







This article has been accepted for publication in Monthly Notices of the Royal Astronomical Society ©: 2019 The Authors. Published by Oxford University Press on behalf of the Royal Astronomical Society. All rights reserved.

AREPO-RT: radiation hydrodynamics on a moving mesh

Rahul Kannan,¹[†] Mark Vogelsberger,²[‡] Federico Marinacci^{1b,2}, Ryan McKinnon^{1b,2},
Rüdiger Pakmor^{1b,3} and Volker Springel^{1b,3,4,5}

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

²Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, 02139 MA, USA

³Heidelberg Institute for Theoretical Studies, Schloss-Wolfsbrunnengasse 35, D-69118 Heidelberg, Germany

⁴Zentrum für Astronomie der Universität Heidelberg, ARI, Mönchhof-str. 12-14, D-69120 Heidelberg, Germany

⁵Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, D-85741 Garching, Germany

Accepted 2019 January 23. Received 2019 January 20; in original form 2018 April 3

ABSTRACT

We introduce AREPO-RT, a novel radiation hydrodynamic (RHD) solver for the unstructured moving-mesh code AREPO. Our method solves the moment-based radiative transfer equations using the M1 closure relation. We achieve second-order convergence by using a slope-limited linear spatial extrapolation and a first-order time prediction step to obtain the values of the primitive variables on both sides of the cell interface. A Harten–Lax–van Leer flux function, suitably modified for moving meshes, is then used to solve the Riemann problem at the interface. The implementation is fully conservative and compatible with the individual time-stepping scheme of AREPO. It incorporates atomic hydrogen (H) and helium (He) thermochemistry, which is used to couple the ultraviolet radiation field to the gas. Additionally, infrared (IR) radiation is coupled to the gas under the assumption of local thermodynamic equilibrium between the gas and the dust. We successfully apply our code to a large number of test problems, including applications such as the expansion of H II regions, radiation pressure-driven outflows, and the levitation of optically thick layer of gas by trapped IR radiation. The new implementation is suitable for studying various important astrophysical phenomena, such as the effect of radiative feedback in driving galactic scale outflows, radiation-driven dusty winds in high-redshift quasars, or simulating the reionization history of the Universe in a self-consistent manner.

Key words: radiative transfer – radiation: dynamics – methods: numerical.

1 INTRODUCTION

Radiation fields are ubiquitous in nature. In the earliest epochs, the Universe was a hot dense soup of matter and radiation. The mean-free path of photons was very small due to Thomson scattering off free electrons. As the Universe cooled to ~ 4000 K about 300 000 yr ($z \simeq 1100$) after the big bang, protons began to capture the free electrons and form atomic hydrogen (H). These recombinations diminished the number density of free electrons allowing the matter and radiation to decouple, making the Universe transparent to light.

The radiation from this epoch is observed today as the nearly uniform ~ 2.72 K cosmic microwave background (CMB; Alpher & Herman 1948; Penzias & Wilson 1965). The CMB is currently our best probe to explore the early Universe. The slight temperature

inhomogeneities present in the CMB ($\Delta T/T \simeq 10^{-5}$, Smoot et al. 1992) correspond to the primordial density fluctuations, which grow through gravitational collapse and form the structures that we see today (White & Rees 1978).

After recombination, the Universe went through a period of darkness with no sources of visible light. Once the gas in the high-density regions became dense enough, it started forming stars and eventually protogalaxies. These early stars and galaxies have an important effect on the surrounding environment. They emit copious amounts of ionizing H radiation (≥ 13.6 eV), which is believed to reionize the neutral intergalactic medium (IGM, Shapiro & Giroux 1987; Haardt & Madau 1996; Gnedin & Ostriker 1997; Madau, Haardt & Rees 1999; Gnedin 2000). Reionization is believed to be initially patchy with pockets of ionized plasma surrounding the most energetic sources. As the pockets grow larger and become more numerous, they overlap and eventually reionize the whole Universe. In addition, the photons heat the IGM, altering its thermal state which in turn affects the observed Lyman- α distribution (Gnedin & Hui 1998; Gnedin 1998; Schaye et al. 2000).

* E-mail: rahul.kannan@cfa.harvard.edu

[†] Einstein Fellow.

[‡] Alfred P. Sloan Fellow.

The epoch of reionization (EoR), carries plenty of information about the process of structure formation in the Universe and provides evolutionary links between the smooth matter distribution at early times revealed by CMB, and the large-scale structure observed at low redshifts. Gaining insights into this epoch is challenging, but recent observations of the steep faint end slope of the ultraviolet (UV) luminosity function indicates that the cosmic ionizing photon budget is dominated by faint galaxies (Bouwens et al. 2011, 2015). Full reionization within $z \sim 6$ can only be achieved if the observed luminosity function is extrapolated by two orders of magnitude below the observational limit of the *Hubble Space Telescope* (Finkelstein et al. 2015). However, photoheating can suppress star formation in low-mass haloes, and it is therefore unclear if such an extrapolation is valid. Furthermore, the escape fraction of ionizing photons from the interstellar (ISM) and the circumgalactic (CGM) media is impossible to measure directly above $z \sim 4$. While some indirect measurements of the escape fraction have been made (Zackrisson, Inoue & Jensen 2013; Dijkstra, Gronke & Venkatesan 2016; Reddy et al. 2016; Mas-Ribas et al. 2017), they are not very constraining and allow the escape fraction to be anywhere between 1 per cent and 30 per cent.

With the imminent launch of the *James Webb Space Telescope* (Gardner et al. 2006), the study of EoR enters a new era. It is expected to increase the quality of high-redshift data and extend it beyond $z = 10$. It is therefore important for theoretical models to achieve enough accuracy to interpret the observational results. Numerical radiation hydrodynamic (RHD) simulations offer the most accurate and realistic theoretical models of reionization. Hence, an implementation of an accurate and efficient radiative transfer (RT) algorithm becomes imperative to study the high-redshift Universe.

Radiation fields also play an important role in many physical processes that occur inside dark matter (DM) haloes. For example, a long-standing puzzle in galaxy formation theory has been the low star formation efficiency in DM haloes (Silk & Mamon 2012). Star formation efficiency peaks at about $M_{\text{halo}} \sim 10^{12} M_{\odot}$ and decreases in both higher and lower mass haloes (Moster et al. 2010; Moster, Naab & White 2013; Behroozi, Wechsler & Conroy 2013). Feedback from stars (Navarro, Frenk & White 1996; Springel & Hernquist 2003; Stinson et al. 2006; Dalla Vecchia & Schaye 2008; Agertz et al. 2013; Vogelsberger et al. 2013; Hopkins et al. 2014, 2017) and the central active galactic nuclei (AGNs; Springel, Di Matteo & Hernquist 2005; Sijacki & Springel 2006; Booth & Schaye 2009; Choi et al. 2012; Kannan et al. 2017; Weinberger et al. 2017) are invoked to explain this low star formation efficiency in low- and high-mass galaxies, respectively.

Early galaxy formation simulations showed that the coupling between supernovae (SNe) feedback energy and the ISM is very inefficient (Katz, Weinberg & Hernquist 1996; Navarro et al. 1996), as most of the injected energy is radiated away very efficiently. Various sub-grid models have been proposed in order to avoid cooling losses such as delaying the cooling of gas particles around a star (Thacker & Couchman 2001; Stinson et al. 2006; Agertz et al. 2013), stochastically injecting energy into the surrounding gas such that it is heated up to the temperatures where gas cooling becomes inefficient (Dalla Vecchia & Schaye 2008; Schaye et al. 2015), and injecting kinetic energy that adds velocity kicks to gas particles to remove them from the inner regions of galactic discs (Springel & Hernquist 2003; Oppenheimer & Davé 2006; Vogelsberger et al. 2013). These ad hoc methods have been quite successful in reproducing the properties of galaxies in a broad sense (Vogelsberger et al. 2014a,b; Schaye et al. 2015). However, they require fine

tuning of free parameters in the model in order to reproduce the low-mass end of the luminosity function (Vogelsberger et al. 2014a; Schaye et al. 2015; Pillepich et al. 2018) and in some cases require unrealistic values of SNe feedback energy ($> 10^{51}$ erg; Guedes et al. 2011; Schaye et al. 2015) or excessively large gas outflow velocities (Pillepich et al. 2018).

Many recent works have pointed out that young massive stars deposit large amounts of energy in the form of photons and stellar winds before they go SNe, which can have a significant dynamical impact on the ISM (Murray, Quataert & Thompson 2010; Walch et al. 2012). Stinson et al. (2013) showed that the high-energy photons emitted by OB stars can ionize and photoheat the surrounding regions helping to regulate star formation especially at high redshifts (Kannan et al. 2014b). However, this requires efficient thermalization in the injected radiation energy close to the source, which is not guaranteed in the high-density regions where stars form. Radiation pressure, both direct UV and multiscattered infrared (IR), is another mechanism hypothesized to drive significant outflows ($\sim 100 \text{ km s}^{-1}$) (Hopkins, Quataert & Murray 2011; Agertz et al. 2013; Hopkins et al. 2014). However, it seems that unphysically large optical depths of IR radiation ($\tau_{\text{IR}} \sim 50$) are required to effectively trap the photons and boost the momentum injection to the levels required to efficiently suppress star formation (Roškar et al. 2014). Alternatively, if enough radiation escapes the star-forming regions and the ISM of galaxies, it can in principle reduce the gas cooling rates of the CGM thereby reducing gas inflows into the centres of galaxies (Cantalupo 2010; Gnedin & Hollon 2012; Kannan et al. 2014a, 2016b). While these works hint towards the importance of radiation fields, the crude nature of these sub-grid models makes it difficult to gauge the exact mechanisms and significance of radiation fields in regulating the star formation rates of low-mass galaxies. Therefore, full RHD simulations are necessary in order to gain a fundamental understanding of stellar feedback (Rosdahl et al. 2015; Kim et al. 2017; Peters et al. 2017).

The impact of stellar feedback decreases as the mass and potential depth of DM haloes increases. A more energetic source of feedback is needed, which is conveniently found in the form of the central AGN. The inability of galaxy formation simulations to resolve the region around the central supermassive black hole, necessitates sub-grid prescriptions to account for this feedback channel. These models generally discriminate between a high accretion rate quasar mode (Springel et al. 2005) feedback and a low accretion rate mechanical radio mode feedback (Weinberger et al. 2017). However, the details of how the AGN energy couples to the gas in and around galaxies is still uncertain, so modelling efforts have so far been necessarily crude.

In general, the AGN deposits energy and momentum into the surrounding gas driving winds. Observations show compelling evidence for galaxy-scale AGN-driven outflows (Heckman et al. 1981; Nelson & Whittle 1995; Greene & Ho 2005; Karouzos, Woo & Bae 2016). These are generally high-mass loaded ($\dot{M}_{\text{out}} > 10^8 M_{\odot}$), fast ($v_{\text{out}} \geq 1000 \text{ km s}^{-1}$), and multiphase winds that can extend over up to several tens of kpc (Cicone et al. 2015; Tombesi et al. 2015; Zakamska et al. 2016). The efficiency with which these winds suppress star formation is still unknown mainly because the exact mechanism of the coupling between the AGN feedback energy and the gas is not understood. Ultrafast outflows ($\geq 10\,000 \text{ km s}^{-1}$) can develop a two-temperature structure, where most of the thermal pressure support is provided by the protons, while the cooling processes operate directly only on the electrons. This significantly slows down inverse Compton cooling, maintaining the thermal structure of the wind and generating large momentum boosts ($\sim 20L/c$) and high kinetic

luminosities (Faucher-Giguère & Quataert 2012; Costa, Sijacki & Haehnelt 2014). The AGN luminosity can also couple directly to the surrounding gas through radiation pressure on dust (Fabian 1999; Murray et al. 2010; Thompson et al. 2016). Observations indicate a large fraction of the optical and UV radiation is absorbed and re-emitted at IR wavelengths by a surrounding envelope of dusty gas before escaping the galactic nucleus (Fabian & Iwasawa 1999). If the column density of the gas is high enough, the optical depth to IR radiation can be rather large, trapping this radiation within the AGN nucleus. The multiscattered IR photons boost the momentum injection by a factor proportional to the optical depth ($\dot{p} \sim \tau_{\text{IR}} L/c$).

Many recent studies have tried to model this mechanism using simple analytic models (Ishibashi & Fabian 2015; Thompson et al. 2016) and idealized high-resolution RT simulations of the dusty AGN torus (Roth et al. 2012; Ishibashi & Fabian 2017; Costa et al. 2018). Recent advances in numerical techniques have allowed for the estimation of the impact of radiation pressure on galactic scale using isolated disc simulations (Bieri et al. 2017; Cielo et al. 2017) and more recently in cosmological simulations (Costa et al. 2017). These simulations suggest that IR radiation pressure can drive fast ($\geq 1000 \text{ km s}^{-1}$), high-mass loaded and short-lived winds during the obscured phase of the AGN. In fact, these winds are able to remove enough material from the centre to completely shut down star formation, indicating the need to employ accurate RHD simulations to understand radiative feedback from AGNs.

On smaller scales, the radiation from massive protostars can unbind gas in its surroundings and create cavities. These cavities can prevent any further accretion on to the star from the direction of the bubble, cutting off fuel supply to the star and stalling the mass growth (Kuiper et al. 2012), although the development of Rayleigh–Taylor instabilities (RTIs) can destroy the cavity and allow gas to fall back on to the star (Rosen et al. 2016). The stellar radiation fields also have a large impact on the structure of the protoplanetary discs (Flock et al. 2016) and on the climate of exoplanets (Heng, Frierson & Phillipps 2011).

To summarize, radiation plays a crucial role in a large variety of astrophysical systems, and its impact ranges all the way from small planetary systems to the large-scale thermal and ionization history of the Universe. The complexity of RT requires accurate RHD simulations to precisely capture and model its impact. Consequently, an accurate and efficient RT implementation is needed to improve current astrophysical simulations by taking into account the effects of radiation. In this work, we present a moment-based RT implementation for the moving mesh hydrodynamics code AREPO.

The paper is structured as follows. In Section 2, we briefly outline the various RT schemes used in literature and discuss the advantages and shortcomings of the scheme used in this paper. Section 3 describes the spatial discretization and time-integration techniques used to solve the RT equations for our scheme. In Section 4, we present several test problems to quantify the accuracy of our implementation. Finally, we present our conclusions in Section 5.

2 THE RADIATIVE TRANSFER EQUATIONS

Here, we first discuss the relevant RT equations that we are going to solve. We start by defining the specific intensity $I_\nu(\mathbf{x}, t, \mathbf{n}, \nu)$, at position \mathbf{x} and time t , as the rate of radiation energy (E_ν) flowing per unit area (dA), in the direction (\mathbf{n}), per unit time (dt), per unit frequency interval ($d\nu$) centred on frequency ν , and per unit solid angle ($d\Omega$)

$$dE_\nu = I_\nu(\mathbf{x}, t, \mathbf{n}, \nu) (\mathbf{n} \cdot d\mathbf{A}) dt d\nu d\Omega. \quad (1)$$

The propagation of the radiation field and interactions with the surrounding medium such as absorption and emission of radiation leads to a change in the radiation energy at that spatial position. Taking these processes into account, we can write down the continuity equation for the specific intensity as (Mihalas & Mihalas 1984)

$$\frac{1}{c} \frac{\partial I_\nu}{\partial t} + \mathbf{n} \cdot \nabla I_\nu = j_\nu - \kappa_\nu \rho I_\nu, \quad (2)$$

where j_ν is the emission term and κ_ν is the absorption coefficient.

RT is a complex process due its high dimensionality. An accurate numerical solution requires discretizing equation (2) in angular and frequency variables in addition to spatial and time discretization. A variety of different numerical algorithms have been proposed to solve the RT problem. The most common method is the long characteristic ray-tracing scheme (Mihalas & Mihalas 1984; Abel, Norman & Madau 1999; Abel & Wandelt 2002; Greif 2014; Jaura et al. 2017) that cast rays from each source through the simulation domain and solve equation (2) along each ray. This method, although very accurate, is computationally expensive ($\mathcal{O}(N^2)$) and requires high angular resolution to capture the correct transport of radiation. Furthermore, parallelizing this algorithm requires significant data exchanges between different processors. In order to reduce the complexity of the problem, some works have resorted to short characteristics methods (Ciardi et al. 2001; Whalen & Norman 2006; Trac & Cen 2007; Pawlik & Schaye 2008; Petkova & Springel 2011) which integrates the RT equation only along lines that connect nearby cells making it easier to parallelize.

