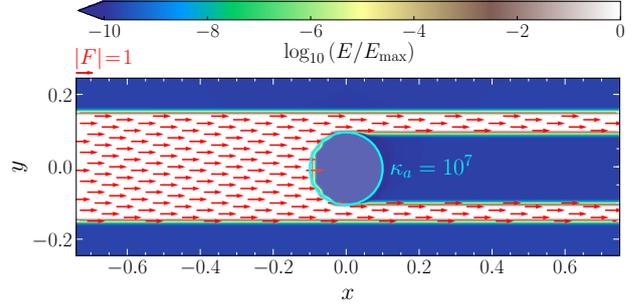


Figure 7. Radiation beam hitting an optically thick obstacle (cyan circle) with $\kappa_a = 10^7$. The colour encodes the radiation energy-density and the red arrows show magnitude and direction of the energy momentum-density.

star, that constantly emits radiation from its surface to the surrounding vacuum.

The system eventually finds a steady state, for which the analytical solution of the distribution function in terms of radial distance r and azimuthal angle θ reads as follows:

$$f(r, \mu) = b(1 - e^{-\kappa_a s(r, \mu)}), \quad (48)$$

where $\mu := \cos \theta$, $b = \kappa_a / \eta = 1$ in our case, and

$$s := \begin{cases} r\mu + Rg(r, \mu) & r < R \text{ and } -1 < \mu < 1, \\ 2Rg(r, \mu) & r \geq R \text{ and } \sqrt{1 - R^2/r^2} < \mu < 1, \end{cases} \quad (49)$$

with

$$g(r, \mu) := \sqrt{1 - \frac{r^2}{R^2}(1 - \mu^2)}. \quad (50)$$

Since the final equilibrium is spherically symmetric, the

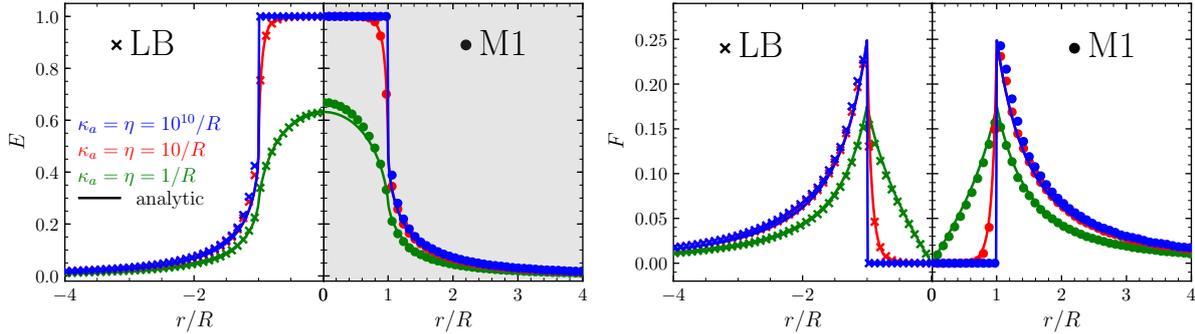


Figure 8. Diagonal cuts for the radiation energy-density (left) and the magnitude of the momentum-density (right) of the equilibrium for the radiating-sphere test. We compare the numerical solution obtained with LB (crosses; left halves) and M1 (circles; right halves) with the analytic solution (solid lines.) for a low (green), moderate (red) and high (blue) value of $\kappa_a = \eta$.

zeroth moment E can be obtained via the integration of the distribution function over μ , i.e.,

$$E(r) = \frac{1}{2} \int_{-1}^1 d\mu f(r, \mu). \quad (51)$$

Likewise, we obtain the first moment F as

$$F(r) = \frac{1}{2} \int_{-1}^1 d\mu \mu f(r, \mu). \quad (52)$$

We run this test using a 3D grid³ with 128^3 uniform sized grid cells and a spherical-design stencil of order 20 ($N_{\text{pop}} = 222$) for three different values of κ_a , i.e., $\kappa_a = R^{-1}$, $10R^{-1}$, $10^{10}R^{-1}$, and where $R = 1/(8n_x)$ is the sphere’s radius.

The discrete intensities are initialised as

$$I_i(r, t = 0) := w_i \begin{cases} 1 & r < R, \\ r^{-2} & r \geq R, \end{cases} \quad (53)$$

where w_i are the weights associated with the underlying velocity stencil.

The results of these simulations are reported in Fig. 8 in terms of the radiation energy-density (left panel) and of the momentum-density (right panel).

As one can appreciate, the LB method works very well in all of the three cases. As already discussed for the previous shadow test, numerical stability in the case of $\kappa = 10^{10}/R$ is only possible thanks to the implicit time-stepper.

Figure 8 also offers a comparison with the corresponding results obtained with an M1 scheme (these are shown in the right portions of the two panels in Fig. 8). It is clear that for small values of κ_a , the LB method performs significantly better than M1. As already reported by Weih et al. (2020b), the M1 code fails in this specific case due to the lack of the second moment, i.e., the correct pressure tensor. While the pressure tensor is exact in the limit of infinite optical depths (see red and blue curves in Fig. 8), it is not so for the intermediate regime between optically thick and thin media. Indeed, the case of $\kappa_a = R^{-1}$ (green curves) falls

exactly in this regime, for which the pressure tensor is interpolated inaccurately (see also Murchikova et al. (2017) for a detailed analysis of this test for the M1 scheme with different closures). The failure of the M1 scheme in this test is particularly evident upon looking at the energy density, which is systematically above the correct analytical solution inside the sphere (note that the tail, which is within the free-streaming regime is reproduced correctly). The same – albeit less visible – is true for the corresponding momentum density, which is everywhere below the analytic solution.

4.2.3 Radiation wave in scattering regime

As a final test of the solution of the RTE with the LB method, we consider the effects of scattering by simulating a Gaussian distribution that diffuses over a static background fluid. The same test has been used also in the validation of various M1 codes (see, e.g., Pons et al. 2000; O’Connor 2015; Weih et al. 2020b). The initial conditions are given by:

$$E(r, t = 0) = A \exp\left(-\frac{(r - r_0)^2}{2\sigma_0^2}\right), \quad (54)$$

where A is the energy-density amplitude and σ_0 the width of the Gaussian at $t = 0$ centered at r_0 . By neglecting emission and absorption ($\eta = \kappa_a = 0$), the evolution of the system is governed by the diffusive equation, which exhibits the following analytic solution:

$$E(r, t) = A \frac{\sigma_0^2}{\sigma_0^2 + \sigma_D^2} \exp\left(-\frac{(r - r_0)^2}{2(\sigma_0^2 + \sigma_D^2)}\right), \quad (55)$$

where $\sigma_D = \sqrt{2Dt}$. The connection between the microscopic parameters κ_0 and κ_1 in Eq. (17) and the diffusion coefficient D is discussed in Appendix A by an asymptotic analysis.

We perform two different simulations at different values of the ‘‘Peclet number’’ ($\text{Pe} := \kappa_0 \Delta x$), tracking the history of the system at different time intervals. In Fig. 9, we compare the analytical solution (lines) with the numerical results (dots) obtained running in 2D, on a 100^2 grid, with $\kappa_0 \Delta x = 1$ (left panel) and $\kappa_0 \Delta x = 10^5$ (right panel). In both cases, we set $\lambda := 3\kappa_1/\kappa_0 = 0.5$. The same test can be performed in three dimensions, leading to results of similarly good quality. Finally, in Fig. 10 we show that, in stark contrast with the results shown for the free-streaming regime

³ As already remarked by Radice et al. (2013) and Weih et al. (2020b), this test requires three dimensions on a Cartesian grid, since fluxes also propagate across grid-cells in the angular directions.

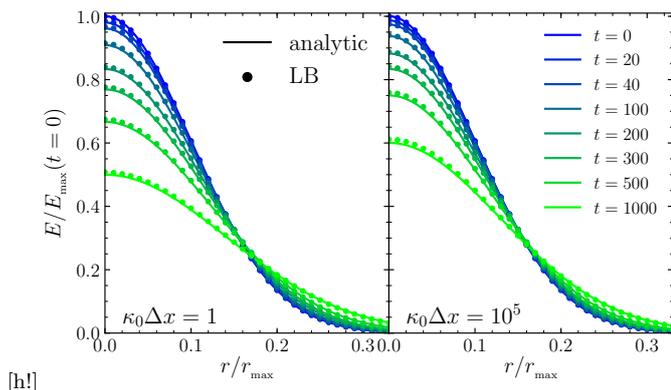


Figure 9. Gaussian wave packet diffusing in a static fluid. The analytic solution (solid lines) is compared with results of simulations (dots) obtained on a two-dimensional grid of size 100^2 , with $\Delta t = 0.1$, and $\lambda = 0.5$. The two panels show results at low and high Peclet numbers.

(Fig. 5), the evolution is well captured even by quadratures of comparatively low orders, featuring a small number of discrete directions. This highlights the fact that the LB method performs extremely well in the diffusion limit, in fact the one it was born for. In general, we have observed that in this regime quadratures with order $p = 5$ are sufficient to correctly recover the correct diffusive dynamics.

5 NUMERICAL TESTS: DYNAMICAL FLUID

After having shown that the proposed radiative LB method performs well in all the regimes of the solution of the RTE on a static fluid, we next move on to test the method for a dynamical fluid, i.e. one coupled to radiation via an RMHD simulation of a realistic astrophysical scenario. As a result, the test presented here is of great importance as it allows us to explore our novel approach under conditions typical of the astrophysical scenarios for which it has been developed in the first place.