The RT equations can also be solved with Monte Carlo methods. These schemes (Oxley & Woolfson 2003; Semelin, Combes & Baek 2007; Dullemond et al. 2012), emit individual photon packets to sample the interaction lengths and scattering angles of the photons from the underlying probability density functions. While they perform remarkably well, they are computationally very demanding and the Poisson noise inherent to the statistical description of the radiation field leads to a signal-to-noise ratio that grows only with the square root of the number of photon packets emitted.

Solving the moments of the RT equation has gained popularity in recent years (Levermore 1984; González, Audit & Huynh 2007; Rosdahl et al. 2013; Rosdahl & Teyssier 2015). A fluid description of the radiation field is obtained by taking the zeroth and first moments of equation (2)

$$\frac{\partial E_r}{\partial t} + \nabla \cdot \mathbf{F}_r = S - \kappa_E \rho \tilde{c} E_r, \quad (3)$$

$$\frac{\partial \mathbf{F}_r}{\partial t} + \tilde{c}^2 \nabla \cdot \mathbb{P}_r = -\kappa_F \rho \tilde{c} \mathbf{F}_r, \quad (4)$$

where the radiation energy density (E_r), flux (\mathbf{F}_r), and pressure (\mathbb{P}_r) are defined as

$$\{\tilde{c} E_r, \mathbf{F}_r, \tilde{c} \mathbb{P}_r\} = \int_{\nu_1}^{\nu_2} \int_{4\pi} \{1, \mathbf{n}, (\mathbf{n} \otimes \mathbf{n})\} I_\nu d\Omega d\nu. \quad (5)$$

Here (in equations 3 and 4), S denotes the source term which quantifies the amount of radiation energy emitted, κ_E and κ_F are the radiation energy density and radiation flux-weighted mean opacities within the frequency range defined by $[\nu_1, \nu_2]$ and ρ is the density of gas in the cell. We note that we have reformulated the equations in terms of signal speed (\tilde{c}) of radiation transport, which can be different from the actual speed of light (c) when the reduced speed of light approximation (RSLA) is used (see Section 3.1 for more details).

This set of hyperbolic conservation equations defines the rate of change of E_r and \mathbf{F}_r as a function of time and position. However, in order to solve these equations an estimate of \mathbb{P}_r is required. We therefore need to obtain the pressure tensor by invoking a closure relation. One popular choice, which works remarkably well in optically thick media, is to recast equations (3) and (4) into a diffusion equation (Lucy 1977; Krumholz & Thompson 2012; González et al. 2015) by assuming that the photon flux is proportional to the gradient of the photon energy density ($\mathbf{F}_r = -\tilde{c}\nabla E_r/3\kappa_F\rho$). While this approximation performs well in highly optically thick media, its accuracy in optically thin cases is not well understood. Furthermore, maintaining the directionality of photon propagation is quite difficult as the photon flux is assumed to be directed along the gradient of the photon energy density, forcing the photons to diffuse isotropically. This makes it difficult to form sharp shadows behind optically thick barriers (Zhang & Davis 2017). A better approximation can be obtained by ignoring the term of the order c^{-1} in equation (4) and set

$$\mathbf{F}_r = -\frac{\tilde{c}}{\kappa_F\rho}\nabla\cdot\mathbb{P}_r, \quad (6)$$

which yields

$$\frac{\partial E_r}{\partial t} + \nabla\cdot\left(-\frac{\tilde{c}}{\kappa_F\rho}\nabla\cdot\mathbb{P}_r\right) = S - \kappa_E\rho\tilde{c}E_r. \quad (7)$$

The Eddington tensor formalism can then be used to equate the radiation energy density (E_r) and the radiation pressure tensor (\mathbb{P}) by defining a proportionality tensor called the Eddington tensor (\mathbb{D})

$$\mathbb{P}_r = E_r\mathbb{D}. \quad (8)$$

This Eddington tensor essentially encodes the direction of photon transport at each point in the domain. In this form, the RT equations transform into an anisotropic diffusion equation. The discretization of this equation is surprisingly non-trivial and widely used methods give rise to unphysical oscillations (Parrish & Stone 2005). Some works get around this problem by adding an isotropic component to the anisotropic diffusion tensor, however, this reduces the accuracy with which the algorithm preserves the directionality of the underlying photon field (Petkova & Springel 2009). Therefore, this method suffers from the same problems as the isotropic diffusion approximation. Furthermore, the time-step limitations imposed by the parabolic diffusion equation requires the implementation of implicit or semi-implicit schemes in order to make the algorithm fast enough (Kannan et al. 2016a, 2017). These time-integration techniques are difficult to implement and parallelize efficiently.

For these reasons, we chose to discard the approximations above and instead solve the coupled hyperbolic conservation laws for the photon energy density (equation 3) and the photons flux (equation 4), coupled with the Eddington closure relation (equation 8). It is straightforward to compute \mathbb{D} in the case of a single or few sources. However, the computation becomes quite arduous when considering galaxy-scale simulations which can have millions of sources within the simulation domain. Many works have tried to derive approximate estimates of \mathbb{D} , such as, for instance, the optically thin variable Eddington tensor (OTVET) formalism (Gnedin & Abel 2001; Finlator, Özel & Davé 2009; Petkova & Springel 2009) or the M1 (Levermore 1984; Dubroca & Feugeas 1999; Ripoll, Dubroca & Duffa 2001) method. The OTVET formalism computes \mathbb{D} by assuming that the intervening material between the radiation source and sink is optically thin. The obvious drawback of this method is that direction of radiation field behind any optically thick material will not be correctly captured. Additionally, the computational cost

associated with estimating the positions of every source relative to every volume element can be quite significant.

On the other hand, the M1 closure requires only local quantities of a given cell to compute the Eddington Tensor \mathbb{D} as

$$\mathbb{D} = \frac{1-\chi}{2}\mathbb{I} + \frac{3\chi-1}{2}\mathbf{n}\otimes\mathbf{n}, \quad (9)$$

where

$$\mathbf{n} = \frac{\mathbf{F}_r}{|\mathbf{F}_r|}, \quad \chi = \frac{3+4f^2}{5+2\sqrt{4-3f^2}}, \quad \text{and } f = \frac{|\mathbf{F}_r|}{\tilde{c}E_r}. \quad (10)$$

Since the radiation flux cannot be larger than the signal speed times the radiation energy density, the reduced flux (f) will always be limited by $0 \leq f \leq 1$. The local nature of the M1 closure implies that the computational cost is independent of the number of sources and only depends on the number of resolution elements within the domain. This has allowed recent works to perform galaxy-scale simulations with relatively low computational cost (Rosdahl et al. 2015; Costa et al. 2017; Bieri et al. 2017). Since we are mostly interested in performing simulations containing a large number of sources, we adopt the M1 closure formalism in our work.

Having established our numerical scheme to evolve the radiation field, we now have to couple the radiation to the gas. The radiation field couples to the gas hydrodynamics via photon absorption and scattering (RHS of equations 3 and 4). These physical processes are quantified using the average opacities (κ_E and κ_F) of the gas. Energy and momentum conservation then dictates that photon absorption introduces source terms into hydrodynamic momentum and energy conservation equations (González et al. 2007):

$$\frac{\partial(\rho\mathbf{v})}{\partial t} + \nabla\cdot(\rho\mathbf{v}\mathbf{v}^T + P\mathbb{I}) = \frac{\kappa_F\rho}{c}\mathbf{F}_r, \quad (11)$$

$$\frac{\partial(\rho E)}{\partial t} + \nabla\cdot[(\rho E + P)\mathbf{v}] = -\Lambda + \kappa_E\rho\tilde{c}E_r + \frac{\kappa_F\rho}{c}\mathbf{F}_r\cdot\mathbf{v}, \quad (12)$$

where Λ is the gas cooling rate which is function of the abundance of the ionic species present in the gas, which is in turn dependent on the incident radiation field, P its thermal pressure, E its total energy per unit mass, and \mathbf{v} the gas velocity field.

3 METHODS

In this section, we describe in detail the spatial discretization and time-integration techniques used to solve the RHD equations introduced in Section 2. The transport equations (setting the RHS of equations 3 and 4) and the source terms are solved separately using an operator split approach. This is achieved using a Strang split scheme (Strang 1968), which involves a half-step update of the primitive RT variables (E_r , \mathbf{F}_r) due to the source terms, a full step update due to the transport operations, and finally another half-step update with the source operations. This makes the solution formally converge at second order.

3.1 The transport equations

Let us first consider the free transport of photons, which is obtained by setting the RHS of equations (3) and (4) to zero

$$\frac{\partial E_r}{\partial t} + \nabla\cdot\mathbf{F}_r = 0, \quad (13)$$

$$\frac{\partial \mathbf{F}_r}{\partial t} + \tilde{c}^2\nabla\cdot\mathbb{P}_r = 0. \quad (14)$$

Equations (13) and (14) are conservation laws for the photon energy density and photon flux that take the form of a system of hyperbolic partial differential equations. They can be written in compact conservative form as

$$\frac{\partial \mathcal{U}}{\partial t} + \nabla \cdot \mathcal{F}(\mathcal{U}) = 0, \quad (15)$$

where

$$\mathcal{U} = \begin{pmatrix} E_r \\ \mathbf{F}_r \end{pmatrix}, \quad (16)$$

and

$$\mathcal{F}(\mathcal{U}) = \begin{pmatrix} \mathbf{F}_r \\ \tilde{c}^2 \mathbb{P}_r \end{pmatrix}. \quad (17)$$

A finite-volume simulation code like AREPO divides the computational domain into a set of control volumes. The fluid's state is described by the cell averages of the conserved quantities, which are obtained by integrating the primitive quantities over the volume of the cell

$$\mathcal{Q} = \int_V \mathcal{U} dV. \quad (18)$$

Using the Gauss' theorem, we can estimate the change in these quantities with time as

$$\frac{\partial \mathcal{Q}}{\partial t} = - \int_{\partial V} [\mathcal{F}(\mathcal{U}) - \mathcal{U} \mathbf{w}^T] dA, \quad (19)$$

where \mathbf{w} is the velocity of each point of the cell boundary. For Eulerian schemes, the mesh is static ($\mathbf{w} = 0$), while in a fully Lagrangian approach, the surface would be allowed to move at every point with the local flow velocity ($\mathbf{w} = \mathbf{v}$). For moving-mesh codes like AREPO, it is not possible to follow the distortions of the shapes of fluid volumes exactly in multidimensional flows and therefore the general formula of equation (19) is used. Practically, this requires then to solve the total flux as a combination of the flux over a static interface ($\mathcal{F}(\mathcal{U})$) and an advection step owing to the movement of the interface ($-\mathcal{U} \mathbf{w}^T$).

The hydro and magnetohydrodynamic schemes solve the respective equations in the reference frame of the moving face. This ensures full Galilean invariance, a property that is of significant importance for cosmological simulations where highly supersonic bulk flows are common. Unfortunately, this approach is not possible for a photon fluid, where Galilean invariance does not have any meaning. We instead choose to modify the Riemann solution such that it takes the motion of the mesh into account.

We therefore write down the total flux on a moving mesh (\mathcal{F}^m) as

$$\mathcal{F}^m = \begin{pmatrix} \mathbf{F}_r - E_r \mathbf{w}^T \\ \tilde{c}^2 \mathbb{P}_r - \mathbf{F}_r \mathbf{w}^T \end{pmatrix}, \quad (20)$$

where the velocity of the cell interface (\mathbf{w}) is calculated using the method outlined in Springel (2010, equations 32 and 33). Godunov's (1959) approach is used to compute \mathcal{F}^m and solve the approximate Riemann problem normal to the interface. Since AREPO uses unstructured Voronoi meshes, the dimensionally operator split framework cannot be applied. Rather the unsplit approach is used as described in Springel (2010). In a nutshell, this involves defining the Riemann problem normal to the cell face by rotating the relevant primitive variables into the coordinate system defined by setting the x -axis normal to the cell face. Once the flux has been calculated, it is transformed back to the lab frame.

We employ the Harten–Lax–van Leer (HLL, Harten, Lax & van Leer 1983) framework that splits the solution of the Riemann

problem at each interface into three possible flux estimates

$$\mathcal{F}^m = \begin{cases} \mathcal{F}_L^m & \text{if } \lambda^- \geq 0, \\ \mathcal{F}_{\text{hll}}^m & \text{if } \lambda^- \leq 0 \leq \lambda^+, \\ \mathcal{F}_R^m & \text{if } \lambda^+ \leq 0, \end{cases} \quad (21)$$

where

$$\mathcal{F}_{\text{hll}}^m = \frac{\lambda^+ \mathcal{F}_L^m - \lambda^- \mathcal{F}_R^m + \lambda^+ \lambda^- (\mathcal{U}_R - \mathcal{U}_L)}{\lambda^+ - \lambda^-}. \quad (22)$$

The subscripts 'L' and 'R' refer to the value of the variables ($\mathcal{U}, \mathcal{F}(\mathcal{U})$) on the left and right states of the cell interface. λ^+ and λ^- are the maximum or minimum eigenvalues of the of the Jacobian $\partial \mathcal{F}^m / \partial \mathcal{U}$ defined as

$$\begin{aligned} \lambda^+ &= \max(\lambda_L^{\text{max}}, \lambda_R^{\text{max}}), \\ \lambda^- &= \min(\lambda_L^{\text{min}}, \lambda_R^{\text{min}}). \end{aligned} \quad (23)$$

The eigenvalues represent the wave speeds of system of equations which in our formulation are estimates for the lower and upper bounds of the signal velocities. The eigenvalues of the system are obtained by solving for λ in

$$\left| \frac{\partial \mathcal{F}^m}{\partial \mathcal{U}} - \lambda_m \mathcal{I} \right| = 0. \quad (24)$$

However, from equation (20)

$$\frac{\partial \mathcal{F}^m}{\partial \mathcal{U}} = \frac{\partial \mathcal{F}^s}{\partial \mathcal{U}} - w \mathcal{I}, \quad (25)$$

where $w = \mathbf{w} \cdot \hat{\mathbf{n}}$ is the component of the velocity of the face along the face normal ($\hat{\mathbf{n}}$) and \mathcal{F}^s is the total flux on a static mesh. Substituting the value of $\partial \mathcal{F}^m / \partial \mathcal{U}$ from equation (25) into equation (24), we get

$$\left| \frac{\partial \mathcal{F}^s}{\partial \mathcal{U}} - (\lambda_m + w) \mathcal{I} \right| = 0. \quad (26)$$

Comparing equations (24) and (26) gives $\lambda_s = \lambda_m + w$, meaning that the eigenvalues of the system on a moving mesh (λ_m) are just a linear combination of the eigenvalues of the system on static mesh (λ_s) and the face velocity (w),

$$\begin{aligned} \lambda^+ &= \lambda_m^+ = -w + \lambda_s^+, \\ \lambda^- &= \lambda_m^- = -w + \lambda_s^-. \end{aligned} \quad (27)$$

This is equivalent to rotating the eigenvectors by an angle $x/t = -w$. We note that if the face velocity is superluminal or greater than the largest signal speed of the static system, then the Riemann solver chooses to purely advect the fluxes in an upwind manner, analogous to supersonic fluid flow in non-relativistic hydrodynamics.

The eigenvalues of the Jacobian matrix are determined by interpolating the tabulated values obtained by González et al. (2007).¹ By setting $\lambda_s^+ = -\lambda_s^- = \tilde{c}$, we obtain the Rusanov (1961) flux function, which is also implemented (also known as the Global–Lax–Friedrichs or GLF flux function).

A conservative time integration of equation (19) is obtained by using the method outlined in Pakmor et al. (2016, equations 19–22), which employs Heun's method, a variant of the second-order Runge–Kutta scheme. The fluxes are computed as an average of fluxes at the beginning and end of the time-step. The mesh geometry of the second half of the current time-step is used for the first half of the next time-step, essentially requiring only one mesh construction

¹We obtain the table from the public version of the RAMSES-RT code.

per time-step. Therefore, the update of the conservative variables are given by

$$\mathcal{Q}_i^{n+1} = \mathcal{Q}_i^n - \frac{\Delta t}{2} \left(\sum_j A_{ij}^n \hat{\mathbf{F}}_{ij}^n(\mathcal{U}^n) + A_{ij}' \hat{\mathbf{F}}_{ij}'(\mathcal{U}') \right). \quad (28)$$

This makes the scheme easily compatible with the hierarchical time-stepping scheme used in AREPO (see Pakmor et al. 2016 for more details).

A higher order accuracy is obtained by replacing the piecewise constant (PC) approximation of Godunov's scheme with a slope-limited piecewise linear spatial extrapolation and a first-order prediction forward in time to obtain the states of the primitive variables on both sides of the interface (van Leer 1979)

$$\mathcal{U}_{\text{L,R}} = \mathcal{U}_{\text{L,R}} + \frac{\partial \mathcal{U}}{\partial \mathbf{r}} \Big|_{\text{L,R}} (\mathbf{f} - \mathbf{s}_{\text{L,R}}) + \frac{\partial \mathcal{U}}{\partial t} \Big|_{\text{L,R}} \Delta t, \quad (29)$$

where \mathbf{s} is the centre of mass of the cell and \mathbf{f} is the centre of the cell face. The time derivatives of the primitive variables ($\partial \mathcal{U} / \partial t$ in equation 29) are expressed in terms of their spatial derivatives using equations (13) and (14)

$$\frac{\partial E_r}{\partial t} = -\nabla \cdot \mathbf{F}_r, \quad (30)$$

$$\frac{\partial \mathbf{F}_r}{\partial t} = -\tilde{c}^2 \nabla \cdot \mathbb{P}_r. \quad (31)$$

We use the local least-squares fit (LSF) method described in Pakmor et al. (2016) to obtain the gradient estimates ($\partial \mathcal{U} / \partial \mathbf{r}$ in equation 29). They are constructed such that they reproduce the cell centred values of the neighbouring cells as well as possible. If ϕ_i is the primitive variable defined at the centre of cell ' i ', and ϕ_j 's are the values of the primitive variable for all neighbouring cells ' j ', then the best linear approximation of $\langle \nabla \phi \rangle_i$ is obtained by minimizing the sum of the deviations for all neighbours $> (S_{\text{tot}})$

$$S_{\text{tot}} = \sum_j \frac{A_{ij}}{|\mathbf{s}_j - \mathbf{s}_i|^2} (\phi_j - \phi_i - \langle \nabla \phi \rangle_i (\mathbf{s}_j - \mathbf{s}_i))^2, \quad (32)$$

where A_{ij} is the area of the interface between cells ' i ' and ' j '. The monotonicity of the gradients is imposed by requiring that the linearly reconstructed quantities on the face centroids are bounded by the maxima and minima of all the cell centred values of the neighbouring cells (equations 28–30, Springel 2010). This general gradient estimate retains the necessary accuracy even for large offsets between the mesh-generating points and the centre of mass of cells.

Ordinarily, the spatial and time extrapolations of the primitive variables are carried out independent of each other. Unfortunately, these independent estimates cannot ensure that the reduced flux at the interface ($f_{\text{L,R}} = (|\mathbf{F}_{\text{r,L,R}}| / \tilde{c} E_{\text{r,L,R}})$) remains bounded between $[0, 1]$. We can enforce this condition by limiting the value of the extrapolated photon flux as $\tilde{\mathbf{F}}_{\text{r,L,R}} = \alpha \mathbf{F}_{\text{r,L,R}}$, where $\alpha = \min(1, \tilde{c} E_{\text{r,L,R}} / |\mathbf{F}_{\text{r,L,R}}|)$. However, this form of a limiter introduces too much noise in the solution and degrades the convergence of the code in addition to increasing the diffusivity.

Instead, the extrapolation of \mathbf{F}_r is made dependent on value of E_r and the reduced flux. Accordingly, the gradients of E_r , $\mathbf{F}_N = \mathbf{F}_r / E_r$ and $|\mathbf{F}_N|$ are calculated instead of just E_r and \mathbf{F}_r . First, the time prediction step of the photon energy density (equation 30) is

performed²

$$E_r' = E_r - \nabla \cdot \mathbf{F}_r \Delta t, \quad (33)$$

where

$$\begin{aligned} \nabla \cdot \mathbf{F}_r &= \text{Tr}(\nabla \mathbf{F}_r) \\ &= \text{Tr} \left(E_r \nabla \mathbf{F}_N + \frac{1}{E_r} \mathbf{F}_r \otimes \nabla E_r \right). \end{aligned} \quad (34)$$

The spatial extrapolation is then carried out by extrapolating the photon energy density using the estimate of the gradient of E_r

$$E_{\text{r,L,R}}' = E_r' + \nabla E_r \cdot (\mathbf{f} - \mathbf{s}_{\text{L,R}}), \quad (35)$$

and then, we extrapolate

$$|\mathbf{F}_{\text{N,L,R}}| = |\mathbf{F}_N| + \nabla |\mathbf{F}_N| \cdot (\mathbf{f} - \mathbf{s}_{\text{L,R}}). \quad (36)$$

Finally, we impose

$$\mathbf{F}_{\text{r,L,R}}' = \psi \mathbf{F}_r, \quad (37)$$

where

$$\psi = \frac{|\mathbf{F}_{\text{N,L,R}}|}{|\mathbf{F}_r|} E_{\text{r,L,R}}'. \quad (38)$$

We take advantage of the fact that the gradients are limited such that they are bounded by the maximum and minimum cell centred values of the neighbouring cells. This method of extrapolating the photon flux ensures that the reduced flux at the interface ($f_{\text{L,R}}$) is always within $[0, 1]$ without affecting the directionality of the underlying photon field. Additionally, the photon energy density and the reduced flux at the interface is not slope limited in any way except for the condition that it does not create local maxima or minima. This ensures stability without adding too much diffusion into the scheme. We note that the prediction step is not performed for the photon flux (\mathbf{F}_r) because, it introduces too much noise into the solution, in addition to making it difficult to preserve the directionality of photon field.

Since the transport equations are solved using an explicit scheme, the time-step of each cell is constrained by the von Neumann stability condition

$$\Delta t_{\text{RT}} \leq \eta \frac{\Delta x}{\tilde{c} + |\mathbf{v}_c|}, \quad (39)$$

where Δx is the cell width, \mathbf{v}_c is the velocity of the cell in the lab frame, and $\eta \sim 0.3$. Since the RT scheme is coupled to hydrodynamics and gravity the final time-step will be

$$\Delta t = \min(\Delta t_{\text{RT}}, \Delta t_{\text{hydro}}, \Delta t_{\text{grav}}). \quad (40)$$

The high speed of light demands very small Δt_{RT} forcing other computationally expensive parts of the code, such as, mesh construction and gravity force calculations, to be called more often than actually required. In many physical problems, this can be overcome by using the RSLA, which is applicable for systems where the characteristic velocity is much smaller than the speed of light. For large-scale cosmological simulations, this is no longer true. For example, to track the ionization fronts (I-fronts) properly in the IGM one must use the full speed of light (Rosdahl et al. 2013; Bauer et al. 2015). Implicit/semi-implicit time-integration schemes can be used to overcome this problem. Unfortunately, scalable implementation

²For the calculation of the fluxes between two cells on different time-steps, the time extrapolation is done for each cell always from the last time the cell was active.

of these schemes is quite involved as it requires inverting large sparse matrices (e.g. Kannan et al. 2016a). The other option is to subcycle the RT steps (Commerçon, Debout & Teyssier 2014), i.e. perform $N_{\text{sub}} \geq 1$ number of RT steps per hydrodynamical step. This effectively reduces the frequency with which the time-consuming routines are called. We chose to implement the latter and the details of this method are described in Appendix A.

3.2 The thermochemistry and cooling equations

In this section, we describe the numerical methods to evaluate the source terms of the RT equations (RHS of equations 3 and 4).

3.2.1 Ultraviolet thermochemistry, photoheating, and radiation pressure

We first look at the single-scattering regime, where a particular photon interacts with the surrounding medium only once i.e. it is absorbed only once, which destroys the photon, and there is little scattering. This occurs mainly in the UV regime, where the photons have enough energy to ionize atoms such as H and helium (He). Since the ionization potential varies quite considerably, an accurate treatment of UV RT requires multiple frequency bins centred around the energy of each ionization state of the gas.

In the single-scattering regime, it is easier to work with photon number densities instead of the radiation energy densities (E_r). Therefore, we define the photon number density N_γ^i , photon number flux (\mathbf{F}_γ^i), and the associated pressure tensor (\mathbb{P}_γ^i) in each frequency bin ‘ i ’ as

$$\{\tilde{c}N_\gamma^i, \mathbf{F}_\gamma^i, \tilde{c}\mathbb{P}_\gamma^i\} = \int_{\nu_{i1}}^{\nu_{i2}} \frac{1}{h\nu} d\nu \int_{4\pi} \{1, \mathbf{n}, (\mathbf{n} \otimes \mathbf{n})\} I_\nu d\Omega, \quad (41)$$

where $\nu_{i1} \leq \nu_i < \nu_{i2}$.

Accordingly, equations (3) and (4) can be reformulated as

$$\begin{aligned} \frac{\partial N_\gamma^i}{\partial t} + \nabla \cdot \mathbf{F}_\gamma^i &= - \sum_j \tilde{c} n_j N_\gamma^i \tilde{\sigma}_{ij} - \kappa_i \rho \tilde{c} N_\gamma^i \\ &+ \sum_j s_{ij}, \end{aligned} \quad (42)$$

$$\frac{\partial \mathbf{F}_\gamma^i}{\partial t} + \tilde{c}^2 \nabla \cdot \mathbb{P}_\gamma^i = - \sum_j \tilde{c} n_j \mathbf{F}_\gamma^i \tilde{\sigma}_{ij} - \kappa_i \rho \tilde{c} \mathbf{F}_\gamma^i, \quad (43)$$

where n_j is the number density of a particular ionic species, which in our case consists of $j \in \{\text{H I}, \text{He I}, \text{He II}\}$. κ_i is the dust opacity of photon group ‘ i ’ and $\tilde{\sigma}_{ij}$ is the mean ionization cross-section of species ‘ j ’ in the frequency bin ‘ i ’

$$\tilde{\sigma}_{ij} = \frac{\int_{\nu_{i1}}^{\nu_{i2}} \frac{4\pi J_\nu}{h\nu} \sigma_{j\nu} d\nu}{\int_{\nu_{i1}}^{\nu_{i2}} \frac{4\pi J_\nu}{h\nu} d\nu}, \quad (44)$$

where

$$J_\nu = \frac{1}{4\pi} \int_{4\pi} I_\nu d\Omega. \quad (45)$$

Finally, s_{ij} is the source term that accounts for recombination radiation. Most of the previous RT implementations employ the On-The-Spot-Approximation (OTSA), which assumes that any radiation emitted by recombinations is immediately absorbed in the surroundings also called case B recombinations. This is a good

approximation in optically thick media, but is not valid in optically thin environments where the case A recombination rates are more relevant. Here, we include the option of not applying OTSA and define

$$s_{ij} = \begin{cases} 0 & \text{if OTSA} \\ \sum_j \delta_{ij} (\alpha_j^A - \alpha_j^B) n_j n_e & \text{Otherwise,} \end{cases} \quad (46)$$

where δ_{ij} is unity if the recombination radiation from species ‘ j ’ emits into the frequency bin ‘ i ’, else it is set to zero and α_j^A and α_j^B are the cases A and B recombination rates.

As commonly done, we use the operator split approach to solve these equations. First, the transport equations (setting RHS of equations 42 and 43 to zero) are solved as described in Section 3.1. Then, the thermochemistry equations are solved, which involves solving the equation for the change in the photon number density and photon number flux

$$\frac{\partial N_\gamma^i}{\partial t} = -\tilde{c} N_\gamma^i \left(\sum_j n_j \tilde{\sigma}_{ij} + \kappa_i \rho \right) + \sum_j s_{ij}, \quad (47)$$

$$\frac{\partial \mathbf{F}_\gamma^i}{\partial t} = -\tilde{c} \mathbf{F}_\gamma^i \left(\sum_j n_j \tilde{\sigma}_{ij} + \kappa_i \rho \right), \quad (48)$$

which are coupled with the equations which govern the number density evolution of the ionic species

$$\frac{dn_{\text{H II}}}{dt} = -\alpha_{\text{H II}} n_{\text{H II}} n_e + \sigma_{\text{eH I}} n_e n_{\text{H I}} + \tilde{c} n_{\text{H I}} \sum_i \tilde{\sigma}_{\text{iH I}} N_\gamma^i, \quad (49)$$

$$\begin{aligned} \frac{dn_{\text{He II}}}{dt} &= \alpha_{\text{He III}} n_{\text{He III}} n_e + \sigma_{\text{eHe I}} n_e n_{\text{He I}} \\ &+ \tilde{c} n_{\text{He I}} \sum_i \tilde{\sigma}_{\text{iHe I}} N_\gamma^i - \alpha_{\text{He II}} n_{\text{He II}} n_e \\ &- \sigma_{\text{eHe II}} n_e n_{\text{He II}} - \tilde{c} n_{\text{He II}} \sum_i \tilde{\sigma}_{\text{iHe II}} N_\gamma^i, \end{aligned} \quad (50)$$

$$\begin{aligned} \frac{dn_{\text{He III}}}{dt} &= -\alpha_{\text{He III}} n_{\text{He III}} n_e + \sigma_{\text{eHe II}} n_e n_{\text{He II}} \\ &+ \tilde{c} n_{\text{He II}} \sum_i \tilde{\sigma}_{\text{iHe II}} N_\gamma^i, \end{aligned} \quad (51)$$

where σ_{ej} and α_j are the collisional ionization and recombination rates of the ionic species ‘ j ’. These equations are supplemented with the following closure relations:

$$n_{\text{H}} = n_{\text{H I}} + n_{\text{H II}}, \quad (52)$$

$$n_{\text{He}} = n_{\text{He I}} + n_{\text{He II}} + n_{\text{He III}}, \quad (53)$$

$$n_e = n_{\text{H II}} + n_{\text{He II}} + 2n_{\text{He III}}. \quad (54)$$

We note that case B recombination rates are used if OTSA is applied otherwise, case A recombination rates are used. Finally, the dust reprocessed optical/UV radiation is added on to the IR energy density

$$\frac{\partial E_{\text{IR}}}{\partial t} = \sum_i^{\text{IR}} \kappa_i \rho \tilde{c} e_i N_\gamma^i, \quad (55)$$

where e_i is the mean energy per photon of the frequency group ‘ i ’ defined as

$$e_i = \frac{\int_{\nu_{i1}}^{\nu_{i2}} 4\pi J_\nu d\nu}{\int_{\nu_{i1}}^{\nu_{i2}} \frac{4\pi J_\nu}{h\nu} d\nu}, \quad (56)$$

and E_{IR} is the energy density of the infrared (IR) radiation field.

In addition to changing the ionization state of the gas, photons can also deposit energy through photoheating. Quantitatively, the photoheating rate (Γ) of a given ionic species is given as

$$\Gamma_j = \int_{\nu_{ij}}^{\infty} \frac{4\pi J_\nu}{h\nu} \sigma_{j\nu} (h\nu - h\nu_{ij}) d\nu, \quad (57)$$

where h is the *Planck* constant and $\mathcal{E}_j = h\nu_{ij}$ is the ionization potential of the ionic species ‘ j ’. In order to be compatible with a multifrequency approach, we need to discretize this equation into finite-frequency bins. We do so by splitting the above integral as follows

$$\Gamma_j = \sum_i \int_{\nu_{i1}}^{\nu_{i2}} \frac{4\pi J_\nu}{h\nu} \sigma_{j\nu} (h\nu - h\nu_{ij}) d\nu, \quad (58)$$

which can in turn be written down as

$$\Gamma_j = \tilde{c} \sum_i N_\gamma^i \bar{\sigma}_{ij} \epsilon_{ij}, \quad (59)$$

where the photoheating rate in the frequency bin ‘ i ’ due to the ionization of species ‘ j ’ is defined as

$$\epsilon_{ij} = \frac{\int_{\nu_{i1}}^{\nu_{i2}} \frac{4\pi J_\nu}{h\nu} \sigma_{j\nu} (h\nu - h\nu_{ij}) d\nu}{\int_{\nu_{i1}}^{\nu_{i2}} \frac{4\pi J_\nu}{h\nu} \sigma_{j\nu} d\nu}. \quad (60)$$

Therefore, the total amount of energy deposited into the gas through photoheating (\mathcal{H}) is then

$$\mathcal{H} = \sum_j n_j \Gamma_j. \quad (61)$$

Finally, the momentum injection rate by photon absorption is discretized in the same way and is given by

$$\frac{\partial \rho v}{\partial t} = \frac{1}{c} \sum_i \mathbf{F}_\gamma^i \left(\sum_j n_j \bar{\sigma}_{ij} p_{ij} + \kappa_i \rho e_i \right), \quad (62)$$

where

$$p_{ij} = \frac{\int_{\nu_{i1}}^{\nu_{i2}} 4\pi J_\nu \sigma_{j\nu} d\nu}{\int_{\nu_{i1}}^{\nu_{i2}} \frac{4\pi J_\nu}{h\nu} \sigma_{j\nu} d\nu}. \quad (63)$$

The numerical integration of this thermochemical network is quite challenging as small changes in the photon density can lead to rapid changes in the ionization state and temperature of the gas. Therefore, an explicit time integration of these equations would require very small time-steps making the thermochemistry step computationally expensive. In order to overcome this problem, we use a variant of the method outlined in Petkova & Springel (2009), which employs a semi-implicit approach, that first solves for the number density evolution of the ionic species implicitly using the values of the n_e and N_γ^i from the previous time-step. n_e is then updated with the revised values of the number density of the ionic species according to equation (54). The change in temperature of

the gas is then calculated explicitly using the updated abundances of the ionic species (a detailed description of this scheme is discussed in Appendix B).