We should also note that, strictly speaking, this is not a genuine test, as the explored scenario does not have an analytic solution to be compared with. Moreover, simulations of this type have been performed before only by [Rivera-Paleo & Guzmán \(2019\)](#) using an M1 scheme for evolving the radiation but with rather different prescriptions for the properties of the radiation field.

Hence, to gain confidence on the reliability of our results and contrast them with those obtained when the RTE is solved using an M1 scheme, we carry out additional simulations of the identical physical scenario but making use of the M1 code [FRAC \(Weih et al. 2020b\)](#). While this does not prove the correctness of our results – both codes could be incorrect despite the many tests passed – it does provide us with the confidence necessary to implement the LB method for even more realistic astrophysical scenarios.

More specifically, we simulate the evolution of a relativistic jet as it propagates through the interstellar medium. Simulations of this type have a long history since such jets are of major importance in the study of active galactic nuclei (see [Perucho 2019](#), for a recent review), where they are

produced through the accretion process onto a supermassive black hole (see [Porth et al. 2019](#), for a recent comparative study).

Highly energetic relativistic jets are known to accompany the phenomenology of short gamma-ray bursts and are associated with the merger of two magnetised neutron stars ([Rezzolla et al. 2011](#)).

Here, we simulate this problem by coupling our LB code to the *Black hole accretion code* (BHAC) ([Porth et al. 2017](#)). BHAC is a finite-volume code that solves the equations of general-relativistic MHD in a fixed and curved spacetime. The results presented here, however, are restricted to a flat background spacetime.

In essence, we perform the coupling between BHAC and the LB code following the strategy indicated below.

1. At every iteration, we pass the fluid rest-mass density ρ , temperature T and three-velocity v_i to the LB code, from which we then compute the emissivity and opacities (see also below).

2. While BHAC advances the conservative RMHD variables in time, the LB code does the same for the N_{pop} populations of the radiation specific intensity, I_i .

3. After performing the streaming and collision step, the LB code computes the zeroth (E), first (F^i) and second moment (P^{ij}) of the radiation distribution function according to Eqs. (27) and (28).

4. From these moments, we compute the radiative source terms $\mathcal{S}_0 = W(\tilde{\kappa}_a J - \tilde{\eta}) + \tilde{\kappa} H_0$ and $\mathcal{S}_j = W(\tilde{\kappa}_a J - \tilde{\eta})v_j + \tilde{\kappa} H_j$, which we return to BHAC. Here W is the Lorentz-factor, $\tilde{\kappa} := \tilde{\kappa}_a + \tilde{\kappa}_s := \tilde{\kappa}_a + (\tilde{\kappa}_0 - 1/3\tilde{\kappa}_1)$, and J and H_j are the radiation energy and momentum density in the comoving fluid frame, to which we transform via

$$J = W^2(E - 2F^i v_i + P^{ij} v_i v_j) \quad (56)$$

$$H_j = W^3(F^i v_i - E)v_j + W h_{ji} F^i - W h_{ji} v_k P^{ik}, \quad (57)$$

where $h_{ij} = W^2 v_i v_j + \delta_{ij}$ is the projection operator orthogonal to the fluid velocity. H_0 can then be computed from $H_\mu u^\mu = 0$.

5. After BHAC has updated its variables and before cycling to 1), we add the sources \mathcal{S}_0 and \mathcal{S}_j to the energy and momentum equation, respectively, which are solved in conservative form within BHAC.

As it is the case in most simulations of high-energy astrophysical phenomena, the fluid moves at relativistic speeds. Eq. (17), is written in an Eulerian (lab) frame so that the opacities and emissivities η , κ_a , κ_0 , κ_1 it employs are to be evaluated in the same Eulerian frame. However, the microphysics used to derive such quantities is well defined only in the fluid’s rest-frame (i.e., the frame co-moving with the fluid) where the corresponding quantities $\tilde{\eta}$, $\tilde{\kappa}_a$, $\tilde{\kappa}_0$, $\tilde{\kappa}_1$ are isotropic and can be written in a compact way. Therefore, care needs to be taken in transforming the opacities and emissivities between the two frames, as we discuss in detail in Appendix B.

For the setup of our simulation, we follow [Martí et al. \(1997\)](#), who have extensively analysed relativistic jets in a purely hydrodynamical context. The jet is simply injected through a circular nozzle at the lower edge of the computational domain and propagating parallel to the coordinate z -axis in a Cartesian grid. The simulation is then char-

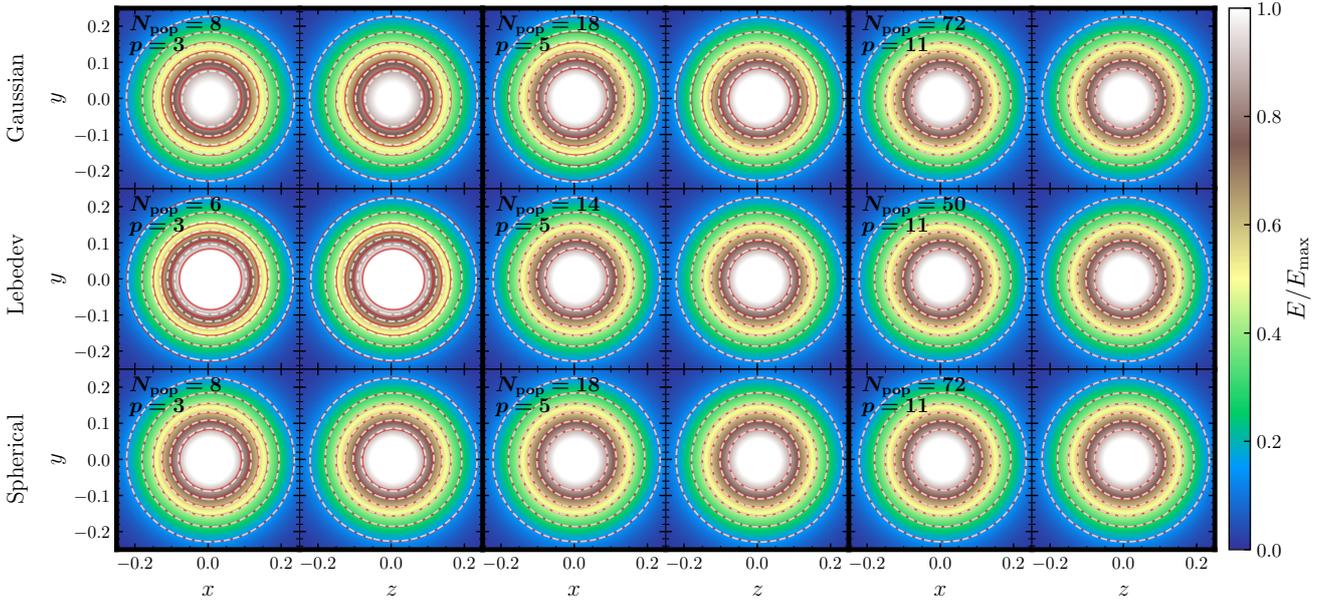


Figure 10. Scattering wave in 3D. The simulations are performed on a grid of size 128^3 , with $\Delta t = 0.1$, $\kappa_0 \Delta x = 1$ and $\lambda = -0.25$. From top to bottom we show the results obtained using stencils derived from (i) Gaussian product quadrature, (ii) Lebedev quadrature and (iii) spherical design. Each column compares these three stencils for a comparable number of discrete velocities. Snapshots show the radiation energy-density in the (x, y) (left) and (y, z) (right) plane after $100/\Delta t$ iterations. Contour lines are used to compare the analytical solution (continuous red lines) with the numerical results (white dotted lines). In each panel, we also report N_{pop} and the quadrature order p .

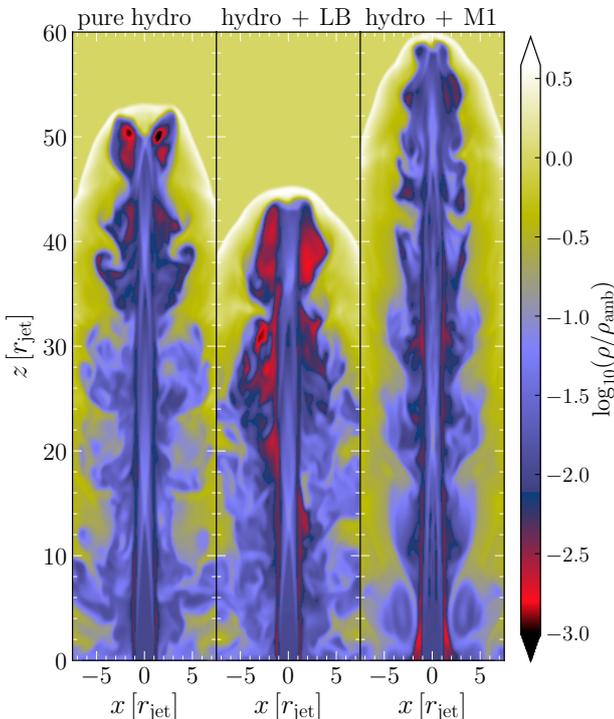


Figure 11. Cut through the (x, z) plane for the relativistic jet after $t = 125 r_{\text{jet}}$. Shown is the rest-mass density for the pure-hydro (left) and the coupled hydro-radiation using LB (middle) and M1 (right).

acterized by four parameters: the Newtonian Mach number $\mathcal{M} := v_{\text{jet}}/c_s$, with c_s the local sound speed, the jet Lorentz factor $W := (1 - v_{\text{jet}}^2)^{-1/2}$ with v_{jet} the jet propagation velocity, the ratio of the jet rest-mass density to that of the ambient medium $\mathcal{R} := \rho_{\text{jet}}/\rho_{\text{amb}}$ and the pressure ratio $\mathcal{K} := p_{\text{jet}}/p_{\text{amb}}$. We simulate a pressure matched jet, i.e., $\mathcal{K} = 1$, with $\mathcal{R} = 0.01$, $W = 7$ and $\mathcal{M} = 42$ on a grid of size $-7.5 < x, y < 7.5$, $0 < z < 60$ covered by $160 \times 160 \times 640$ grid cells, where the jet is injected at $z = 0$ through a nozzle with radius $r_{\text{jet}} = 1$.