If the temperature or one of the abundances changes by more than 10 per cent during a time-step, then we resort to using the publicly available ordinary differential equation solver SUNDIALS CVODE (Hindmarsh et al. 2005) which employs a variable order, variable step, and multistep backward differencing scheme to compute the temperature and chemical abundances. This scheme is stable enough to allow us to set the thermochemistry to the transport time-step given by equation (39).

3.2.2 Infrared dust–gas coupling

Dust grains are solid, macroscopic particles composed of dielectric and refractory materials. Many of the physical details are empirical as we do not yet know the precise composition of dust grains, nor do we know their precise physical properties. They do however, play a major role in the physics of the ISM. Although, they only make up about 1 per cent of the ISM (Gilmore, Wyse & Kuijken 1989; Zubko, Dwek & Arendt 2004), they absorb and reprocess almost 50 per cent of the starlight in the galaxy (Battisti, Calzetti & Chary 2016). The surface of dust grains host a variety of chemical reactions that changes the chemical composition of the ISM significantly, which in turn leads to changes in the star formation rate. They are also thought to be in radiative equilibrium with the local IR radiation field (Krumholz & Thompson 2012).

Modelling the coupling between IR radiation, gas, and dust can be quite challenging. Ideally, this requires treating the dust as a separate fluid and accurately accounting for the energy and momentum exchange between gas, dust, and the IR radiation field, which can get rather complicated. We instead choose to couple the IR RT scheme to the semi-empirical dust model of McKinnon, Torrey & Vogelsberger (2016). This model accounts for the stellar production of dust, accretion of gas-phase metals on to existing grains, destruction of dust through local SN activity, thermal sputtering of dust, and dust driven by winds from star-forming regions. It reproduces the dust content in low-redshift galaxies (McKinnon et al. 2017).

In this model, the dust is treated as a passive scalar, whose motion is coupled to the gas, which is a good approximation for short stopping time-scales found in the ISM. We further assume that the system is close to local thermodynamic equilibrium, which is good approximation for cold high-density regions of the ISM (Goldsmith 2001; Krumholz & Thompson 2013). Under these conditions, the gas emits as a blackbody. If the IR photon frequency bin covers a sufficiently large range, then the source function can be approximated by the frequency integral of a *Planck* spectrum

$$S_{\text{IR}} = \kappa_p \rho c a T^4 \quad (64)$$

where a is the radiation constant, κ_p is the *Planck* mean opacity, and T is the temperature of the gas. If the IR spectral energy distribution (SED) is dominated by reprocessed radiation, then it can be approximated to a Planckian, implying that $\kappa_E = \kappa_p$. Additionally, for a system with a large IR optical depth, we can assume that the flux-weighted mean opacity is similar to the Rosseland mean opacity $\kappa_F \sim \kappa_R$. Using these assumptions, we can then write equations (3) and (4) as

$$\frac{\partial E_{\text{IR}}}{\partial t} + \nabla \cdot \mathbf{F}_{\text{IR}} = \kappa_p \rho (c a T^4 - \tilde{c} E_{\text{IR}}) \quad (65)$$

$$\frac{\partial \mathbf{F}_{\text{IR}}}{\partial t} + \tilde{c}^2 \nabla \cdot \mathbb{P}_{\text{IR}} = -\kappa_{\text{R}} \rho \tilde{c} \mathbf{F}_{\text{IR}}. \quad (66)$$

The dust opacities (for both IR and UV radiation bins) are calculated self-consistently based on the empirical dust model (see Appendix C for more details). The destruction of dust by the UV radiation field is not included as it is sub-dominant compared to other dust destruction mechanisms such as SNe and thermal sputtering and is only important in highly luminous systems with an extremely hard spectra such as gamma-ray bursts (Waxman & Draine 2000; Draine & Hao 2002).

Using the usual operator split approach, we first solve the pure transport equations (setting the RHS of equations 65 and 66 to zero) using the algorithm described in Section 3.1. The source terms are then

$$\frac{\partial E_{\text{IR}}}{\partial t} = \kappa_{\text{P}} \rho (c a T^4 - \tilde{c} E_{\text{IR}}), \quad (67)$$

$$\frac{\partial \mathbf{F}_{\text{IR}}}{\partial t} = -\kappa_{\text{R}} \rho \tilde{c} \mathbf{F}_{\text{IR}}, \quad (68)$$

and the change in the internal energy (u) and momentum of the gas are given by

$$\frac{\partial u}{\partial t} = -\kappa_{\text{P}} \rho (c a T^4 - \tilde{c} E_{\text{IR}}), \quad (69)$$

$$\frac{\partial \rho v}{\partial t} = \frac{\kappa_{\text{R}} \rho}{c} \mathbf{F}_{\text{IR}}. \quad (70)$$

Equations (67) and (69) form a set of coupled equations. The fourth power dependence on the temperature makes the equations stiff requiring very small time-steps to maintain stability of the solution. We use a variant of the approach presented in Rosdahl & Teyssier (2015, R15) to solve these equations. First, we get a semi-implicit estimate of the change in $U_{\text{E}} \in (E_{\text{IR}}, u)$

$$\Delta U_{\text{E}} = \dot{U}_{\text{E}} \Delta t (\mathbf{I} - \mathbf{J} \Delta t)^{-1} \quad (71)$$

where Δt is the time-step and $\mathbf{J} = \frac{\partial \dot{U}_{\text{E}}}{\partial U_{\text{E}}}$ is the Jacobian matrix. Using the symmetry of the problem ($\Delta E_{\text{IR}} = -\Delta u$) the change over time-step Δt can be written as

$$\Delta E_{\text{IR}} = -\Delta u = \frac{c a T^4 - \tilde{c} E_{\text{IR}}}{(\kappa_{\text{P}} \rho \Delta t)^{-1} + \tilde{c} + 4 c a T^3 C_v^{-1}} \quad (72)$$

where $C_v = k_{\text{B}}/((\gamma - 1)\mu)$ is the specific heat at constant volume, γ is the adiabatic index, and μ is the mean molecular weight. If the relative change in both E_{IR} and u is less than 10 per cent we keep this estimate for the solution. Otherwise, we discard it and switch to solving these equations using the SUNDIALS CVODE library as mentioned in the previous section. This method is quite stable and avoids large iterative loops even when sudden changes in the photon energy density occur.

4 TEST PROBLEMS

This section presents various test problems of our RHD implementation. The relevant hydrodynamic equations are evolved using the moving mesh finite-volume scheme outlined in Springel (2010), with the improved time-integration and gradient estimation techniques described in Pakmor et al. (2016).

In Section 4.1, we test the accuracy and the convergence order of our scheme. Sections 4.2 and 4.3 discuss a test for the ionizing chemistry scheme (Section 3.2.1). Section 4.4 tests the UV chemistry and cooling implementation and the ability of the scheme to capture and maintain the directionality of the underlying photon field. The simulations presented in Sections 4.5 and 4.6 are evolved using the full RHD equations and test the coupling between the photons and the gas. Section 4.7 investigates how the M1 closure relation performs in a multisource set-up. Sections 4.8 and 4.9 analyse the performance of the multiscattered IR+coupled dust implementation (Section 3.2.2), specifically, the ability of our scheme to capture accurate results in optically thick media. Finally, Section 4.10 tests the temperature coupling multiscattering, coupling to the hydrodynamics module, and explores the competition between gravity and radiation pressure. The simulations in Sections 4.1–4.4, and 4.7–4.9 are run without hydrodynamics and are intended to test the accuracy and the stability of the RT implementation only. In Sections 4.5, 4.6, and 4.10 we additionally test the accuracy of the our scheme when hydrodynamics and ionizing RT are coupled. In full RHD tests, we initially start with a regular Cartesian grid, which is then allowed to move and adapt to the local fluid motion (Springel 2010). In addition, the mesh is regularized where needed using the scheme outlined in Vogelsberger et al. (2012).

In many of these tests, we also vary the numerical scheme of the basic radiation transport equations. We explore the differences between the flux functions used to solve the Riemann problem at the cell interfaces, specifically the differences between the GLF and the HLL flux functions. The HLL flux function uses accurately calculated eigenvalues that represent the wave speeds of system of transport equations. These depend highly on the reduced flux and the angle between the photon flux and the cell interface. The GLF flux function on the other hand uses a single wave speed set equal to the reduced speed of light in our calculations. So theoretically, the HLL flux function is less diffusive and better maintains the directionality of the photon field. Additionally, we also compare the accuracy and stability of the PC and the linear gradient extrapolation methods. The PC scheme assumes that the primitive variables do not vary within a volume element and therefore the input values for the Riemann solver are the cell centred values. On the other hand, the gradient extrapolation scheme linearly extrapolates the cell centred values on to the face of the cell and then solve the Riemann problem with these extrapolated values (see Section 3.1 for more information). In summary, we test the following four numerical transport schemes:

- (i) **PC-GLF**: piecewise constant approximation with GLF flux function.
- (ii) **PC-HLL**: piecewise constant approximation with HLL flux function.
- (iii) **GLF**: linear extrapolation with GLF flux function.
- (iv) **HLL**: linear extrapolation scheme with HLL flux function.

We note that our fiducial scheme is the HLL scheme.

Finally, we also explore the performance of our code on different mesh configurations. We mainly use three different mesh geometries, namely a regular Cartesian mesh, a regular staggered mesh constructed by two Cartesian meshes that are displaced from each other by $0.45 \Delta x$ in every direction, where Δx is the cell size for the given resolution and an irregular mesh, which is obtained from the Cartesian mesh by adding a random offset of up to $0.2 \Delta x$, mimicking the typical maximum deviation between mesh-generating points and cell centres in real problems (Vogelsberger




Figure 1. Radiation wave propagation: L^1 error after one complete wave period as a function of the number of resolution elements (N) for different numerical schemes. While the PC approximation converges very slowly with a slope of $\lesssim 1.0$, our fiducial scheme converges much faster with a slope of ~ 2.0 , irrespective of the flux function used.

et al. 2012). Finally, please note that we use the labels ‘regular mesh’ and ‘staggered mesh’ interchangeably.

4.1 Radiation wave propagation

We start with a test to gauge the accuracy of our RT scheme. We investigate the propagation of small-amplitude, free-streaming radiation wave in a purely absorbing, homogeneous medium with low optical depth, like the test described in Gardiner & Stone (2005) and Skinner & Ostriker (2013). A 2D box of sidelength $L_{\text{box}}\{x, y\} = \{2, 1\}$ is initialized with a photon number density as follows

$$E(r) = E_{\text{bg}} + \epsilon \sin\left(\frac{2\pi(x+2y)/\sqrt{5}}{\lambda}\right), \quad (73)$$

where E_{bg} is a uniform background photon density field, $\epsilon = 10^{-6}$, and $\lambda = 2/\sqrt{5}$. The radiation flux points in the direction of the wave and has a value $|\mathbf{F}_r| = cE_r$, with $c = 1$. The optical depth per wavelength is set to $\tau_\lambda = \rho\kappa\lambda = 0.1$ and periodic boundary conditions are used everywhere.

We employ four different background grids a regular Cartesian grid, a Cartesian grid with 1 per cent deviation between the mesh-generating point and the mesh centroid, Cartesian grid with 10 per cent deviation, and finally a Cartesian grid with 20 per cent deviation, which we denote as an irregular mesh as it mimics the typical maximum deviation between mesh-generating points and cell centres in real problems. We start the simulation at time $t = 0$ and evolve the system until $t = 0.89$ so that there is one complete wave period in each of the x - and y -directions. The solution after one complete wave period is then given by

$$E(r) = E_{\text{bg}} + \epsilon e^{-\rho\kappa t} \sin\left(\frac{2\pi(x+2y)/\sqrt{5}}{\lambda}\right). \quad (74)$$




Figure 2. Strömgren sphere: a projected map of the ionized H fraction in the Strömgren sphere test for the highest resolution (2×128^3) simulations performed using HLL (top panels) and GLF (bottom panels) functions with underlying regular (left-hand panels) and irregular (right-hand panels) meshes. We obtain quite spherical Strömgren regions, because the photon injection region is well resolved.

The L^1 norm measured for different resolutions for the PC and fiducial schemes is shown in Fig. 1. The PC-GLF and PC-HLL schemes converge rather slowly with a convergence order of $\lesssim 1.0$. The fiducial schemes on the other hand start out with much smaller errors and converge significantly faster towards the analytic solution with an order of ~ 2.0 . There is essentially no dependence of the accuracy of the solution on the type of the flux function used. The convergence order decreases as we move to more distorted meshes, especially at higher resolutions, which shows the importance of using mesh regularization schemes in AREPO.

These results illustrate the advantages of our algorithm. It is less diffusive and highly accurate even at relatively low resolutions. A gradient extrapolated predictor–corrector scheme coupled to a Strang split approach for sources terms is necessary to achieve second-order accuracy.

4.2 Strömgren sphere

We now test our scheme for the classical problem of an H II region expanding in a constant density and temperature medium (Strömgren 1939; Spitzer 1978). This tests the radiation transport scheme and the H chemistry implementation in parallel. A single monochromatic radiation source is placed at the centre of a domain of sidelength 16 kpc, which is emitting H ionizing photons with an energy 13.6 eV. The source outputs a constant stream of photons at a rate of $\dot{N}_\gamma = 5 \times 10^{48} \text{ photons s}^{-1}$. The density of the surrounding gas is set to $n_{\text{H}} = 10^{-3} \text{ cm}^{-3}$ and the gas has




Figure 3. Strömgren sphere: profiles of the neutral (solid curves) and ionized (dashed curves) H fractions as a function of radius at the end of the Strömgren sphere test for the simulations with 2×128^3 (red curves), 2×64^3 (green curves), and 2×32^3 (blue curves) runs. The simulations match the analytic results (black curves) quite well.

temperature $T = 10^4$ K. The case B recombination rate of H II under these conditions is $\alpha_B = 2.69 \times 10^{-13} \text{ cm}^3 \text{ s}^{-1}$. Assuming all the emitted photons are used to ionize the surrounding H the maximal extent of the Strömgren radius ($r_s, 0$) is given by

$$r_s = \left(\frac{3\dot{N}_\gamma}{4\pi\alpha_B n_H^2} \right)^{1/3} = 5.38 \text{ kpc}. \quad (75)$$

The evolution of the radius of the I-front can be obtained by assuming that it is infinitely thin

$$r_{I,0}(t) = r_s [1 - \exp(-t/t_{\text{rec}})]^{1/3}, \quad (76)$$

where

$$t_{\text{rec}} = \frac{1}{n_H \alpha_B} = 125.13 \text{ Myr}, \quad (77)$$

is the recombination time for our choice of the gas density.

The radial profiles of the neutral ($\tilde{n}_{\text{HI}} = n_{\text{HI}}/n_{\text{H}}$) and ionized H (\tilde{n}_{HII}) fractions can analytically be computed from (Osterbrock & Ferland 2006)

$$\frac{\tilde{n}_{\text{HI}}(r)}{4\pi r^2} \int d\nu \dot{N}_\gamma(\nu) e^{-\tau_\nu(r)} \sigma_\nu = \tilde{n}_{\text{HII}}^2(r) n_{\text{H}} \alpha_B, \quad (78)$$

where

$$\tau_\nu(r) = n_{\text{H}} \sigma_\nu \int_0^r dr' \tilde{n}_{\text{HI}}(r'). \quad (79)$$

We perform the simulation with 2×128^3 , 2×64^3 , and 2×32^3 resolution elements using our fiducial scheme on a regular mesh. We use the RSLA in these runs, with $f_r = \tilde{c}/c = 0.01$. In order to understand the effect of reducing the speed of light, we perform additional simulations with $f_r = 10^{-4}$, 2×10^{-4} and 10^{-3} . Fig. 2

shows the projected maps of ionized H, at 0.5 Gyr for the run with 2×128^3 resolution elements. There are small departures from a perfect spherical symmetry due to the geometry of the injection region and to a lesser extent due to the geometry of the underlying mesh. Resolving the injection radius with multiple cells improves the sphericity of the solution.