Since we are only interested in a proof-of-concept simulation, we limit ourselves to this simple setup, but refer the reader to [Fromm et al. \(2018\)](#) for an extension that also includes a non-homogeneous background, as is to be expected near a gravitational source like a black hole. However, we include a helical perturbation as in [Aloy et al. \(1999\)](#), so that the jet deviates from axisymmetry, leading to a truly three-dimensional structure that allows us to test all terms of the LB code. In all cases, we consider the magnetic field to be zero.

For the solution of the RTE within our LB scheme we choose a Lebedev stencil with $N_{\text{pop}} = 154$ discrete velocity directions. We initialise the populations at zero and let the radiation evolve self-consistently during the simulation. To this purpose, we set $\tilde{\kappa}_a = \rho^2 T^{-3.5}$ and $\tilde{\eta} = \sigma_{\text{SB}}/\pi \tilde{\kappa}_a T^4$, where the fluid temperature T is computed from an ideal gas equation of state. This absorption opacity is motivated by the Rosseland mean opacity for thermal bremsstrahlung ([Rybicki & Lightman 1986](#)) and the emissivity simply follows from Kirchhoff's law upon assuming black-body radiation.

We should note that in contrast to the pure-RMHD simulation, which is scale invariant, the coupled RMHD-radiation simulation fixes the length scale via the value of

the Stefan-Boltzmann constant σ_{SB} and the density assumed for the ambient medium ρ_{amb} . Here, we simply use $\rho_{\text{amb}} = 1$ and $\sigma_{\text{SB}} = 0.1$, which does not lead to a physically realistic setup, but ensures that a moderate amount of radiation is produced, that neither dominates the fluid nor is dominated by it.

Finally, we also add scattering using $\tilde{\kappa}_0 = 10^{-3}\rho$ and $\tilde{\kappa}_1 = 0$. This choice is motivated by the microphysical process of Thomson scattering, which is proportional to the number of scatters in the medium (hence, the choice for $\tilde{\kappa}_0$) and has no preferred direction (hence, the choice for $\tilde{\kappa}_1$).

We compare the results of the pure-RMHD and the RTE-coupled simulations in Fig. 11, whose left panel refers to the pure-RMHD jet, the central panel to the RTE-coupled solution obtained with the LB method, and the right panel to the corresponding evolution when the RTE-coupled solution is obtained with the M1 scheme. A quick comparison of the pure-RMHD jet morphology in the left panel shows that it is in good agreement with the one presented by [Martí et al. \(1997\)](#) and [Aloy et al. \(1999\)](#), where this type of jets has been studied extensively. On the other hand, the solutions employing a coupling with the radiation are considerably different. In particular, the RTE-coupled simulation with the LB method shows that the jet propagates more slowly, i.e., the Lorentz factor is $\sim 15\%$ smaller than for the pure-RMHD case. This is due to the fact that the fluid making up the jet loses energy via the emission of radiation. By stark contrast, the RTE-coupled solution obtained with the M1 scheme shows that the jet propagates more rapidly. While this behaviour is similar to the one reported by [Rivera-Paleo & Guzmán \(2019\)](#) – who, in addition, continuously injected energy in the radiation field – we believe it is actually incorrect.

The origin of this substantially different dynamical behaviour is due to the poor handling by the M1 scheme of radiation interacting with itself. We find this to cause significantly different distributions of the energy-density. In the case of the M1-evolution, the radiation energy density accumulates mostly in a narrow region along the z -axis, while it is more evenly spread in the case of LB-evolution. The actual cause of this radiation focussing is the same at the origin of the poor performance of the M1 scheme in the beam-crossing test reported in Fig. 4. Also in this case, in fact, different beams of radiation originating from the recollimation shock produced near the injection region of the jet ([Mizuno et al. 2015](#)), intersect along the z -axis and lead to an inconsistent combination of radiation fluxes.

This is illustrated in the top panels of Fig. 12, which reports a cut through the (x, z) plane of the radiation energy density E at an early time during the jet evolution, namely at $t = 30 r_{\text{jet}}$, using either the LB method (left) or the M1 scheme (right). Note that in both cases there is a triangular region of small E originating from the recollimation shock, and ultimately due to the contact discontinuity between propagating jet and ambient medium (see [Aloy & Rezzolla 2006](#), for a discussion of the role of the contact discontinuity in accelerating the jet). This region of small E is surrounded on both sides by high- E streams, which appear as white in the colorcode used. In the case of the M1-evolution (right panel), these two streams merge at the top of the triangular low- E region leading to the same pathology discussed in Sec. 4.1.1 (cf., Fig. 4) and to the artificial

accumulation of radiation along the z -axis. The latter then provides additional momentum, pushing the jet forward and overcompensating the linear momentum lost in the production of the radiation. In the LB-evolution, on the other hand, the two beams do not merge but cross correctly. As a result, the radiation energy density is not artificially focused and E spreads out over a larger region, leading to a broader bow shock ahead of the jet, which is also propagating more slowly.

Note that these problems in the M1-evolution affect the dynamics of the jet only downstream of the first recollimation shock. This is very clearly illustrated in the bottom panels of Fig. 12, which refer instead to the rest-mass densities in the two cases and to the same time as the top panels, i.e., $t = 30 r_{\text{jet}}$. As one would expect, the differences are in this case very small since the radiative effects have not yet been able to play a role, which they will instead do for $t \gtrsim 30 r_{\text{jet}}$.

In summary, these results go well-beyond our intention of providing a proof-of-principle evidence for a correct implementation of the LB method in a fully coupled relativistic-fluid configuration. In particular, they clearly show the ability of the LB method to handle correctly scenarios with physical conditions that are very close to those encountered in relativistic astrophysical phenomena. More importantly, they highlight that under these very same conditions, the M1 approach commonly employed – even by us ([Weih et al. 2020a](#)) – may suffer from artefacts that may affect the dynamics of relativistic jets, such as those expected to be generated in a binary system of merging neutron stars, and whose accurate description is essential to gain insight in the launching of relativistic jets in short gamma-ray bursts.

6 PERFORMANCE AND SCALABILITY

Novel computational methods are expected to i) feature good scalability on parallel computers, ii) perform better than existing ones in accuracy, efficiency or both.

In what follows we show that the LB method proposed here possesses both these qualities, i.e. parallel scalability and efficiency.

Indeed, one key ingredient in the success and widespread adoption of LB is parallel efficiency. Thanks to the synchronous algorithmic flow stemming from the stream-collide paradigm, LB schemes exhibit high amenability to parallel coding, making them natural targets for efficient implementations on modern computer architectures ([Godenschwager et al. 2013](#); [Bernaschi et al. 2009](#); [Mazzeo & Coveney 2008](#); [Bernaschi et al. 2010](#); [Succi et al. 2019](#)).

In this section, we provide a brief performance evaluation of the numerical scheme designed in this paper. With respect to classical LB models, our scheme presents two main differences: *i*) we employ off-lattice quadratures, which require interpolation to implement the streaming phase; *ii*) the number of velocity components forming the stencil is significantly larger with respect to standard LB schemes.

We have implemented our numerical scheme using directive-based programming environments, such as *OpenMP* and *OpenACC*, to expose parallelism ([Calore et al. 2016](#)). The advantage of this approach is that the code is portable and can therefore be compiled and executed on diverse architectures, from commodity CPU-based processors to GPU accelerators. We test the code on an Intel Skylake

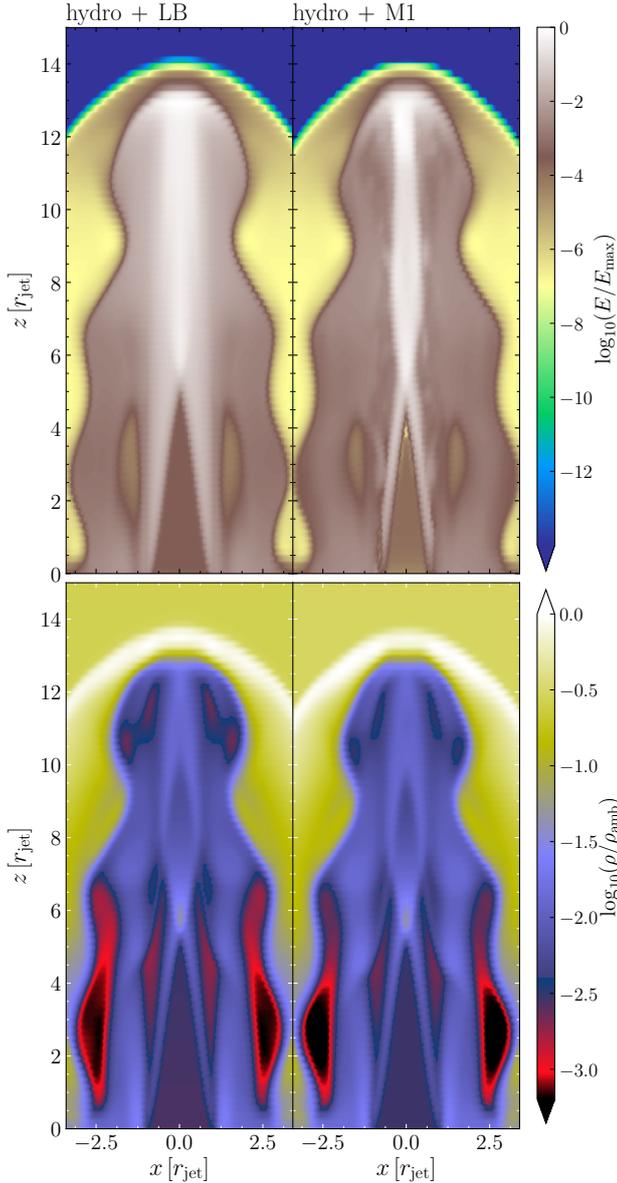


Figure 12. *Top panel:* Cut through the (x, z) plane of the radiation energy density at $t = 30r_{\text{jet}}$ for the coupled RMHD-radiation simulations using the LB method (left) and the M1 scheme (right). *Bottom panel:* The same as in the top panel, but for the rest-mass density.