The profiles of neutral (solid curve) and ionized (dashed curve) H at 0.5 Gyr are shown in Fig. 3. The analytic solution obtained from equations (78) and (79) is plotted in black. We can see that our method reproduces the analytic solution and the accuracy of the solution increases with resolution as expected.

In Fig. 4, we show the time evolution of the ionizing front, for the runs with different light speeds. The ionization front is defined as the radius at which the ionization fraction equals 0.5. We note that at late times, the analytic expectations derived from equation (76, solid black curve) and equation (78, dashed black curve) diverge because equation (76) assumes an infinitely thin transition region and fully ionized gas within r_s which is not an accurate description of the Strömgren sphere. The simulation results match the more accurate analytic expectation described by equation (78).

The position of the I-front after 0.5 Gyr is quite similar in all the runs, implying that, given time, the simulations will converge to the right solution irrespective of the speed of light used. However, this is only true if we are interested in the final state of the Strömgren sphere and not in its evolution, which is clearly different for the different speeds of light used. As shown in R13, the behaviour of the solution can be quantified by comparing the light crossing time ($t_{\text{cross}} = r_s/\tilde{c}$) to the recombination time (t_{rec} , equation 77); $q = t_{\text{cross}}/t_{\text{rec}}$. When, $q \geq 1$, the ionization front expands at the reduced speed of light and reaches r_s at $t = t_{\text{cross}}$ and this is insensitive to the recombination timescale. For $q < 1$, the reduced speed of light is larger than the analytic expectation of the speed of the ionization front and hence the simulated ionization front travels at the reduced speed of light till it reaches the analytic solution and starts following the analytic results. Of course, these considerations are only valid if we assume full ionization within r_s . In most realistic situations, there is a complex ionization structure given by equation (78) and shown in Fig. 3. In these realistic situations, the time to reach the analytic solution is longer. For example, in our simulation with $f_r = 10^{-4}$ ($q = 1.4$), the ionization front should start to follow the analytic solution by $t/t_{\text{rec}} = 1.4$, but it takes $t/t_{\text{rec}} > 3.5$ to achieve this because of the complicated internal ionization structure of the Strömgren sphere.

If τ_{min} is the shortest relevant time-scale of a simulation, then ensuring that $t_{\text{cross}} < \tau_{\text{min}}$ will ensure that the right solution for the evolution of the I-front will be recovered without affecting other time-scales of the system. Therefore, we can typically set

$$f_r = \min\{1, \eta t_{\text{cross}}/\tau_{\text{min}}\} \quad (80)$$

where $\eta \simeq 10$. The value of τ_{min} is of course problem dependent. For reionization simulations, where the speed of the ionization front can reach speeds of 10^4 km s^{-1} (Bauer et al. 2015), there is very little room to reduce the speed of light. However, if we want to simulate the ISM, then the high densities and low gas velocities allows for large reductions in the speed of light $f_r \lesssim 10^{-3}$. Since most of the tests in the paper are performed for relatively high-density initial




Figure 4. Strömgren sphere: radius of the ionization front as a function of time for simulations with 2×32^3 – resolution elements performed with $f_r = 10^{-4}$ (cyan curve), $f_r = 2 \times 10^{-4}$ (blue curve), $f_r = 10^{-3}$ (green curve), and $f_r = 0.01$ (red curve). The analytic solutions obtained from equations (76) and (78) are plotted as solid black and dashed black curves, respectively. The run with the highest speed of light matches the analytic expectations well, while the runs with lower light speeds take longer to reach the expected analytic r_s .

conditions, a value of $f_r = 0.01$ is good enough to capture accurate results.³

4.3 Multifrequency H–He Strömgren sphere

As an extension of the previous test, we simulate the ionization structure around an O4V star. For this test, we include both atomic H and He chemistry. In order to appropriately model the absorption of photons at different frequencies by the different ionic species, a multifrequency approach is required. Qualitatively, the ionization structure in a nebula with H and He depends on the He abundance as well as the spectrum of the ionizing star. The first ionization potential of He is 24.6 eV and it can be doubly ionized when photons with energies greater than 54.4 eV are present. The photons with energies between 13.6 and 24.6 eV can ionize only neutral H while photons with energies above 24.6 eV can ionize both neutral H and He. In order to accurately capture this behaviour, we split the spectrum into three bins with frequency ranges corresponding to the ionization potentials of H I [13.6 eV, 24.6 eV], He I [24.6 eV, 54.4 eV], and He II [54.4 eV, 100.0 eV].

The star is assumed to emit a blackbody spectrum with $T_{\text{eff}} = 4.87 \times 10^4$ K and a luminosity of $7.6 \times 10^5 L_{\odot}$, which translates to a Lyman continuum photon rate of 5×10^{49} photons s^{-1} . We simulate a domain of size 3 pc with 2×128^3 resolution elements and a regular staggered grid. The surrounding gas has a constant density of 10^3 cm^{-3} and a temperature of 10^4 K.

³See R13 and Skinner & Ostriker (2013) for a more detailed discussion on the RSLA and its applicability.




Figure 5. Multifrequency H–He Strömgren sphere: profiles of the neutral H (solid red curve), ionized H (dashed red curve), neutral He (dot dashed green curve), singly ionized He (solid green curve) and doubly ionized He (dashed green curve) fractions as a function of radius around a O4V star at the end of the multifrequency H–He Strömgren sphere test. The analytic expectations of doubly ionized He and neutral He are over plotted as black curves.

The approximate size of the He II zone is given by (Tielens 2005)

$$r_s(\text{He}) = \left(\frac{3\dot{N}_{\gamma, \text{Lyc}}(\text{He})}{4\pi y(1+y)\alpha_B(\text{He})n_H^2} \right)^{1/3}, \quad (81)$$

where ‘y’ is the He number fraction and for this test $r_s(\text{He}) \sim 1$ pc. This estimate is only correct, if there is no H ionization. For realistic conditions, we need to solve the RT equation and the coupled chemistry equations as described in Section 3.2.1.

The fractional ionization profiles H I (solid red curve), H II (dashed red curve), He I (dotted–dashed green curve), He II (solid green curve), and He III (dashed green curve) are plotted in Fig. 5. He is doubly ionized in the central 0.1 pc, while both H and He are singly ionized within the central 1 pc. Even a very hot O4V star is only able to produce a small doubly ionized He region as the amount of photons above 54.4 eV is rather low.

As a comparison, we also plot the analytic fractional ionization profiles of He III and He I in black. While the position of the transition region matches the analytic solution well, the width is larger than the expected value. This is because the width of the transition region is of the order of the mean-free path of ionizing photons,

$$l = 1/(n_H \sigma_e) \sim 10^{-3} \text{ pc}, \quad (82)$$

which is much smaller than the cell size in our simulation. As the ionization cross-section is highly peaked towards the ionization potential, the stellar radiation field is least attenuated at the highest frequencies. Since the H cross-section for highly energetic photons is smaller than that of He, He can stay ionized slightly further out than H as seen in the figure. We have only plotted the analytic expectation for two of the five ionic species in the simulation, in order to increase the readability of the plot. We note that we match




Figure 6. I-front trapping in a dense clump and the formation of a shadow: The H I maps showing slices at $z = 0.5L_{\text{box}}$ after 1 Myr of evolution for HLL (first column), PC-HLL (second column), GLF (third column), and PC-GLF (fourth column) schemes. The top panels depict the results for simulations using a underlying Cartesian mesh, while the bottom panels depict the results for a staggered mesh.




Figure 7. I-front trapping in a dense clump and the formation of a shadow: same as Fig. 6 but at simulation time of $t = 15$ Myr. The fiducial HLL and GLF schemes are able to form sharp shadows irrespective of the mesh geometry used. The PC-HLL scheme is only able to obtain sharp shadows if the mesh interfaces are exactly parallel/perpendicular to the photon propagation direction as is the case in a Cartesian mesh, but fails to do so in a staggered mesh which has cross-mesh transport. The PC-GLF scheme is unable to form sharp shadows irrespective of the mesh geometry.

the profiles of the other ionic species equally well (see figure 7.2 of Tielens 2005 for more details). This test confirms the accuracy of our multifrequency RT scheme coupled to the H–He chemistry.

4.4 I-front trapping in a dense clump and the formation of a shadow

Next we simulate the trapping of a plane-parallel I-front by a dense, uniform, spherical clump. This test mimics self-shielding within a high-density gas cloud illuminated by ionizing UV photons. Ideally, the other side of the clump should be shielded from the ionizing

radiation, producing sharp shadows. The ability of a clump to trap an ionization front depends on the strength of the ionizing flux (F), the clump density (n_{H}), the radius of the clump (r_{clump}), and the case B recombination rate (α_{B}). We can define the ‘Strömgren Number’ for a clump as $L_{\text{s}} = 2r_{\text{clump}}/l_{\text{s}}(0)$, where $l_{\text{s}}(0)$ is the Strömgren length at zero impact parameter

$$l_{\text{s}}(0) = \frac{F}{\alpha_{\text{B}} n_{\text{H}}^2}, \quad (83)$$

$$L_{\text{s}} = \frac{2r_{\text{clump}} \alpha_{\text{B}} n_{\text{H}}^2}{F}. \quad (84)$$




Figure 8. I-front trapping in a dense clump and the formation of a shadow: the temperature maps showing slices at $z = 0.5L_{\text{box}}$ after 1 Myr of evolution.




Figure 9. I-front trapping in a dense clump and the formation of a shadow: same as Fig. 8 but at simulation time of $t = 15$ Myr. The temperature maps replicate the trends seen in the H I maps.

If $L_s > 1$, the clump is able to trap the I-front, while if $L_s < 1$, the clump would be unable to trap the I-front and would instead be flash ionized by its passage.

In our setup, the UV ionizing radiation has a blackbody spectrum with an effective temperature $T_{\text{eff}} = 10^5$ K and an ionizing flux of $F = 10^6$ photons $\text{s}^{-1}\text{cm}^{-2}$. This flux is a plane-parallel wave travelling in the $+x$ -direction and is incident on the $x = 0$ boundary. The domain size is 6.6 kpc on a side with 128^3 resolution elements. A spherical high-density clump of radius $r_{\text{clump}} = 0.8$ kpc is placed within the domain, with the centre of the clump at $(x, y, z) = (5, 3.3, 3.3)$ kpc. The density and temperature of the clump are $n_{\text{H}}^{\text{clump}} = 4 \times 10^{-2} \text{ cm}^{-3}$ and $T^{\text{clump}} = 40$ K, respectively. The remaining domain is filled with a hot, low-density gas with a temperature of $T^{\text{out}} = 8000$ K

and density of $n_{\text{H}}^{\text{out}} = n_{\text{H}}^{\text{clump}}/200 = 2 \times 10^{-4} \text{ cm}^{-3}$. Initially, the domain is neutral with an ionization fraction of $\bar{n}_{\text{H II}} = 10^{-6}$. For these parameters and assuming a case B recombination rate of $\alpha_{\text{B}}(T) = 2.59 \times 10^{-13} (T/10^4 \text{ K})^{-3/4} \text{ cm}^3 \text{ s}^{-1}$, we obtain $l_s \simeq 0.78 (T/10^4 \text{ K})^{3/4} \text{ kpc}$ and $L_s \simeq 2.05 (T/10^4 \text{ K})^{-3/4}$. Therefore, along the axis of symmetry the I-front should be trapped approximately at the centre of the clump for $T = 10^4$ K. In reality, the temperature could be expected to be somewhat different and spatially varying, but to a rough first approximation this estimate should hold. We run the simulation for 15 Myr with $f_r = 0.1$. A single frequency approximation is used, so that there is no leakage of photons beyond the I-front.

We perform this simulation with four different types of numerical schemes, HLL, PC-HLL, GLF, and PC-GLF, which are described




Figure 10. I-front trapping in a dense clump and the formation of a shadow: the ionization fraction (top panels) and temperature (bottom panels) profiles calculated in a thin cylindrical shell around the axis of symmetry after 1 Myr (left-hand panels) and 15 Myr (right-hand panels) of evolution. The solid curves depict the results for a simulation with an underlying Cartesian mesh, while the dashed curves are for a staggered grid. The red, green, blue, and yellow curves show the result for HLL, PC-HLL, GLF, and PC-GLF schemes respectively. Numerical photon diffusion increases the temperature and H II fraction in the shadow region if the PC schemes are used.

in Section 4. We note that ‘PC’ denotes a piecewise-constant approximation, which is used in many recent works (Rosdahl & Teyssier 2015; Bieri et al. 2017; Costa et al. 2017; Cielo et al. 2017; Lupi et al. 2017). However, we have shown in Section 4.1 that this approximation is exceedingly diffusive. Just switching to a less diffusive HLL flux function from GLF will only slightly improve the numerical diffusivity of the scheme. We have shown that it is important to have gradient extrapolated values at the interface, in order to reduce the numerical diffusivity and increase the convergence order of the scheme. We use this test to further show the limitations of the ‘PC’ approximation and demonstrate the importance of using higher order schemes for RT. We use a reduced light speed fraction of $f_r = 0.1$. This high speed is needed for the light to have reached the cloud in the first snapshot under consideration, at 1 Myr.

The I-front travels fast through the diffuse medium outside the cloud, but moves much more slowly inside of it, and a shadow is cast behind it. As the UV radiation slowly ionizes and heats the cloud, the shadow very slowly diminishes in width because some photons manage to cross through the edges of the cloud. Fig. 6 presents the H I fraction maps in the simulations using HLL (first column), PC-HLL (second column), GLF (third column), and PC-GLF (fourth column) schemes after 1 Myr of evolution. The top panel denotes the results where the underlying mesh is Cartesian. Fig. 7 shows the same after 15 Myr of evolution. These maps reveal interesting differences between the numerical schemes. On a regular Cartesian mesh, the HLL, PC-HLL, and GLF schemes are able to maintain the directionality of the photons, even after 15 Myr of evolution resulting in accurate shadows, which is not true for the PC-GLF




Figure 11. I-front trapping in a dense clump and the formation of a shadow: the time evolution of the average ionization fraction (top panels) and temperature (bottom panels) of the gas clump for simulations using an HLL (red curve), PC-HLL (green curves), GLF (blue curves), and PC-GLF (yellow curves) schemes with an underlying Cartesian (left-hand panels) and staggered (right-hand panels) grids. The more diffusive PC schemes tends to increase the temperature and ionization fraction of the high-density clump.

scheme. This verifies the results presented in Rosdahl et al. (2013; hereafter referred to as R13). Based on this test R13 and Lupi et al. (2017) argued that switching to an HLL flux function will reduce the diffusion of the scheme and maintains the directionality of the photons. The directionality of the photons is only exactly maintained when the photon flux is either completely parallel or perpendicular to the cell interface. The eigenvalues of the system are exactly $\pm c$ and 0 when the angle between the photon flux and the interface is 90 and 0 deg, respectively, and the value of the reduced flux $f = 1$ (see fig. 1 of González et al. 2007). In this test, the photon flux is exactly perpendicular to the cell interfaces along the photon propagation direction and parallel to the other faces, in a Cartesian mesh. This necessitates that the perpendicular numerical diffusion is exactly 0 by construction.

However, this geometry can only be achieved in test problems where the radiation propagation direction is known beforehand and will almost never occur in realistic simulations. To elucidate this point, we perform the simulations also on a regular staggered mesh (bottom panels of Figs 6 and 7) with 2×96^3 resolution elements. The photon flux is no longer parallel/perpendicular to the cell interfaces, therefore the PC-HLL scheme no longer produces accurate and sharp shadows, and this is particularly evident after 15 Myr. It still performs better than the PC-GLF scheme, but performs unfavourably compared to fiducial schemes, which use gradient extrapolations, regardless of the flux function used. These results reflect the findings of Section 4.1 for a more realistic setup. Since, the temperature of the gas has a sharp dependence on the photon density of the cell, the temperature maps (Figs 8 and 9) show similar trends.