20-cores processor, a commonly adopted architecture in the HPC-market. We measure performances in terms of Million Lattice Updates per Second (MLUPS):

$$\text{MLUPS} = 10^6 \frac{L^3 N_{\text{iter}}}{t_{\text{exe}}}, \quad (58)$$

where t_{exe} measures the execution time (in seconds) required to simulate N_{iter} timesteps on a grid of L^3 points. In essence, for a problem of fixed size and iterations L and N_{iter} , a parallel implementation should lead to an execution time that decreases linearly with the number of cores and hence a linearly growing MLUPS. This is shown in Fig. 13, where we assess the scalability of the code on a single Intel Skylake board. We solve the emitting-sphere benchmark with

an explicit Euler stepper, on an $L = 128^3$ grid, and with a spherical-design quadrature of order $p = 20$ with $N_{\text{pop}} = 200$ discrete components. The code scales up to 20 threads with a parallel efficiency above 70%.

These figures are already quite good, but could be further improved. A limiting performance factor is given by inefficient memory accesses, which in turn leads to a sub-optimal use of the cache and of the vector unit of the processor. Omitting details, we simply mention here that a careful optimisation of the data-layout used to store the grid in memory should be taken into consideration. The two most common data layouts used in several stencil applications are the so-called array of structures (AoS) and structure of arrays (SoA) schemes; in the AoS layout, all populations associated to one lattice site are stored in contiguous memory locations. Conversely, in the SoA scheme all the populations having the same index i are stored contiguously, while populations belonging to the same lattice site are stored far from each other at non unit-stride addresses. Recently, more sophisticated data-layouts have been designed explicitly for LB applications (Shet et al. 2013) yielding strong benefits, especially when targeting multi-core architectures with wide vector units and large cache memories (Calore et al. 2019). The application of these optimisation and a more in-depth analysis will be reported elsewhere.

In Fig. 14, we show the performance obtained on a dual-Skylake processor as a function of N_{pop} . We show results for an explicit and also for the implicit stepper described in Sec. 3.3, which extends the stability of the method also to the case of very stiff terms. To simplify the comparison, we have executed the *Lambda iteration* loop exactly 5 times at each grid-cell. Since the execution time of the current implementation is completely bounded by memory accesses, the difference between the implicit and the explicit scheme is almost negligible.

The figure also presents a comparison with the performance of the special-relativistic version of the M1 code FRAC (Weih et al. 2020b), featuring a similar level of optimisation. For M1 the major bottleneck is the closure relation, which requires a complex root-finding algorithm for obtaining the pressure tensor. In principle, the performance of FRAC can be arbitrarily good or bad depending on the desired accuracy of the root-finding step. Here we consider a standard setup (i.e., a relative accuracy of 10^{-9} for root-finding and computing the closure on cell-centres as well as cell-faces). As can be seen from the dashed line in Fig. 14, the LB method outperforms the M1 code for $N_{\text{pop}} \lesssim 400$, the latter being a rather generous number of discretized directions to solve most astrophysical problems. We also consider somewhat higher (10^{-14}) and lower (10^{-4}) relative accuracy for the root-finding, which is shown as a blue band in Fig. 14.

Finally, we comment that the LB method is particularly well suited for GPU implementations (Bernaschi et al. 2010; Calore et al. 2017). For this reason we also test a GPU-optimized version of our code on an NVIDIA V100 GPU. We observe the performances to be systematically one order of magnitude higher than those reported on the CPU, both for stencils with 10 – 100 components, suitable for simulations in which scattering terms are relevant, as well as for stencils with > 200 components, which instead allow to extend the applicability of the solver to free-streaming regimes. The present version of the code allows for efficient memory ac-

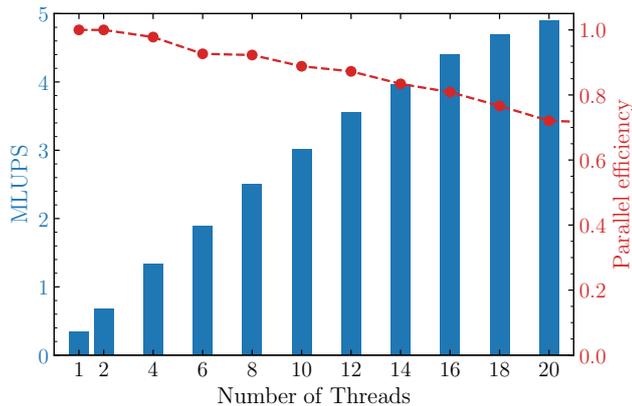


Figure 13. Analysis of performance scalability on a single Intel Skylake board. The results refer to simulations of the emitting-sphere benchmark with an explicit Euler stepper, on an $L = 128^3$ grid, using a stencil with $N_{\text{pop}} = 200$ components. The figures reported correspond to the best out of 10 trials for each set of parameters.

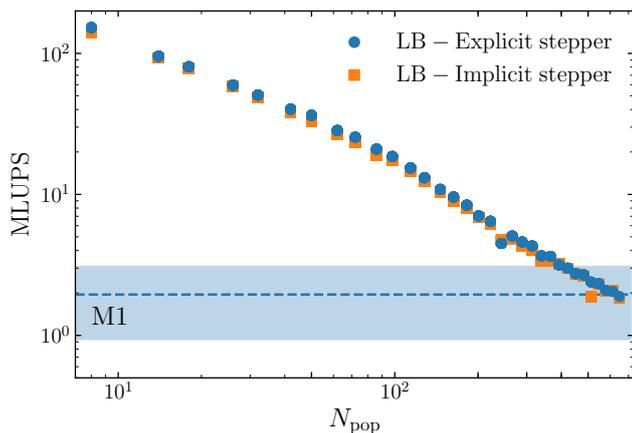


Figure 14. Comparison of the performance (measured MLUPS) achieved on a dual-Skylake processor (40 threads) as a function of N_{pop} , the number of the discrete components forming the stencil. The results refer to simulations of the emitting-sphere benchmark, on an $L = 128^3$ grid. Dots correspond to simulations based on an explicit Euler stepper, while squares correspond to simulations using an implicit stepper. The dashed line shows the performance of the M1 code `FRAC` with the blue-shaded band corresponding to different accuracies for the M1 closure. The figures reported are the best out of 10 trials for each set of parameters.

cesses on GPU architectures, in turn exposing the higher costs of the implicit solver, which for large values of N_{pop} is found to be 2–2.5× slower than the explicit Euler method.

7 CONCLUSIONS AND OUTLOOK

We have presented an extension of the LB method, which is commonly used in classical fluid dynamics, to the solution of the RTE in special relativity, thus making it applicable to three-dimensional simulations of high-energy astrophysical phenomena.

After implementing the new method in flat spacetime and under the grey approximation, we analysed its perfor-

mance in a number of code tests. In this way, we have shown that while the LB method performs extremely well in the diffusion limit, the free-streaming regime represents the most difficult one to treat accurately. In this regime, a large number of discrete velocity directions is required since the propagation of radiation is restricted to these directions only. Nevertheless, even in this optically thin limit, the LB method proves superior to the commonly used M1 method, as witnessed by the fact that radiation beams cross correctly.

This feature is important for astrophysical systems with an accretion disk and torus, from which radiation is expected to cross and lead to the phenomenology observed in short gamma-ray bursts. Besides this important advantage, the new LB method also outperforms the M1 method in the intermediate regime, between diffusion and free-streaming, mostly because it does not need to rely on a closure relation to compute higher moments, as it is the case for the M1 scheme, where the pressure tensor is simply interpolated via the closure.

The accuracy in the calculation of the moments also depends on the underlying quadrature, for which we have compared three possibilities. While in our code tests only minor differences could be seen among these quadratures, it is mathematically clear that the Lebedev and spherical-design quadratures are more accurate than the Gauss-Legendre quadrature commonly used in direct Boltzmann solvers (see e.g., Nagakura et al. 2018).

We have also coupled our new radiation code to the GRMHD-code `BHAC` and simulated a relativistic jet, where the radiation back-reacts dynamically onto the fluid. Not only does this represent the first such simulation using an LB scheme, but also proves that our new method is indeed applicable to high-energy astrophysics. Furthermore, when comparing the corresponding results obtained with the more standard M1 code `FRAC`, that employs a moment-based scheme, we have shown that the LB-method solution does not suffer from the inaccuracies that plague the M1 method.