Fig. 10 quantifies the difference between the various schemes by plotting the H ionization fraction (top panels) and temperature

(bottom panels) profiles calculated in a thin cylindrical shell around the axis of symmetry after 1 Myr (left-hand panels) and 15 Myr (right-hand panels) of evolution. The temperature of the surrounding gas increases slightly from 8000 K to about 3×10^4 K very quickly. Inside the clump, however, the I-front is trapped and its velocity slows down considerably. In the pre-ionization zone ahead of the main I-front, all schemes agree quite well at all times. Slight differences emerge in the position of the I-front (defined as the point of 50 per cent ionized fraction) at 15 Myr, with the less diffusive fiducial schemes propagating to a slightly smaller distance into the cloud.

The main differences arise in the shadow region behind the clump. The temperature and the ionization fraction, in particular, of this low-density gas is extremely sensitive to any photons leaking into this region due to numerical diffusion. On a Cartesian mesh (solid curves), the HLL (red curves), GLF (blue curves), and the PC-HLL (green curves) schemes all perform quite well in the shadow region. The ionized H fraction stays at 10^{-6} indicating very little numerical diffusion perpendicular to the direction photon flow. The PC-GLF (yellow curves) scheme, however, has difficulty maintaining the directionality of the photons and this drastically increases the ionization fraction and temperature in the shadow region. In fact the gas in the shadow region becomes fully ionized behind the clump at 15 Myr. The ability of the PC-HLL scheme to produce accurate shadows is diminished as soon as we move to a staggered mesh (dashed curves), which mimics cross-mesh transport of photons. The ionization fraction in the shadow region increases almost as dramatically as the PC-GLF scheme implying that there is no longer a sharp shadow. Using an HLL flux function only slightly improves the numerical accuracy over using a GLF flux function if a PC approximation is used. Only our fiducial gradient extrapolated schemes are able to maintain low numerical diffusivity even when the photon propagation direction is across the mesh interfaces.

Finally, Fig. 11 shows the time evolution of the average ionization fraction and temperature of the clump. The evolution matches in general the range of results seen in tests performed using various RT schemes presented in Iliev et al. (2006). This confirms that the position and velocity of the I-front match the expected solution very well. The evolution in the average quantities only slightly differs between the schemes, with the more diffusive PC-GLF (on both meshes) and PC-HLL (only on a staggered mesh) schemes showing slightly higher ionization fractions and temperatures as expected.

We conclude that this test demonstrates the importance of implementing higher order, low diffusion schemes for RT in order to produce accurate results that are independent of the geometry of the problem and that of the underlying mesh.

4.5 Expansion of a H II region

In this test, we explore the problem of the expansion of an I-front due to photoheating from a point source (tests 5 and 6 of Iliev et al. 2009). The temperature of the gas is allowed to vary and the hydrodynamics is switched on. The photons heat the region around the source through photoheating (equations 12 and 61), producing an overpressurized region which drives the gas out with a certain velocity. The I-fronts are generally classified by comparing their speed, to two critical speed of the gas: R-critical, defined as $v_R = 2c_{s,1,2}$, and D-critical, given by $v_D \sim c_{s,1,2}^2 / (2c_{s,1,2})$, where $c_{s,1,2} = (p_1/\rho_1)^{1/2}$ and $c_{s,1,2} = (p_2/\rho_2)^{1/2}$ are the isothermal sound speeds in the gas ahead of and behind the I-front, respectively.

Typically, the I-front is initially R-type ($V_I \geq V_R$), where it expands supersonically with respect to the neutral gas ahead, which means RT post-processing is a fairly good approximation. The I-front then begins to slow down once it approaches the Strömgren radius. When $v_D < v_I < v_R$ (sometimes referred to as an M-type I-front), the I-front is necessarily led by a shock which compresses the gas entering the I-front sufficiently to slow it down and guarantees that it is converted to a D-type front ($v_I \leq v_D$).

We test the performance of our implementation under two different physical conditions, first, we simulate the expansion of the I-front in a uniform density medium and then move on to a more realistic situation where a source is at the centre of a spherically symmetric, steeply decreasing power-law density profile with a small flat central core. For the first simulation, we initialize a box of size $2L_{\text{box}} = 30$ kpc on a side, which is resolved initially with 2×80^3 resolution elements placed in a regular staggered grid. As the gas starts to move, the mesh is allowed to move and distort according to the local fluid motion. The box is initialized with a pure H gas of density and temperature $n_H = 10^{-3} \text{ cm}^{-3}$ and $T = 100$ K, respectively. A constant luminosity source is placed at the centre of the domain that emits a blackbody spectrum with $T_{\text{eff}} = 10^5$ K at the rate of $\dot{N}_\gamma = 5 \times 10^{48} \text{ photons s}^{-1}$. We use a reduced speed of light with $f_r = 0.01$ and the run the simulation for 500 Myr. A multifrequency RT scheme is employed, where, the emitted photons are grouped into three separate bins with frequency ranges [13.6 eV, 24.6 eV], [24.6 eV, 54.4 eV], and [54.4 eV, 100.0 eV].

Fig. 12 presents cross-section (at $z = 15$ kpc) maps of the H I fraction (first column), density (second column), temperature (third column), pressure (fourth column), and Mach number ($\mathcal{M} = v/c_s$; fifth column) at 10 Myr (top panels), 200 Myr (middle panels), and 500 Myr (bottom panels). Although the radiation is able to ionize and heat the gas in the central ~ 2 kpc within 10 Myr, there is very little change in the density of the gas. The gas velocity is still very low implying that it has not yet started evacuating the central regions. By 200 Myr, the overpressurized region manages to push enough gas out of the central regions and a high-density expanding gas shell is formed around a low-density region. Interestingly, a second transient density peak forms beyond the I-front, which is replicated in the pressure and Mach number maps. The I-front expands almost to L_{box} by 500 Myr. We note that the maps show some asymmetric artifacts that can be attributed to the geometry of the underlying mesh. In our experiments, these kind of features emerge only when there is a large-scale coherence in the geometry of the underlying mesh, and disappear as the mesh becomes more unstructured as we will demonstrate later in this section.

In Fig. 13, we plot the ionization fraction ($\tilde{n}_{\text{H I}}, \tilde{n}_{\text{H II}}$), density (ρ), pressure (P), temperature (T), and Mach number ($\mathcal{M} = v/c_s$) as a function of radius (r) at $t = 10$ Myr (blue curves), $t = 200$ Myr (green curves), and $t = 500$ Myr (red curves). By 10 Myr, the temperature increases behind the I-front due to photoheating which in turn increases the pressure. The density has not changed by much because the gas has not had time to react to these changes. This is a classic R-type front which moves supersonically to about the Strömgren radius (~ 5.4 kpc), within a single recombination time-scale ($t_{\text{rec}} \simeq 125$ Myr). By 200 Myr, the expansion of the I-front has progressed beyond what is expected from a pure R-type expansion. The I-front is now D-type meaning that the front moves along with the gas. The density behind the front is reduced, as the gas reacts to the pressure jump inside the front causing it to flow radially outwards. The gas piles up at the position of the I-front inducing a small density peak. It is important to note, however, that there is




Figure 12. Expansion of an H II region: maps showing slices of the domain at $z = 15$ kpc for the simulation of the expansion of an H II region in a constant density medium. The H I fraction (first column), density (second column), temperature (third column), pressure (fourth column), and Mach number (fifth column) are plotted at 10 Myr (top panels), 200 Myr (middle panels), and 500 Myr (bottom panels).

a second density peak beyond the I-front, which is reproduced in the temperature, pressure, and Mach number profiles. This is the temporary effect of photoheating by high-energy photons. The front then expands slowly outwards till a pressure equilibrium is reached. The final radius of the H II region (r_f) is given by,

$$r_f \simeq \left(\frac{2T}{T_e} \right)^{2/3} r_s \sim 220 \text{ kpc}, \quad (85)$$

where T is the temperature inside the H II region, T_e is the background temperature, and r_s is the Strömgren radius (equation 75). This indicates that the domain needs to be much larger in order to simulate the equilibrium state of the solution. Therefore, we stop the simulation at 500 Myr, a time at which the I-front is still within the domain and plot the profiles as shown. We note that the positions of the fronts and profiles match previous results in Iliev et al. (2009) and R13 very well.

Fig. 14 shows the position and velocity of the I-front (defined as where the radial average of $\tilde{n}_{\text{H II}}$ is equal to 0.5). For comparison, we also overplot the position of the front obtained using the RAMSES-RT

code presented in R13. The curves for the two codes are virtually identical, with a very slight difference in the speed at late times, which can be attributed to improper boundary conditions as the size of the I-front becomes comparable to the box size.

We now turn our attention to an I-front expansion created by a point source at the centre of a spherically symmetric, steeply decreasing power-law density profile with a small flat central core of density n_0 and radius r_0 :

$$n_{\text{H}}(r) = \begin{cases} n_0 & \text{if } r < r_0 \\ n_0(r_0/r)^2 & \text{if } r \geq r_0. \end{cases} \quad (86)$$

Within the central core, the I-front propagates similarly as in the previous simulation. The propagation of an I-front in r^{-2} density profiles with full gas dynamics does not have an exact analytical solution, but some insights can be gained by comparing the core radius (r_0) to the Strömgren radius (r_s). For $r_s < r_0$, the I-front stalls within the core, converts to D-type, but starts to re-accelerate upon entering the steep density gradient. Alternatively, when $r_s \geq r_0$, the source flash ionizes the cloud on time-scales shorter than the dynamical time of the gas. We simulate the first case, which is more




Figure 13. Expansion of an H II region: the ionization (top left panel), density (top middle panel), pressure (top right panel), temperature (bottom left panel), and Mach number (bottom middle panel) profiles at 10 Myr (blue curves), 200 Myr (green curves), and 500 Myr (red curves) in the simulation of the expansion of an H II region in a constant density medium. The profiles generally match the results from previous simulations of the same test reported in Iliev et al. (2009, test 5).




Figure 14. Expansion of an H II region: the radius (top panel) and velocity (bottom panel) of the ionization front as a function of time in the expansion of an H II region in a constant density medium. The evolution of the ionization front matches well with results obtained for the same test with RAMSES-RT (dashed black curve) as reported in R13.

interesting and is a better test of the coupling between the radiation field and the hydrodynamics.

For this test, the domain of side length $2L_{\text{box}} = 1.6 \text{ kpc}$ is resolved with 2×80^3 resolution elements placed on a regular staggered grid. As before, the mesh is allowed to move and distort

according to local fluid flow. The central core has a density of $n_0 = 3.2 \text{ cm}^{-3}$ and a radius of $r_0 = 91.5 \text{ pc}$. The central source is a blackbody spectrum with $T_{\text{eff}} = 10^5 \text{ K}$ and emits at a rate of $10^{50} \text{ photons s}^{-1}$. The initial temperature of the gas is 100 K. For these parameters, $r_s \simeq 70 \text{ pc}$, meaning that the I-front changes from R-type to D-type within the core. The simulation is run for 25 Myr, which is much larger than the recombination time-scale within the core, which is about $t_{\text{rec}} \sim 0.04 \text{ Myr}$.

Fig. 15 shows the cross-section (at $z = 0.8 \text{ kpc}$) maps of the H I fraction (first column), density (second column), temperature (third column), pressure (fourth column), and the Mach number (fifth column) at 3 Myr (top panels), 10 Myr (middle panels), and 25 Myr (bottom panels). The photoheated gas pressure is able to evacuate most of the gas from the central regions, completely changing the background density profile. We see that the I-front is more spherical than in the previous simulation because the gas velocities are higher which distorts the mesh and reduces the large-scale coherence in its geometry.

Fig. 16 plots the ionization fraction ($\tilde{n}_{\text{H I}}$, $\tilde{n}_{\text{H II}}$), density (ρ), pressure (P), temperature (T), and Mach number ($\mathcal{M} = v/c_s$) as a function of radius (r) at $t = 3 \text{ Myr}$ (blue curves), $t = 10 \text{ Myr}$ (green curves), and $t = 25 \text{ Myr}$ (red curves). By 3 Myr, the I-front has moved out of the central core and is expanding rapidly as a D-type front due to the steep density gradient outside the core. We obtain very sharp I-fronts and the profiles of the various hydrodynamic quantities match quite well with same test performed with an array of numerical schemes presented in Iliev et al. (2009).

Finally, Fig. 17 shows the position and velocity of the I-front (defined as where the radial average of $\tilde{n}_{\text{H II}}$ is equal to 0.5). For comparison, we also overplot the position of the front obtained using the RAMSES-RT code presented in R13.

These tests verify the accuracy and reliability of our RT scheme on randomly oriented, moving meshes. They also validate the accuracy




Figure 15. Expansion of an H II region: maps showing slices of the domain at $z = 0.8$ kpc for the simulation of the expansion of an H II in an r^{-2} initial density profile. The H I fraction (first column), density (second column), temperature (third column), pressure (fourth column), and Mach number (fifth column) are plotted at 3 Myr (top panels), 10 Myr (middle panels), and 25 Myr (bottom panels).

of the coupling between the hydrodynamics and the radiation field in our implementation.

4.6 Radiation pressure-driven outflows

Next we assess the accuracy of momentum injection into the gas due to photon absorption, i.e. radiation pressure (equation 62). During photon–matter interactions, both energy and momentum are conserved. Photoheating is the product of thermalization of the leftover energy above the ionization threshold of the ionic species. The photon’s momentum is also transferred to the ion, which receives a kick ($\Delta v = E/c$) in the direction of the absorbed photon. An estimate of the change in momentum of an optically thick shell of gas under spherically symmetric geometry can be written as $\dot{P} = L/c$, where L is the luminosity of the source. If L is invariant then the change in velocity of the shell in time Δt is

$$\Delta v_{\text{shell}} = \frac{L \Delta t}{m_{\text{shell}} c}, \quad (87)$$

where m_{shell} is the mass of the shell. In case of a shell expanding into a uniform background of density ρ_0 , the velocity of the shell

v_{shell} at any time t is given by (Wise et al. 2012)

$$v_{\text{shell}} = t A (r_i^4 + 2 A t^2)^{-3/4}, \quad (88)$$

where $A = 3L/4\pi\rho_0 c$ and r_i is the starting position of the shell. The shell will first form once ionization balances recombinations at the Strömgren radius, therefore $r_i = r_s$.

Here, we follow the test setup presented in Sales et al. (2014). Specifically, we look at the outflow velocities generated when a constant monochromatic ($E = 13.6 \text{ eV}$) source is placed in the centre of a uniform medium. The luminosity of the source is $L = 10^6 L_\odot$ which translates to a photon injection rate of $\dot{N}_\gamma = 1.8 \times 10^{50} \text{ photons s}^{-1}$. The initial density and temperature of the gas are $n_{\text{H}} = 1 \text{ cm}^{-3}$ and $T = 100 \text{ K}$ respectively, with the gas composed of only neutral H atoms. A monochromatic source with the energy equal to the ionization potential of the H atom implies that there is no photoheating and therefore the internal energy per unit mass of the cells, u , experiences no change due to the presence of a luminous source. This translates into an approximately constant temperature in the cells, except the change in temperature due to the change in the mean molecular weight

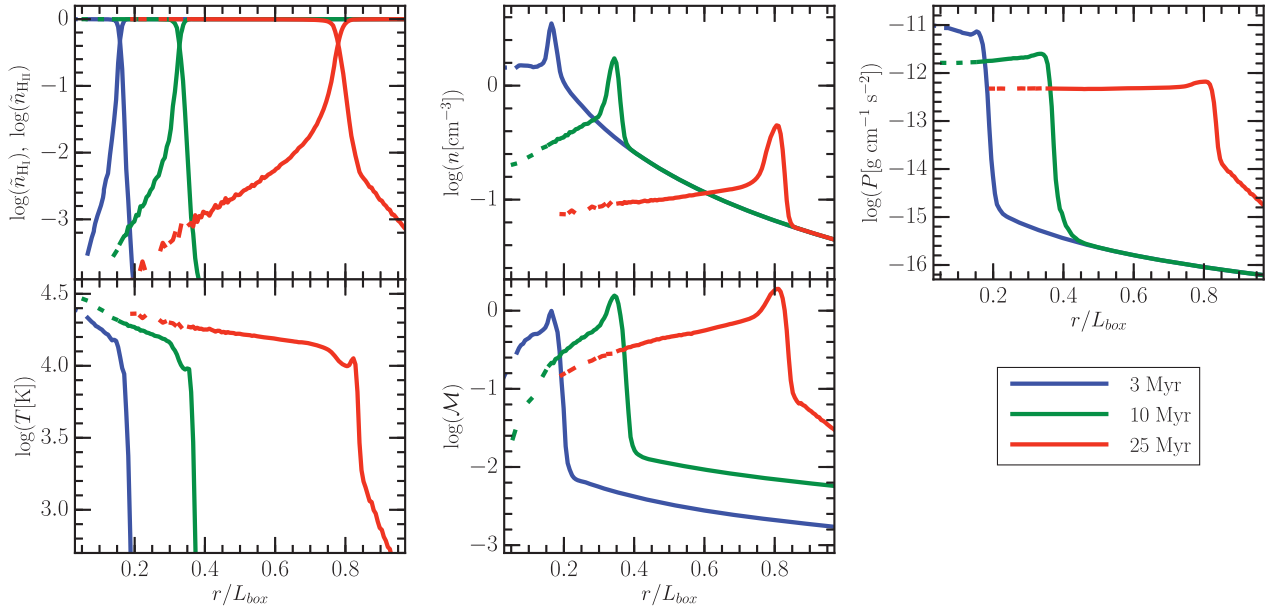


Figure 16. Expansion of an H II region: the ionization (top left panel), density (top middle panel), pressure (top right panel), temperature (bottom left panel), and Mach number (bottom middle panel) profiles at 3 Myr (blue curves), 10 Myr (green curves), and 25 Myr (red curves) in the simulation of the expansion of an H II region in an r^{-2} density profile. The profiles generally match the results from previous simulations of the same test reported in Iliev et al. (2009, test 6).




Figure 17. Expansion of an H II region: the radius (top panel) and velocity (bottom panel) of the ionization front as a function of time in the of the expansion of an H II region in a r^{-2} density profile. The evolution of the ionization front matches well with results obtained for the same test with RAMSES-RT (dashed black curve) as reported in R13.

of the gas due to ionization. For these parameters, $r_s = 51.7$ pc and the corresponding recombination time is $t_{\text{rec}} = 2.43 \times 10^3$ yr. The simulation domain is $L_{\text{box}} = 200$ pc on a side, initialized with 2×80^3 resolution elements placed in a regular staggered grid. As with previous RHD tests, the mesh is allowed to move and distort according to local fluid flow. The simulation is run for 5 Myr.




Figure 18. Radiation pressure-driven outflows: velocity as an function of time for gas with an initial density $n_{\text{H}} = 1 \text{ cm}^{-3}$ at a fixed temperature of $T = 100$ K. The simulation (red curve) agrees well with analytical estimations from equation (88, dashed black curve).

Fig. 18 shows the velocity (red curve) of the ionized shell defined as mass-weighted mean velocity calculated within a radius at which almost all the gas is fully ionized. For comparison, we also plot (black dashed curve) the expected analytic result from Wise et al. (2012, equation 88), which is in very good agreement with the simulated gas velocity. Initially, the I-front has a larger velocity compared to the analytic result as the light travels to the Strömgen radius before the gas has time to react to the changes. At later times ($t > 1$ Myr), mass entrainment will slow down the gas, an

effect which can be seen in the Wise et al. (2012) formulation as well. However, the simulation curves turn over at larger velocities compared to the analytic estimate which is also seen in Sales et al. (2014). This test confirms the accuracy of our radiation pressure implementation.

4.7 Free-streaming radiation from a thin disc

We have demonstrated in previous tests that the M1 closure relation performs remarkably well for a wide range of problems. However, the previous tests only simulate the propagation of radiation from a single source or a plane-parallel wave of radiation. A well-known disadvantage of the M1 closure is the inability to accurately model convergent rays (González et al. 2007; R13). The closure relation tends to produce spurious perpendicular flux when two rays are converging to a point, instead of passing through each other. Therefore, it is important to understand the limitations of this approximation and to test it under realistic conditions.

We therefore perform the test described in R15 to check the validity of our scheme in a multiple source geometry. Specifically, we test the distribution of the radiation field intensity, when a luminous thin disc is surrounded by an optically thick torus. The rest of the domain is assumed to contain a tenuous optically thin gas that allows for free transport of radiation. The 2D box of size $L_{\text{box}} = 1$ on a side is resolved by 128^2 resolution elements. A thin luminous disc of height $1/128$ (corresponding to one cell width) and width $L = 0.125$ is placed parallel to the x -axis and is centred at $(x, y) = (0.5, 0.1)$. Surrounding this disc is a one cell high torus that is optically thick and for the purposes of this simulation acts as a radiation sink. The disc has a constant photon energy density (E_0), which is imposed at every time-step. The disc and the torus are made up of regular Cartesian cells, while the rest of the domain is made up of an irregular mesh obtained by randomly deviating the cell centres of a Cartesian mesh by a $0.2\Delta x$, where Δx is the cell width, mimicking a typical deviation between mesh-generating points and cell centres in real problems (Vogelsberger et al. 2012).

The field morphology can be obtained analytically for this set-up as shown in R15

$$E(x, y) = \frac{E_0}{2\pi} \left[\arctan\left(\frac{L/2 - x}{y}\right) + \arctan\left(\frac{L/2 + x}{y}\right) \right], \quad (89)$$

in a coordinate system, whose centre is defined at the centre of the disc, i.e. at $(x, y) = (0.5, 0.1)$ cm.

In Fig. 19, we present the histogram of the radiation field intensity in the box, with the dashed contour plotting the simulation results and the solid lines representing the analytic solution. The small irregularity in the contours arises because of the irregularity of the underlying mesh. We see that the simulation matches the analytic result at least qualitatively. However, there are a few important differences which highlight the shortcomings of the M1 closure approximation. Namely, the radiation field intensity contours overshoot the analytic solution in the y -direction by a small amount and undershoot the solution in the x -direction. This is because the photons from the right-hand side of the disc are not able to propagate to the left-hand side of the domain, because they interact with the photons that emanate from the left-hand side of the disc going in the opposite direction. This causes spurious perpendicular fluxes, thus overshooting the solution in the y -direction and undershooting the solution in the x -direction. We note that R15 finds similar results with their M1 closure scheme and that we have a qualitative agreement with the analytic estimate for the radiation field morphology.




Figure 19. Free-streaming radiation from a thin disc: the radiation field morphology around a thin disc surrounded by an optically thick torus. The solid lines depict the analytic radiation field contours that are expected from this set-up, while the dashed lines plot the simulation results. The simulation slightly overshoots the analytic solution in the y -direction and undershoots it in the x -direction.

4.8 Dust absorption in an optically thick medium

In this section, we test the coupling between dust, gas, and the IR radiation field. Specifically, we examine how well the IR dust–gas coupling performs in the case of absorption in an optically thick regime. It is well known that in highly optically thick media, the photons propagate in a random walk reducing the radiation transport equations to an isotropic diffusion equation.

Previous works (Liu 1987; Bouchut 2004) have shown that if the numerical diffusion of the scheme becomes larger than the true radiation diffusion then the operator split approach to solve the RT equations is not valid anymore as the source terms become extremely stiff compared to the hyperbolic transport terms. Therefore, the ability of a numerical scheme to model radiation transport in an optically thick regime is very sensitive to the inherent numerical diffusivity of the scheme. There are a couple of ways to overcome this problem. Berthon, Charrier & Dubroca (2007) proposed to modify the Riemann solution such that it explicitly takes care of the source terms. This solution to the Riemann problem becomes much more complicated but it does recover the right asymptotic limit in the optically thick regime. However, it is unclear how radiation pressure can be accounted for in such a scheme.

R15 on the other hand propose an alternative method based on the isotropic diffusion source approximation methodology (Liebendörfer, Whitehouse & Fischer 2009). The IR photon group is split into two components, a trapped component (E_t) and a free streaming component (E_s). The trapped radiation energy is assumed to be strictly isotropic in angular space, corresponding to the asymptotic limit of vanishingly small mean-free path. The amount of trapped photons within a cell can be obtained by comparing the numerical diffusivity of the scheme with the expected analytic diffusivity in the diffusion limit. Specifically, the free streaming flux




Figure 20. Dust absorption in an optically thick medium: the maps showing the stationary results from 2D runs with a constant flux of photons into the box from the left. The colour represents the radiation temperature, T_r , as indicated by the colour bar and contours indicated by the solid black lines. The left-hand panel shows the result using a PC-GLF scheme, which matches very well with the solution obtained in the 32^2 simulation of R15 (dashed green curves). The plot on the right-hand side shows the results from the same simulation but now run with our fiducial scheme. The radiation temperatures are about two times higher and is able to match the 1024^2 simulation of R15.

in the diffusion limit is

$$\mathbf{F}_s \simeq \frac{\tilde{c}}{3\kappa_R \rho} \nabla E_t, \quad (90)$$

and the numerical diffusion for a PC-GLF scheme is

$$\mathbf{F}_s \simeq \frac{\tilde{c} \Delta x}{3} \nabla E_s. \quad (91)$$

Equating these two equations gives $E_t = 3\tau_c E_s/2$, where τ_c is the optical depth of the cell. The trapped photons are advected along with the gas and the radiation pressure from trapped photons is accounted for by adding an additional non-thermal pressure component from the radiation field, $P = P_{\text{therm}} + P_{\text{rad}}$, where $P_{\text{rad}} = \tilde{c} E_t/(3c)$, to the momentum conservation equation (11). This method produces accurate results in the diffusion limit.

There are, however, a few drawbacks with this scheme. First, we note that equation (91) is only valid if the Riemann problem at the interface is solved using a GLF flux function and more importantly, the left- and right-state inputs to the Riemann solver must be the cell centred values and no longer works if the left and right states of the interface are the gradient extrapolated values. This forces the underlying numerical scheme to follow a PC approach, which, as seen before, is very diffusive and has suboptimal convergence properties (see Section 4.1 for more details). Secondly, even at relatively low optical depths (~ 1), 60 per cent of the total radiation flux is deposited into the trapped component. This component is isotropic in angular space and hence the diffusion is isotropic. Therefore, the M1 scheme reverts back to a flux-limited diffusion (FLD; Lucy 1977; Krumholz et al. 2007) scheme, thereby, erasing the directionality of the initial photon field. By the same rationale, the radiation pressure will be isotropic and will generate isotropic

velocities even when the underlying radiation field has an inherent directionality.

For these reasons, we chose not to implement sub-grid schemes to model the diffusion limit. Our fiducial scheme is extremely accurate and has very low diffusivity. Moreover, the convergence order of our scheme is ~ 2.0 . This implies that a small improvement in the resolution will decidedly improve the accuracy of the solution. Since, AREPO is a moving mesh code, which automatically refines the high-density regions (and correspondingly high opacity regions), we are able to obtain accurate solutions in many realistic problems as we will show in the present and forthcoming tests.

To elucidate the points made in this section, we perform a quantitative test, proposed in R15, as a simple demonstration of the accuracy of our scheme. We initialize a 2D domain of size 1 pc on a side. The opacity is $\kappa_R \rho = 6.48 \times 10^{-17} \text{ cm}^{-1}$, which sets $\tau_{\text{box}} = 200$. The underlying mesh is a regular Cartesian grid with 32 resolution elements on a side, translating to an optical depth in the cell of $\tau_c = 6.25$. We chose a Cartesian grid so as to compare our scheme to that of R15. The left boundary emits a constant IR flux of $5.44 \times 10^4 \text{ erg s}^{-1} \text{ cm}^{-1}$. The rest of the boundaries act as sinks for radiation. The hydrodynamics is turned off and the only source of cooling and heating of gas is through the gas–dust IR coupling. For this test, we use the full speed of light. The initial temperature of the gas is 10 K. The simulation is run till the result has converged to a steady-state solution.

We perform two different simulations, one with a PC approximation and a GLF flux (PC-GLF) function and another with our fiducial scheme. The PC-GLF run mimics the scheme outlined in R15. The left-hand panel of Fig. 20 shows the temperature map in the PC-GLF run. The solid black lines show the contours of the

temperature map for the simulation and the dashed green lines are the contours from the same test performed by R15. The PC–GLF run matches quite nicely with the 32^2 run of R15, confirming that our PC approximation performs equally well with their scheme. The maximum temperature reached is about ~ 260 K.

As the optical depth of the cell is quite high, the numerical diffusion of the PC–GLF scheme exceeds that of the analytic value and hence the IR heating of the gas is massively underestimated. We rerun the same low-resolution 32^2 simulation with our fiducial scheme. The right-hand panel shows the temperature map in our fiducial run with black solid contours indicating the temperature obtained in our simulation and the dashed green line now instead shows the results from R15’s high-resolution 1024^2 simulation. We see that our low-resolution 32^2 simulation using the fiducial scheme accurately reproduces R15’s 1024^2 simulation. This is because our scheme is able to reduce the numerical diffusivity by a large amount, thereby accurately capturing the diffusion limit without the need for sub-grid diffusion models.

4.9 Diffusion of constant luminosity source

In this subsection, we redo the test proposed in R15 (Section 3.6), with the aim to quantify the numerical diffusivity of our scheme in the radiation diffusion regime. A 3D domain of $L_{\text{box}} = 500$ pc on a side is initialized with 32^3 resolution elements arranged in a Cartesian mesh. A source with a constant luminosity of $L = 10^{50}$ photons s^{-1} is placed in the centre of the domain. The gas is assumed to have an opacity of $\kappa_{\text{R}} = 10 \text{ cm}^2 \text{ g}^{-1}$. The hydrodynamics is turned off, but the radiation is allowed to propagate radially outward and the simulation is stopped when a steady-state solution is reached. The density of the gas is varied from $n_{\text{H}} = 5 - 10^4 \text{ cm}^{-3}$ corresponding to cell optical depths of $\tau_{\text{c}} = 0.004 - 8$.

The diffusion equation in a uniform optically thick medium is given as

$$\frac{\partial N}{\partial t} - \frac{\tilde{c}}{3\kappa_{\text{R}}\rho} \nabla^2 N + \mathcal{L} = 0, \quad (92)$$

where ρ is the density of the gas, \mathcal{L} is the luminosity per unit volume, and N is the number density of photons. The steady-state solution of this equation is then

$$\tilde{c}N(r) = \frac{3\rho\kappa_{\text{R}}L}{4\pi r}, \quad (93)$$

where r is the distance from the source. So, ideally, we should expect the photon number density to diminish as $1/r$.

The analytic expectation is derived assuming an infinite homogeneous medium and the diffusion equation only formally achieves a steady-state solution at $t = \infty$. Both these approximations are broken in our simulation set-up. In order, to approximate the infinite spatial medium, the values of the radiation variables are set such that they roughly match the expected slope of equation (93), i.e.

$$\mathcal{U}_{\text{s}} = \mathcal{U}_{\text{f}} \left(1 - \frac{\Delta x}{L_{\text{box}}} \right), \quad (94)$$

where $\mathcal{U} = (N, \mathbf{F})$, Δx is the distance between the cell centres at the boundary and the subscripts s and f refer to the solid layer cell and the fluid layer cell at the boundary, respectively. The steady-state solution is assumed to have been reached once the change in \mathcal{U} is less than 1 per cent throughout the domain.

Fig. 21 shows the simulated (solid curves) and analytic (dashed curves) photon number density profiles for runs with gas densities,




Figure 21. Diffusion of constant luminosity source: the plots show time-converged radiation profiles from the source at the centre of the box in radiation tests with optical depths of 0.004 (red curves), 0.4 (green curves), 4 (blue curves), and 8 (cyan curves). The simulation results (solid curves) match the analytic solution (dashed curves) up to a cell optical depth of about $\tau_{\text{c}} \sim 4$ above which the numerical diffusion starts to dominate.

$n_{\text{H}} = 5 \text{ cm}^{-3}$ (red curves), $n_{\text{H}} = 5 \times 10^2 \text{ cm}^{-3}$ (green curves), $n_{\text{H}} = 5 \times 10^3 \text{ cm}^{-3}$ (blue curves), and $n_{\text{H}} = 10^4 \text{ cm}^{-3}$ (cyan curves). The simulation with the lowest density (and hence the lowest optical depth) gas is optically thin to the radiation and hence the photons stream out of the box without any hindrance. This leads to a $N \propto 1/r^2$ dependence which is reproduced by the simulation as expected.