Finally, we have also shown that the LB method is faster than the M1 method for a number of discrete directions $N_{\text{pop}} \lesssim 400$, which has been shown to be more than enough for accurately simulating the diffusion limit and gives reasonable results also in the free-streaming regime.

Depending on the system, one might need higher accuracy in the latter and thus $N_{\text{pop}} \approx 600 - 800$ might be necessary. In this case the LB method is slightly more expensive than the M1 scheme. Considering, however, the high amenability of LB to GPU implementations, a major speed-up can certainly be achieved along this line.

While this paper is meant to provide a presentation of the LB method for the solution of the radiative-transfer equation in computational astrophysics, a number of improvements are possible, both in terms of astrophysical applications, and in terms of mathematical and numerical developments. The former involves a more detailed and realistic investigation of the role played by radiation in impacting the dynamics and imaging of astrophysical relativistic jets. The latter necessarily involves the extension of the method to a general-relativistic framework, which requires the extension of the streaming-step to curved spacetimes.

Finally, the numerical developments will include the possibility of using numerical grids with various form of static and dynamic mesh-refinements. Also in this case, an

adjustment of the streaming-step will be needed to allow for the streaming from coarse to fine grid cells and vice-versa. Work along these lines is in progress.

ACKNOWLEDGMENTS

LRW acknowledges support from HGS-HIRE. AG acknowledges funding by "Contributo 5 per mille assegnato all'Università degli Studi di Ferrara-dichiarazione dei redditi dell'anno 2017". DS has been supported by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 765048. SS acknowledges funding from the European Research Council under the European Union's Horizon 2020 framework programme (No. P/2014-2020)/ERC Grant Agreement No. 739964 (COPMAT). He also wishes to thank A. Ferrara, P. Mocz, D. Spergel and J. Stone for illuminating discussions. Support also comes in part from "PHAROS", COST Action CA16214; LOEWE-Program in HIC for FAIR; the ERC Synergy Grant "BlackHoleCam: Imaging the Event Horizon of Black Holes" (Grant No. 610058). The simulations were performed on the SuperMUC and SuperMUC-NG clusters at the LRZ in Garching, on the LOEWE cluster in CSC in Frankfurt, on the HazelHen cluster at the HLRS in Stuttgart, and on the COKA computing cluster at Università di Ferrara.

REFERENCES

- Ahrens C., Beylkin G., 2009, *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 465, 3103
- Aloy M. A., Rezzolla L., 2006, *Astrophys. J.*, 640, L115
- Aloy M. A., Ibáñez J. M. ., Martí J. M. ., Gómez J.-L., Müller E., 1999, *Astrophys. J.*, 523, L125
- Ambruş V. E., Blaga R., 2018, *Phys. Rev. C*, 98, 035201
- Asinari P., Mishra S., Borchellini R., 2010, *Numerical Heat Transfer, Part B: Fundamentals*, 57, 126
- Atkinson K., Han W., 2012, *Spherical Harmonics and Approximations on the Unit Sphere: An Introduction*. Springer Berlin Heidelberg, doi:10.1007/978-3-642-25983-8, <http://dx.doi.org/10.1007/978-3-642-25983-8>
- Baiotti L., Rezzolla L., 2017, *Rept. Prog. Phys.*, 80, 096901
- Beentjes C. H. L., 2015, *International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering*, pp 1–15
- Bernaschi M., Melchionna S., Succi S., Fyta M., Kaxiras E., Sircar J., 2009, *Computer Physics Communications*, 180, 1495
- Bernaschi M., Fatica M., Melchionna S., Succi S., Kaxiras E., 2010, *Concurrency and Computation: Practice and Experience*, 22, 1
- Bhatnagar P. L., Gross E. P., Krook M., 1954, *Phys. Rev.*, 94, 511
- Bindra H., Patil D. V., 2012, *Phys. Rev. E*, 86, 016706
- Bovard L., Martin D., Guercilena F., Arcones A., Rezzolla L., Korobkin O., 2017, *Phys. Rev. D*, 96, 124005
- Broadwell J. E., 1964, *The Physics of Fluids*, 7, 1243
- Bruenn S. W., 1985, *Astrophys. J., Supp.*, 58, 771
- Calore E., Gabbana A., Kraus J., Schifano S. F., Tripicciono R., 2016, *Concurrency and Computation: Practice and Experience*, 28, 3485
- Calore E., Gabbana A., Kraus J., Pellegrini E., Schifano S. F., Tripicciono R., 2017, arXiv e-prints, p. arXiv:1703.00185
- Calore E., Gabbana A., Schifano S., Tripicciono R., 2019, *The International Journal of High Performance Computing Applications*, 33, 124
- Cardall C. Y., Endeve E., Mezzacappa A., 2013, *Phys. Rev. D*, 87, 103004
- Chapman S., Cowling T. G., 1970, *The Mathematical Theory of Non-Uniform Gases*, 3rd ed. Cambridge University Press, doi:10.1119/1.1942035
- Coelho R. C., Mendoza M., Doria M. M., Herrmann H. J., 2018, *Computers & Fluids*, 172, 318
- Delsarte P., Goethals J. M., Seidel J. J., 1977, *Geometriae Dedicata*, 6, 363
- Dietrich T., Ujevic M., 2017, *Classical and Quantum Gravity*, 34, 105014
- Dünweg B., Ladd A. J. C., 2009, *Lattice Boltzmann Simulations of Soft Matter Systems*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp 89–166, doi:10.1007/978-3-540-87706-6_2
- Event Horizon Telescope Collaboration et al., 2019, *Astrophys. J. Lett.*, 875, L1
- Fernández R., Tchekhovskoy A., Quataert E., Foucart F., Kasen D., 2019, *Mon. Not. R. Astron. Soc.*, 482, 3373
- Foucart F., 2018, *Mon. Not. R. Astron. Soc.*, 475, 4186
- Foucart F., et al., 2015, *Phys. Rev. D*, 91, 124021
- Fragile P. C., Gillespie A., Monahan T., Rodriguez M., Anninos P., 2012, *Astrophys. J., Supp.*, 201, 9
- Fragile P. C., Olejar A., Anninos P., 2014, *Astrophys. J.*, 796, 22
- Frisch U., Hasslacher B., Pomeau Y., 1986, *Phys. Rev. Lett.*, 56, 1505
- Fromm C. M., Perucho M., Porth O., Younsi Z., Ros E., Mizuno Y., Zensus J. A., Rezzolla L., 2018, *Astron. Astrophys.*, 609, A80
- Fujibayashi S., Kiuchi K., Nishimura N., Sekiguchi Y., Shibata M., 2018, *Astrophys. J.*, 860, 64
- Gabbana A., Mendoza M., Succi S., Tripicciono R., 2018, *Computers & Fluids*, 172, 644
- Gabbana A., Simeoni D., Succi S., Tripicciono R., 2020, *Physics Reports*, 863, 1
- Gairola A., Bindra H., 2017, *Annals of Nuclear Energy*, 99, 151
- Galeazzi F., Kastaun W., Rezzolla L., Font J. A., 2013, *Phys. Rev. D*, 88, 064009
- Gamba I. M., Haack J. R., Hauck C. D., Hu J., 2017, *SIAM Journal on Scientific Computing*, 39, B658
- Godenschwager C., Schornbaum F., Bauer M., Köstler H., Rüdiger U., 2013, in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis. SC 13*. Association for Computing Machinery, New York, NY, USA, doi:10.1145/2503210.2503273
- Grad H., 1949a, *Communications on Pure and Applied Mathematics*, 2, 325
- Grad H., 1949b, *Communications on Pure and Applied Mathematics*, 2, 331
- Gross B., Atzberger P., 2018, *Journal of Computational Physics*, 371, 663
- Hardy J., Pomeau Y., de Pazzis O., 1973, *Phys. Rev. Lett.*, 31, 276
- He X., Chen S., Doolen G. D., 1998, *Journal of Computational Physics*, 146, 282
- Higuera F. J., Jiménez J., 1989, *EPL (Europhysics Letters)*, 9, 663
- Higuera F. J., Succi S., 1989, *Europhysics Letters (EPL)*, 8, 517
- Higuera F. J., Succi S., Benzi R., 1989, *EPL (Europhysics Letters)*, 9, 345
- Hildebrand F. B., 1956, *Introduction to numerical analysis*. New York : McGraw-Hill, <https://trove.nla.gov.au/work/3478395>
- Janka H.-T., Langanke K., Marek A., Martínez-Pinedo G., Müller B., 2007, *Physics Reports*, 442, 38

- Just O., Obergaulinger M., Janka H. T., 2015, *Mon. Not. R. Astron. Soc.*, **453**, 3386
- Karlin I. V., Sichau D., Chikatamarla S. S., 2013, *Phys. Rev. E*, **88**, 063310
- Krüger T., Kusumaatmaja H., Kuzmin A., Shardt O., Silva G., Vigger E. M., 2017, *The Lattice Boltzmann Method*. Springer International Publishing, doi:10.1007/978-3-319-44649-3, http://dx.doi.org/10.1007/978-3-319-44649-3
- Kuroda T., Takiwaki T., Kotake K., 2016, *Astrophys. J., Supp.*, **222**, 20
- Lebedev V., 1975, *USSR Computational Mathematics and Mathematical Physics*, **15**, 44
- Lebedev V., 1976, *USSR Computational Mathematics and Mathematical Physics*, **16**, 10
- Lebedev V. I., 1977, *Siberian Mathematical Journal*, **18**, 99
- Levermore C. D., Pomraning G. C., 1981, *Astrophys. J.*, **248**, 321
- Lutsko J. F., Lam J., 2018, *Phys. Rev. E*, **98**, 012604
- Martí J. M., Müller E., Font J. A., Ibáñez J. M., Marquina A., 1997, *Astrophys. J.*, **479**, 151
- Massioli F., Benzi R., Succi S., 1993, *Europhysics Letters (EPL)*, **21**, 305
- Mazzeo M., Coveney P., 2008, *Computer Physics Communications*, **178**, 894
- McCulloch R., Bindra H., 2016, *Computers & Fluids*, **124**, 261
- McHardy C., Horneber T., Rauh C., 2016, *Optics Express*, **24**, 16999
- McKinney J. C., Tchekhovskoy A., Sadowski A., Narayan R., 2014, *Mon. Not. R. Astron. Soc.*, **441**, 3177
- McNamara G. R., Zanetti G., 1988, *Phys. Rev. Lett.*, **61**, 2332
- Melon Fuksman J. D., Mignone A., 2019, *Astrophys. J., Supp.*, **242**, 20
- Mendoza M., Karlin I., Succi S., Herrmann H. J., 2013, *Phys. Rev. D*, **87**, 065027
- Mezzacappa A., Liebendörfer M., Messer O. E. B., Hix W. R., Thielemann F.-K., Bruenn S. W., 2001, *Phys. Rev. Lett.*, **86**, 1935
- Mihalas D., Auer L. H., 2001, *Journal of Quantitative Spectroscopy and Radiative Transfer*, **71**, 61
- Miller J. M., Ryan B. R., Dolence J. C., 2019, *Astrophys. J., Supp.*, **241**, 30
- Mink A., McHardy C., Bressel L., Rauh C., Krause M. J., 2020, *Journal of Quantitative Spectroscopy and Radiative Transfer*, **243**, 106810
- Mishra S. C., Poonia H., Vernekar R. R., Das A. K., 2014, *Heat Transfer Engineering*, **35**, 1267
- Mizuno Y., Gómez J. L., Nishikawa K.-I., Meli A., Hardee P. E., Rezzolla L., 2015, *Astrophys. J.*, **809**, 38
- Most E. R., Papenfort L. J., Rezzolla L., 2019, *Mon. Not. R. Astron. Soc.*, **490**, 3588
- Murchikova E. M., Abdikamalov E., Urbatsch T., 2017, *Mon. Not. R. Astron. Soc.*, **469**, 1725
- Nagakura H., et al., 2018, *ApJ*, **854**, 136
- O'Connor E., 2015, *Astrophys. J., Supp.*, **219**, 24
- Pareschi L., Russo G., 2005, *Journal of Scientific Computing*, **25**, 129
- Paschalidis V., 2017, *Classical and Quantum Gravity*, **34**, 084002
- Peng Y., Shu C., Chew Y. T., 2003, *Phys. Rev. E*, **68**, 026701
- Perego A., Rosswog S., Cabezón R. M., Korobkin O., Käppeli R., Arcones A., Liebendörfer M., 2014, *Mon. Not. R. Astron. Soc.*, **443**, 3134
- Perego A., Radice D., Bernuzzi S., 2017, *Astrophys. J. Lett.*, **850**, L37
- Perucho M., 2019, *Galaxies*, **7**, 70
- Philippi P. C., Hegele L. A., dos Santos L. O. E., Surmas R., 2006, *Phys. Rev. E*, **73**, 056702
- Pomraning G. C., 1981, *J. Quant. Spectrosc. Radiative Transfer*, **26**, 385
- Pons J. A., Ibáñez J. M., Miralles J. A., 2000, *Mon. Not. R. Astron. Soc.*, **317**, 550
- Porth O., Olivares H., Mizuno Y., Younsi Z., Rezzolla L., Moscibrodzka M., Falcke H., Kramer M., 2017, *Computational Astrophysics and Cosmology*, **4**, 1
- Porth O., Chatterjee K., Event Horizon Telescope Collaboration 2019, *Astrophys. J. Supp.*, **243**, 26
- Radice D., Abdikamalov E., Rezzolla L., Ott C. D., 2013, *Journal of Computational Physics*, **242**, 648
- Radice D., Galeazzi F., Lippuner J., Roberts L. F., Ott C. D., Rezzolla L., 2016, *Mon. Not. R. Astron. Soc.*, **460**, 3255
- Rahman N., Just O., Janka H. T., 2019, *Mon. Not. R. Astron. Soc.*, **490**, 3545
- Rampp M., 1997, PhD thesis, -
- Rezzolla L., Miller J. C., 1994, *Class. Quantum Grav.*, **11**, 1815
- Rezzolla L., Zanotti O., 2013, *Relativistic Hydrodynamics*. Oxford University Press, Oxford, UK, doi:10.1093/acprof:oso/9780198528906.001.0001
- Rezzolla L., Giacomazzo B., Baiotti L., Granot J., Kouveliotou C., Aloy M. A., 2011, *Astrophys. J. Letters*, **732**, L6
- Rivera-Palao F. J., Guzmán F. S., 2019, *Astrophys. J., Supp.*, **241**, 28
- Rivet J.-P., Boon J. P., 2001, *Lattice Gas Hydrodynamics*. Cambridge Nonlinear Science Series, Cambridge University Press, doi:10.1017/CBO9780511524707
- Roedig C., Zanotti O., Alic D., 2012, *Mon. Not. R. Astron. Soc.*, **426**, 1613
- Rosswog S., Liebendörfer M., 2003, *Mon. Not. R. Astron. Soc.*, **342**, 673
- Rosswog S., Korobkin O., Arcones A., Thielemann F.-K., Piran T., 2014, *Mon. Not. R. Astron. Soc.*, **439**, 744
- Ruffert M., Janka H.-T., Schaefer G., 1996, *Astron. Astrophys.*, **311**, 532
- Rybicki G. B., Lightman A. P., 1986, *Radiative Processes in Astrophysics*. Wiley-VCH
- Sądowski A., Narayan R., Tchekhovskoy A., Zhu Y., 2013, *Mon. Not. R. Astron. Soc.*, **429**, 3533
- Shan X., 2016, *Journal of Computational Science*, **17**, 475
- Shan X., He X., 1998, *Phys. Rev. Lett.*, **80**, 65
- Shan X., Yuan X.-F., Chen H., 2006, *Journal of Fluid Mechanics*, **550**, 413
- Shet A. G., Sorathiya S. H., Krithivasan S., Deshpande A. M., Kaul B., Sherlekar S. D., Ansumali S., 2013, *Phys. Rev. E*, **88**, 013314
- Shibata M., Kiuchi K., Sekiguchi Y., Suwa Y., 2011, *Progress of Theoretical Physics*, **125**, 1255
- Siegel D. M., Ciolfi R., 2016, *Astrophys. J.*, **819**, 14
- Siegel D. M., Metzger B. D., 2017, *Physical Review Letters*, **119**, 231102
- Skinner M. A., Dolence J. C., Burrows A., Radice D., Vartanyan D., 2019, *ApJS*, **241**, 7
- Smit J. M., Cernohorsky J., Dullemond C. P., 1997, *Astron. Astrophys.*, **325**, 203
- Sobolev S., 1962, *Soviet Mathematics Doklady*, **3**, 1307
- Stepán, Jirí Bestard, Jaume Jaume Trujillo Bueno, Javier 2020, *A&A*, **636**, A24
- Succi S., 2015, *EPL (Europhysics Letters)*, **109**, 50001
- Succi S., 2018, *The Lattice Boltzmann Equation: For Complex States of Flowing Matter*. OUP Oxford, doi:10.1093/oso/9780199592357.001.0001
- Succi S., Amati G., Bernaschi M., Falcucci G., Lauricella M., Montessori A., 2019, *Computers & Fluids*, **181**, 107
- Thorne K. S., 1981, *Mon. Not. R. Astron. Soc.*, **194**, 439
- Vernekar R. R., Mishra S. C., 2014, *International Journal of Heat and Mass Transfer*, **77**, 218
- Wang Y., Ma Y., Xie M., 2019, *Progress in Nuclear Energy*, **110**, 341

Weih L. R., Hanauske M., Rezzolla L., 2020a, *Phys. Rev. Lett.*, **124**, 171103

Weih L. R., Olivares H., Rezzolla L., 2020b, *MNRAS*, **495**, 2285

Womersley R. S., 2018, in , *Contemporary Computational Mathematics - A Celebration of the 80th Birthday of Ian Sloan*. Springer International Publishing, pp 1243–1285, doi:10.1007/978-3-319-72456-0_57

Yi H.-L., Yao F.-J., Tan H.-P., 2016, *Phys. Rev. E*, **94**, 023312
Zanotti O., Roedig C., Rezzolla L., Del Zanna L., 2011, *Mon. Not. R. Astron. Soc.*, **417**, 2899

d'Humières D., Lallemand P., Frisch U., 1986, *Europhysics Letters (EPL)*, **2**, 291

APPENDIX A: LINK BETWEEN MICROSCOPIC AND MACROSCOPIC PARAMETERS

In this section of the appendix we perform an asymptotic analysis to link the microscopic parameters with the macroscopic ones. In particular we consider the specific limit of zero emission and absorption ($\kappa_a = \eta = 0$), for which the radiative Lattice Boltzmann equation is shown to recover the diffusion equation

$$\partial_t E = D \Delta E. \quad (\text{A1})$$

In Appendix A1 the link between the diffusive coefficient D and the scattering parameters (κ_0, κ_1) is established analytically through a Chapman-Enskog expansion (Chapman & Cowling 1970). In Appendix A2 the analytic expressions are extended to account for extra numerical corrections by fitting results from numerical simulations.

A1 Chapman-Enskog Analysis

We start from Eq. (29) assuming zero absorption and emission, i.e.,

$$I_i(\mathbf{r} + c\hat{\mathbf{n}}_i \Delta t, t + \Delta t) - I_i(\mathbf{r}, t) = -c\kappa_0 \Delta t \left(I_i(\mathbf{r}, t) - I_i^{\text{eq}}(\mathbf{r}, t) \right). \quad (\text{A2})$$

with $I_i^{\text{eq}}(\mathbf{r}, t)$ given by Eq. (30). Note that throughout this section we will write c explicitly and write vector components with Greek indices rather than using boldface vectors. We also adopt the Einstein convention of summing over repeated indices.

Taking a Taylor expansion of the left-hand side of Eq. (A2), and including terms up to the second order, gives:

$$I_i(\mathbf{x} + \mathbf{n}_i \Delta t, t + \Delta t) - I_i(\mathbf{x}, t) = \Delta t \left(\partial_t + c n_i^\alpha \partial_\alpha \right) I_i + \frac{1}{2} \Delta t^2 \left(\partial_t + c n_i^\alpha \partial_\alpha \right)^2 I_i + \mathcal{O}(\Delta t^3). \quad (\text{A3})$$

We also expand the differential operator with respect to time

$$\partial_t = \epsilon \partial_t^{(1)} + \epsilon^2 \partial_t^{(2)} + \mathcal{O}(\epsilon^3), \quad (\text{A4})$$

and space:

$$\partial_\alpha = \epsilon \partial_\alpha^{(1)} + \mathcal{O}(\epsilon^2), \quad (\text{A5})$$

where $\epsilon \ll 1$.

Next, we expand the specific intensity around its equilibrium:

$$I_i = I_i^{(0)} + \epsilon I_i^{(1)} + \epsilon^2 I_i^{(2)} + \mathcal{O}(\epsilon^3), \quad (\text{A6})$$

where $I_i^{(0)} \equiv I_i^{\text{eq}}$. We also recall the definition of the first and second moment of the distribution:

$$E = \sum_i I_i, \quad (\text{A7})$$

$$F^\alpha = \sum_i n_i^\alpha I_i. \quad (\text{A8})$$

Assuming the most basic level of isotropy for the stencil used in the numerical method we have

$$\begin{aligned} \sum_i w_i &= 1, & \sum_i w_i n_i^\alpha &= 0, \\ \sum_i w_i n_i^\alpha n_i^\beta &= \frac{1}{d} \delta_{\alpha\beta}, & \sum_i w_i n_i^\alpha n_i^\beta n_i^\gamma &= 0, \end{aligned} \quad (\text{A9})$$

where α, β , and γ run over the spatial indexes in d dimensions.

By integrating Eq. (30), in combination with Eq. (A9), we get the following definitions for the moments of the equilibrium distribution:

$$\sum_i I_i^{\text{eq}} = E, \quad (\text{A10})$$

$$\sum_i I_i^{\text{eq}} n_i^\beta = \frac{\lambda}{d} F^\beta, \quad (\text{A11})$$

$$\sum_i I_i^{\text{eq}} n_i^\beta n_i^\gamma = \frac{1}{d} \delta_{\beta\gamma} E. \quad (\text{A12})$$

Since we are neglecting absorption and emission and consider the diffusion limit, where the radiation is in thermodynamic equilibrium with the underlying fluid, we find due to conservation

$$\int (I - I^{\text{eq}}) d\Omega = 0 \Rightarrow \int I d\Omega = \int I^{\text{eq}} d\Omega, \quad (\text{A13})$$

which in the discretized form leads to the condition

$$\sum_i I_i = \sum_i I_i^{\text{eq}} = E, \quad \sum_i I_i^{(k)} = 0 \quad \forall k \geq 1. \quad (\text{A14})$$

We do not write down any condition for the first moment, since its preservation depends on the choice of λ .

Replacing the left-hand side of Eq. (A2) with its second order Taylor expansion, i.e., Eq. (A3), we get

$$\begin{aligned} \Delta t \left(\partial_t + c n_i^\alpha \partial_\alpha \right) I_i(\mathbf{x}, t) + \frac{1}{2} \Delta t^2 \left(\partial_t + c n_i^\alpha \partial_\alpha \right)^2 I_i(\mathbf{x}, t) \\ = -c \Delta t \kappa_0 \left(I_i(\mathbf{x}, t) - I_i^{\text{eq}}(\mathbf{x}, t) \right). \end{aligned} \quad (\text{A15})$$

We now plug in the above Eqs. (A4), (A5), and (A6) and perform a multi-scale expansion in which we keep track separately of terms up to order ϵ and ϵ^2 . The resulting equations are

$$\mathcal{O}(\epsilon) : \left(\partial_t^{(1)} + c n_i^\alpha \partial_\alpha^{(1)} \right) I_i^{(0)} \approx -c \kappa_0 I_i^{(1)}, \quad (\text{A16})$$

$$\mathcal{O}(\epsilon^2) : \left(1 - \frac{\Delta t}{2} c \kappa_0 \right) \left(\partial_t^{(1)} + c n_i^\alpha \partial_\alpha^{(1)} \right) I_i^{(1)} + \partial_t^{(2)} I_i^{(0)} \approx -c \kappa_0 I_i^{(2)}. \quad (\text{A17})$$

We then take into consideration Eq. (A16) and integrate (i.e. sum over all N_{pop} populations), getting

$$\sum_i \left(\partial_t^{(1)} I_i^{(0)} + c n_i^\alpha \partial_\alpha^{(1)} I_i^{(0)} \right) \approx -c \kappa_0 \sum_i I_i^{(1)}, \quad (\text{A18})$$

which using Eq. (A14) and the definition of the first order moment leads to

$$\partial_t^{(1)} E + c \frac{\lambda}{d} \partial_\alpha^{(1)} F_\alpha = 0. \quad (\text{A19})$$

Next, starting again from Eq. (A16), we multiply by n_i^β and integrate, which yields

$$\frac{\lambda}{d} \partial_t^{(1)} F^\beta + \frac{1}{d} c \partial_\beta E = -c \kappa_0 \sum_i n_i^\beta I_i^{(1)}. \quad (\text{A20})$$

Note that the RHS of the above equation vanishes when the conservation of the first moment is ensured; for the moment we leave it in a general form and evaluate this term later on.

We now integrate Eq. (A17) and obtain

$$\sum_i \left(1 - \frac{\Delta t}{2} c \kappa_0 \right) \left(\partial_t^{(1)} + c n_i^\alpha \partial_\alpha^{(1)} \right) I_i^{(1)} + \sum_i \partial_t^{(2)} I_i^{(0)} \approx -\kappa_0 \sum_i I_i^{(2)}, \quad (\text{A21})$$

$$\left(1 - \frac{\Delta t}{2} c \kappa_0 \right) c \partial_\alpha^{(1)} \sum_i n_i^\alpha I_i^{(1)} + \partial_t^{(2)} E = 0, \quad (\text{A22})$$

$$\partial_t^{(2)} E = - \left(1 - \frac{\Delta t}{2} c \kappa_0 \right) c \partial_\alpha^{(1)} \sum_i n_i^\alpha I_i^{(1)}. \quad (\text{A23})$$

The RHS can be derived from Eq. (A20), leading to

$$\partial_t^{(2)} E = \frac{1}{c \kappa_0} \left(1 - \frac{\Delta t}{2} c \kappa_0 \right) c \partial_\alpha^{(1)} \left(\frac{\lambda}{d} \partial_t^{(1)} F^\alpha + \frac{1}{d} c \partial_\alpha E \right), \quad (\text{A24})$$

which can be re-arranged as

$$\partial_t^{(2)} E = D \left(\Delta^{(1)} E + \lambda \frac{1}{c} \partial_\alpha^{(1)} \partial_t^{(1)} F_\alpha \right) \quad (\text{A25})$$

with

$$D = \frac{c^2}{d c \kappa_0} \left(1 - \frac{\Delta t}{2} c \kappa_0 \right). \quad (\text{A26})$$

We note that Eq. (A25) is the diffusion equation with D its diffusion coefficient plus an extra error term. This error term is the same one that arises for the equations of the M1 system, when considering the optically thick limit and neglecting absorption and emission. It is remarkable that despite the M1 system being a set of macroscopic equations, where the connection between D and the system's scattering coefficient is directly evident, in contrast to the microscopic equation of our LB method, we arrive at the exact same macroscopic equation for the diffusion limit.

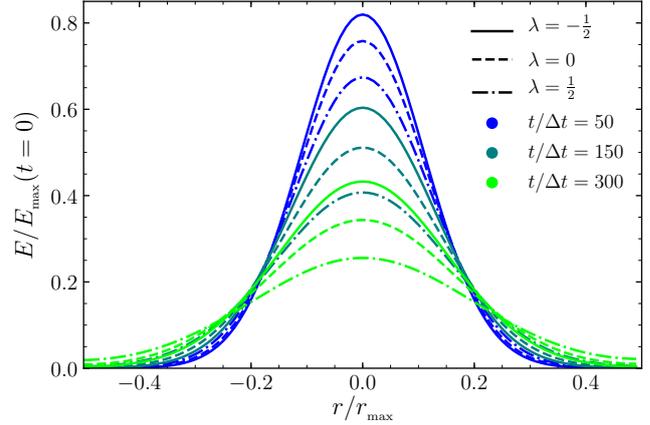


Figure A1. Effect of the parameter λ on the diffusion speed of a Gaussian pulse. The simulations are performed in a bidimensional grid of size $L = 100$, with $\Delta t = 1$ and $k_0 \Delta x = 1$.

A2 Numerical Fit

In the previous section we have established a link between the microscopic parameters and the diffusion coefficient by performing the Chapman-Enskog expansion. In principle one would need to extend the expansion to include corrections coming from higher order terms. Besides, one should also account for extra dissipative effects coming from the fact that the LB method described in the present work is not based on a space-filling Cartesian lattice and consequently requires interpolation. Since the overall analysis would become rather tedious from an analytical point of view, in this section we numerically evaluate the corrections to Eq. (A26) that so far have been neglected.

We start by taking into consideration the effect of varying the parameter λ . We consider the same numerical setup discussed in Sec. 4.2.3, working in 2D, on a 100×100 grid, with $\Delta t = \Delta x$ and $\kappa_0 \Delta x = 1$. In Fig. A1 we show the results of numerical simulations for a few selected values of λ ; the results clearly show how λ impacts the diffusion speed. Since Eq. (A26) does not depend on λ , but only on κ_0 , we extend it by assuming a dependency on the parameter

$$\chi = \kappa_0 \left(1 + \alpha_1 \frac{1}{d} \lambda \right), \quad (\text{A27})$$

i.e., a linear combination of κ_0 and κ_1 , with α_1 a coefficient left to be determined.

In Sec. 4.1 we have discussed the artificial diffusivity introduced by the interpolation scheme used to implement the propagation step. In order to try to capture these effects we propose the following extension for Eq. (A26):

$$D = \frac{c}{d \chi} \left[1 - \left(\frac{1}{2} + \alpha_2 \right) \Delta t c \chi \right] + \frac{\Delta x^2}{\Delta t} \alpha_3. \quad (\text{A28})$$

In the above, α_2 introduces a correction to the leading term of the Taylor expansion of the propagation step, which is also present in other off-grid LB schemes (see e.g. Coelho et al. (2018)). This correction becomes vanishingly small as one takes smaller timesteps. We attribute this correction to the specific interpolation scheme employed. Besides, we also introduce an extra coefficient α_3 which introduces a background diffusivity which depends on the specific stencil taken into consideration.

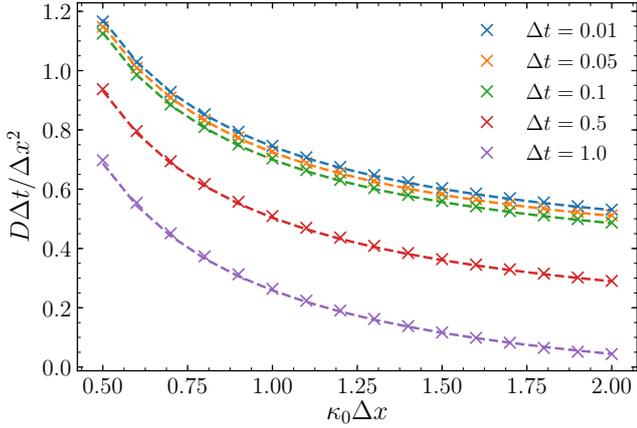


Figure A2. Example of the numerical analysis performed to fit the expression for the diffusion coefficient as in Eq. (A28). The results are shown for the specific case $\lambda = 0.25$. Crosses represent estimates of the numerical diffusion coefficient which ensures that the L2-norm of the relative error is within 1% with respect to the analytic solution. Dashed lines show the predicted diffusion coefficient using Eq. (A28) with $\alpha_1 = 3/4$, $\alpha_2 = 0.48$ and $\alpha_3 = 0.65$.

Our task consists now in determining $\alpha_1, \alpha_2, \alpha_3$. To this aim we perform several numerical simulations, in which we once again reproduce the benchmark described in Sec. 4.2.3, varying $\kappa_0 \Delta x$, Δt and λ . We follow the evolution of each simulation up to $t = 200\Delta t$. Next, we estimate the numerical diffusion coefficient of each simulation. To do so we compare the numerical results with the analytic solution (Eq. (55)), and fit a numerical value of the diffusion coefficient which leaves the L2-norm of the relative error below 1% over several finite time-steps ($t/\Delta t = 5, 10, 20, \dots, 200$). The crosses in Fig A2 provide an example of the results obtained following this procedure, for the specific case $\lambda = 0.25$.

Finally, we attempt to fit the dataset with Eq. (A28). We find that to good accuracy $\alpha_1 \approx -\frac{4}{3}$ in both 2 and 3 dimensions. The parameters α_2 and α_3 , however, depend both on the dimensionality and on the specific stencil taken into consideration, although both tend to stabilise when considering stencils formed by a sufficiently large number of components. To give an example, conducting the analysis in 2D with $N_{\text{pop}} = 120$ we get $\alpha_2 \approx 0.48$, $\alpha_3 \approx 0.65$. In 3D, instead, using a spherical-design quadrature with order $p = 20$ and $N_{\text{pop}} = 222$, we obtain $\alpha_2 \approx 0.58$, $\alpha_3 \approx 0.77$.

APPENDIX B: EMISSIVITY AND OPACITIES IN THE LAB FRAME

For radiative transfer simulations on a moving fluid background it is simpler to express the microscopic quantities describing emission, absorption and scattering processes in the comoving fluid-frame, in which the fluid is at rest. Since our LB scheme is designed in the lab frame, it is then necessary to transform these fluid-frame quantities to their lab-frame counterparts. For doing so we follow the derivation in Mihalas & Auer (2001), which for completeness we summarise here for the grey approximation.

For a fluid moving with three-velocity v_i and Lorentz factor W , the frequency ν in the lab frame of a radiation particle propagating in direction \hat{n}_i is transformed to the comoving fluid frame via

$$\tilde{\nu} = W\nu(1 - v_i \hat{n}^i). \quad (\text{B1})$$

It was then shown that the frequency-dependent quantities, which are isotropic in the fluid frame, transform like

$$\eta_\nu = \frac{\nu^2}{\tilde{\nu}^2} \tilde{\eta}_{\tilde{\nu}} = \frac{\tilde{\eta}_{\tilde{\nu}}}{W^2(1 - v_i \hat{n}^i)^2} \quad (\text{B2})$$

$$\kappa_{a,\nu} = \frac{\tilde{\nu}}{\nu} \tilde{\kappa}_{a,\tilde{\nu}} = W(1 - v_i \hat{n}^i) \tilde{\kappa}_{a,\tilde{\nu}}. \quad (\text{B3})$$

The frequency-averaged absorption opacity defined by Eq. (21) then simply follows as

$$\kappa_a = \frac{\int_0^\infty W(1 - v_i \hat{n}^i) \tilde{\kappa}_{a,\tilde{\nu}} I_\nu d\nu}{\int_0^\infty I_\nu d\nu} = W(1 - v_i \hat{n}^i) \tilde{\kappa}_a, \quad (\text{B4})$$

where $\tilde{\kappa}_a$ is the frequency-integrated fluid-frame absorption opacity. The transformation of the frequency-integrated emissivity is

$$\begin{aligned} \eta &= \int_0^\infty \nu^3 \eta_\nu d\nu = \frac{\int_0^\infty \tilde{\nu}^3 \tilde{\eta}_{\tilde{\nu}} d\tilde{\nu}}{W^2(1 - v_i \hat{n}^i)^2} = \frac{\int_0^\infty \tilde{\nu}^3 \tilde{\eta}_{\tilde{\nu}} d\tilde{\nu}}{W^3(1 - v_i \hat{n}^i)^3} \\ &= \frac{\tilde{\eta}}{W^3(1 - v_i \hat{n}^i)^3}, \end{aligned} \quad (\text{B5})$$

where $\tilde{\eta}$ is the frequency-integrated fluid-frame emissivity.

The scattering opacities are more complicated to transform and we here only consider iso-energetic isotropic scattering like the Thomson scattering process, which we use in Sec. 5. We then recognize that for κ_0 we have two terms on the RHS of Eq. (17). The one proportional to I acts like the absorption term and transforms correspondingly, i.e.,

$$\kappa_0 I = W(1 - v_i \hat{n}^i) \tilde{\kappa}_0 I. \quad (\text{B6})$$

The second term proportional to E acts like an emission term, for which we use the same transformation that we used in Eq. (B5). We then find

$$\kappa_0 E = \frac{\tilde{\kappa}_0 J}{W^3(1 - v_i \hat{n}^i)^3}, \quad (\text{B7})$$

where J is the radiation energy density in the fluid frame, which can be computed from the lab-frame moments according to Eq. (56).

Taking all of the above transformations together means that for the simulation of the relativistic jet in Sec. 5, we solve the mixed-frame equation

$$\frac{1}{c} \frac{\partial I}{\partial t} + \hat{n} \cdot \nabla I = -W(1 - v_i \hat{n}^i) (\tilde{\kappa}_a + \tilde{\kappa}_0) I + \frac{\tilde{\eta} + \tilde{\kappa}_0 J}{W^3(1 - v_i \hat{n}^i)^3} \quad (\text{B8})$$

instead Eq. (17).