As the density increases the gas gradually becomes optically thick and the radiation field starts to diffuse through the medium rather than streaming out and the photon density profiles gradually starts to follow the $1/r$ relation given by equation (93). In the simulations with $n_{\text{H}} = 5 \times 10^2 \text{ cm}^{-3}$, the cells have optical depths of $\tau_{\text{c}} = 0.4$, which can just be resolved by the R15 scheme. We do recover the results with great accuracy in our scheme as well. However, this is the limit to which the R15 scheme can resolve the photon diffusion and it has troubles capturing the correct solution for higher optical depths as evidenced by very little change in the simulation result for higher densities (left two panels of Fig. 8 in R15). This is because, beyond this opacity the numerical diffusion of their scheme becomes larger than the expected analytic diffusion. This is corrected by using a sub-grid diffusion model which has its own drawbacks as discussed in the previous section (Section 4.8). However, in our simulation with a gas density of $n_{\text{H}} = 5 \times 10^2 \text{ cm}^{-3}$ having an optical depth of $\tau_{\text{c}} \sim 4$, our scheme is able to reproduce the analytic solution very well. This implies that the numerical diffusion of the fiducial scheme is so low that it captures the right solution even when the photon mean-free path is under resolved. This is also seen in the previous section, where the scheme was able to obtain the right results even when the cell optical depth was $\tau_{\text{c}} \sim 6.25$. There is a slight discrepancy between the analytic and simulation

results at large radii which is attributed to the fact that the boundary conditions are imperfect.

Obviously, the numerical diffusion of our scheme cannot be identically zero. The results from the highest density ($n_{\text{H}} = 5 \times 10^2 \text{ cm}^{-3}$) show the limitation of the scheme. The cell optical depth is about $\tau_c \sim 8$, and the simulation undershoots the photon number density at all radii in this regime, indicating that the numerical diffusion is larger at such high optical depths.

To conclude, our numerical algorithm performs well up to about a cell optical depth of $\tau_c \sim 6$. Above this threshold, the numerical diffusion dominates, degrading the accuracy of the solution. We note that this threshold is about an order of magnitude larger than the one for the scheme presented in R15. This allows us to capture the accurate solutions even at high optical depths without resorting to sub-grid diffusion models. Additionally, our scheme converges at a much higher rate ~ 2.0 and can therefore efficiently achieve higher accuracy by improving the spatial resolution.

4.10 Levitation of optically thick gas

Radiation pressure, both from direct UV and multiscattered IR, has been hypothesized to drive significant galactic-scale outflows ($\sim 100 \text{ km s}^{-1}$; Hopkins et al. 2011; Agertz et al. 2013; Hopkins et al. 2014). While the momentum injection rate of the single-scattered UV photons is just $\dot{P}_{\text{UV}} = L/c$, the reprocessed IR radiation field can be trapped in a high optical depth medium boosting the momentum injection, i.e. $\dot{P}_{\text{IR}} = \tau_{\text{IR}} L/c$. The efficiency with which the gas can trap the IR radiation field is unknown and different RT schemes seem to produce different results (Krumholz & Thompson 2012; Davis et al. 2014; Rosdahl & Teyssier 2015; Zhang & Davis 2017).

In this section, we perform the experiment first outlined in Krumholz & Thompson (2012), which tests the radiation pressure, radiation–temperature coupling, and the multiscattering of IR radiation. A thin layer of gas is placed in an external gravitational potential, and a certain amount of IR flux is injected in the direction opposite to gravity. This set-up mimics the physical conditions found in stellar nurseries or in the central plane of an optically thick galactic disc. It allows us to study how gravitationally bound gas responds to multiscattering IR radiation. Krumholz & Thompson (2012) argue that as the gas lifts, it becomes Rayleigh–Taylor unstable, which leads to a significant reduction in the coupling between the IR radiation and the gas. The RTI creates chimneys through which the radiation escapes rather than coherently lifting the gas. In their work, the radiation field has been modelled using the FLD approach, which assumes that the radiation flux always points in the direction of gradient of the photon energy density. As discussed in Section 2, the directionality of the underlying photon field is washed out by using this scheme. It is therefore possible that the radiation field diffuses out through the path of least resistance.

Davis et al. (2014) performed the same experiment using a more accurate VET closure scheme, which constructs the radiation flux vector for every volume element by sweeping the whole domain with short characteristics rays, effectively incorporating the contribution from all sources and sinks. Their results show that the VET closure relation coherently lifts the gas even in the presence of RTI. The velocity of the gas also significantly increases with the VET scheme. This is despite the fact that the average optical depths and the radiation force on the gas is quite similar to the FLD runs. The difference stems from the fact that VET scheme manages to maintain the average Eddington ratio to just above one, while the FLD only achieves Eddington ratios which are below unity.

This discrepancy between the different numerical schemes can be explained by their varying accuracy in estimating the direction of the underlying photon field. The diffusion approximation only transports photons in the direction of energy gradient, which is a very good approximation in highly thick media, but fails in optically thin or even slightly optically thick systems. R15 performed the same test but using the M1 closure relation and a PC–GLF scheme (called M1–R15 scheme from now on) and found results which were closer to the FLD results rather than the VET. They argue that although the M1 closure locally stores the bulk direction of the radiation field, its inability to accurately capture the propagation direction in the presence of multiple sources creates artificial diffusion. This causes the radiation to escape out of the chimneys created by the RTI. As we also employ the M1 closure relation, we expect to get similar results as R15. However, it is interesting to see if our low diffusion, higher order scheme will perform better. Moreover, the quantitative results using the FLD, VET, and M1 closures in terms of optical depths, Eddington ratios and gas velocities are close enough that this experiment can act as a good test of our implementation.

The simulation set-up consists of a 2D domain of boxsize $L(x, y) = (L_{\text{box}}/2, 2L_{\text{box}}) = (512, 2048)h_*$, where h_* is the characteristic scale height of the initial density profile. A layer of gas is placed at the bottom of the box and given an exponential density profile $\rho(y) = \rho_* \exp(-y/h_*)$, where $h_* = 2 \times 10^{15} \text{ cm}$ and $\rho_* = 7.1 \times 10^{-16} \text{ g cm}^{-3}$. The column density of the box is then $\Sigma_* = 1.4 \text{ g cm}^{-2}$. On top of this initial density profile a perturbation of the form

$$\frac{\partial \rho}{\rho} = 0.25(1 \pm \chi) \sin(2\pi x/L_{\text{box}}), \quad (95)$$

is added, where χ is random number uniformly distributed between $[-0.25, 0.25]$. This mimics the turbulent nature of the gas present in the star-forming birth cloud and in the ISM. The density profile has a floor at $10^{-10} \rho_*$ and the initial temperature of the gas is uniformly set to $T_* = 82 \text{ K}$. The gas is acted upon by a homogeneous gravitational acceleration pointing downwards with a magnitude of $g = 1.46 \times 10^{-6} \text{ cm s}^{-2}$. We note that under these conditions the system is unstable as gas pressure cannot counteract gravity.

The initial mesh consists of a high-resolution Cartesian mesh at the bottom in order to resolve the high-density gas, with the volume of the resolution element set to $(0.5h_*)^2$ (similar to the resolution used in Davis et al. 2014 and R15). The resolution is then degraded slowly till a minimum resolution with a volume of $(8h_*)^2$ is reached. The initial configuration of the mesh is not that important, as the mesh moves and distorts according to local fluid flow. The mesh is regularized where needed and refined and derefined such that the minimum $((0.5h_*)^2)$ and maximum $((8h_*)^2)$ cell sizes are approximately maintained throughout the simulation run.

The bottom boundary of the domain emits a constant IR radiation flux of $F_* = 2.54 \times 10^{13} \text{ L}_{\odot} \text{ kpc}^{-2} (1.03 \times 10^4 \text{ erg cm}^{-2} \text{ s}^{-1})$ and the box is initialized to contain an upwards radiation flux of the same magnitude, with $\bar{c}E = F_y = F_*$ and $F_x = 0$. This sets the radiation temperature to

$$T_{\text{r}*} = \left(\frac{F_*}{ca} \right)^{1/4} = T_*, \quad (96)$$

where a is the radiation constant. The radiation is coupled to the gas–dust fluid using the equations described in Section 3.2.2. The radiation energy density and the radiation flux are coupled to the gas using the *Planck* (κ_{p}) and *Rosseland* (κ_{r}) mean opacities, given

by

$$\kappa_p = 0.1 \left(\frac{T}{10 \text{ K}} \right)^2 \text{ cm}^2 \text{ g}^{-1}, \quad (97)$$

$$\kappa_R = 0.0316 \left(\frac{T}{10 \text{ K}} \right)^2 \text{ cm}^2 \text{ g}^{-1}, \quad (98)$$

which sets the initial value of the Rosseland mean opacity to $\kappa_{R*} = 2.13 \text{ cm}^2 \text{ g}^{-1}$. These opacity functions are the same as the ones used in previous works (Krumholz & Thompson 2012; Davis et al. 2014; Rosdahl & Teyssier 2015) and mimic the observed dust opacity functions in cold ($T \leq 150 \text{ K}$) gas obtained by Semenov et al. (2003). We note that the only non-adiabatic source of heating-cooling for the gas is through IR radiation-dust-gas coupling

$$\frac{\partial u}{\partial t} = -\frac{\partial E}{\partial t} = \kappa_p \rho (\tilde{c}E - c a T^4). \quad (99)$$

The gravitational force and the radiation pressure compete with each other and the eventual motion of the gas will be determined by the Eddington ratio

$$f_E = \frac{f_{y,\text{rad}}}{g\rho}, \quad (100)$$

where $f_{y,\text{rad}}$ is the radiation force in the vertical direction

$$f_{y,\text{rad}} = \frac{\kappa_R \rho F_y}{c}. \quad (101)$$

Therefore, the Eddington ratio at the start of the simulation is $f_{E*} = 0.5$, implying that the radiation is initially unable to overcome gravity and lift the gas. However, the gas is optically thick to the IR radiation with an initial optical depth of

$$\tau_* = \kappa_{R*} \Sigma_* = 3, \quad (102)$$

and the radiation can get trapped by the gas, increasing the Eddington ratio and driving the gas upwards.

The simulation domain is periodic in the x -direction for both gas and the radiation field. The top layer boundary cells have fixed values of temperature and density of $\rho = 10^{-13} \rho_*$ and $T = 10^{-3} T_*$ respectively. The velocity of the gas is set to zero at this boundary and the energy density and flux of the radiation field is also set to zero. This allows for free flow of gas and radiation field out of the top boundary. For the gas, the bottom boundary is reflective and allows for no escape or entry of gas. The bottom boundary should also emit radiation vertically at a rate of F_* . This is accomplished by setting the radiation energy density and flux at the bottom solid layer to E_s and $F_s = (0, F_*)$, where

$$\tilde{c}E_s = F_* - F_{y,f} + \tilde{c}E_f, \quad (103)$$

where the subscripts ‘s’ and ‘f’ refer to the solid layer and fluid layer cells respectively. We use our fiducial scheme to solve the transport equations. The simulation is run for $t = 150t_*$, where $t_* = h_*/c_*$ is the characteristic sound crossing time and $c_* = \sqrt{k_B T / (\mu m_p)} = 0.54 \text{ km s}^{-1}$ is the sound speed. The mean molecular weight is set to $\mu = 2.33$. We run the simulation with a reduced speed of light fraction set to $f_r = 0.01$.

Fig. 22 shows the maps of normalized density (top panels), radiation temperature (middle maps), and the radiation flux in the vertical direction (F_y ; bottom panels) at $t = 25t_*$ (first column), $t = 50t_*$ (second column), $t = 75t_*$ (third column), $t = 100t_*$ (fourth column), $t = 100t_*$ (fifth column), and $t = 150t_*$ (sixth column). Initially, the trapped radiation field greatly increases the radiation temperature (T_*), which in turn increases the gas temperature due

to the close coupling between the dust, gas, and the radiation field. The increased gas temperature leads to an increase in κ_R , resulting in increasing the radiation force and pushing $f_E > 1$. The gas becomes super-Eddington and it can be driven upwards in a thin shell. By $t = 25t_*$, the gas becomes Rayleigh-Taylor unstable, creating dense filamentary structures punctuated with diffuse chimneys through which the radiation escapes (as seen in the bottom left panel). However, the gas continues to lift and a significant fraction (~ 15 per cent) of the gas lifts beyond $1024h_*$ by $75t_*$. Some of the most dense filaments stall and fall back, but the radiation is still trapped within turbulent medium as evidenced by the temperature map at this time. However, as the chimneys widen the radiation eventually escapes and the gas falls back down to the bottom ($\lesssim 500h_*$) where it is kept turbulent by the competition between the radiation pressure and gravity.

Fig. 23 shows the mass-weighted mean vertical velocity (top panel), the mass-weighted vertical (middle panel), and horizontal (bottom panel) velocity dispersions as a function of the simulation time in units of the characteristic time-scale (t_*). The velocity and velocity dispersions are plotted in units of the characteristic sound speed c_* . The initial gas acceleration obtained using our simulation set-up (red curves) is much higher than either of the VET (blue curves), FLD (cyan curves), or the M1–R15 (green curves) schemes. The maximum average velocity reached by the gas in our simulation is about $\langle v_y \rangle_{\text{max}} \sim 15c_*$ which is about 50 per cent larger than the velocities achieved by the other schemes. This velocity is maintained for a brief amount of time ($\sim 25t_*$) and then the velocity starts to decrease quickly as the gas starts to fall back to the bottom of the domain. The chimneys created by the RTI also allow the radiation to escape, reducing the radiation pressure on the gas. This causes strong deceleration and the minimum velocity reaches $\langle v_y \rangle_{\text{min}} \sim -10c_*$. As the gas falls back the radiation pressure starts to build up again and eventually reaches a turbulent state, with velocity dispersions larger than the ones obtained by the M1–R15 and FLD schemes.

To understand this behaviour quantitatively, we plot the volume averaged Eddington ratio

$$f_{E,v} = \frac{\langle f_{y,\text{rad}} \rangle}{F_y}, \quad (104)$$

as a function of characteristic time in the top panel of Fig. 24. As mentioned earlier, this ratio reveals the competition between the radiation pressure and gravity with $f_{E,v} > 1$ corresponding to the case where the radiation pressure wins and drives outflows, while $f_{E,v} < 1$ indicates that gravity wins and the gas just falls back to the bottom of the box. The Eddington ratio quickly jumps to $f_{E,v} \sim 2$ from an initial value of 0.5 within the one t_* and maintains that value for about $15t_*$. This behaviour is different from the VET and M1–R15 schemes, where the gas motions cause their Eddington ratio to drop below 1.5 very quickly. The reason for this discrepancy can be understood by looking at the volume averaged optical depth from top to bottom (middle panel of Fig. 24),

$$\tau_v = L_{\text{box}} \langle \kappa_R \rho \rangle. \quad (105)$$

The optical depth rises from the initial value of $\tau_v = 3$ to ~ 12 within one t_* similar to the VET simulations. However, while the optical depth falls below ~ 10 quickly in their simulations, it is maintained for about $15t_*$ in our simulation. This increased optical depth results in larger values of the Eddington ratio, which in turn drives the gas to larger velocities. This indicates that our scheme is much more efficient in trapping the photons initially compared to the other schemes used in literature. The bottom panel of Fig. 24




Figure 22. Levitation of optically thick gas: maps showing the normalized density (top panels), radiation temperature (middle maps), and the radiation flux in the vertical direction (F_y ; bottom panels) at $t = 25t_*$ (first column), $t = 50t_*$ (second column), $t = 75t_*$ (third column), $t = 100t_*$ (fourth column), $t = 125t_*$ (fifth column), and $t = 150t_*$ (fourth column).




Figure 23. Levitation of optically thick gas: plot showing the mass-weighted y -velocity (top panel), the vertical velocity dispersion (middle panel), and the horizontal velocity dispersion (bottom panel) as a function of characteristic time t_* in our simulation (}; red curves) compared to M1–R15 (green curves), VET (blue curves), and FLD (cyan curves) schemes. Initially the mean y -velocity increases to about $15c_*$ due to the efficient trapping of photons. However, RTIs form low-density chimneys through which the radiation escapes and the gas falls back to the bottom of the box and settles into a turbulent state.

shows the ratio between the flux-weighted mean optical depth

$$\tau_F = L_{\text{box}} \frac{\langle \kappa_R \rho F_{ry} \rangle}{\langle F_{ry} \rangle}, \quad (106)$$

and τ_V . τ_F gives the momentum per unit area transferred from the radiation to the gas. Therefore, the ratio gives us an estimate of the fraction of actual momentum in the radiation field transferred to the gas. This quantity hovers around unity for the initial $15t_*$ implying that there is almost perfect coupling between radiation field and the gas at these early times. This is again better than the VET and M1–R15 schemes which only manage to couple about ~ 85 per cent of the radiation momentum into the gas. After about $15t_*$, all three quantities decline sharply, as the RTI instabilities build up. There is a slight rebound at about $25t_*$ but declines again quite quickly. All three quantities rebound slowly as the gas falls back to the bottom, the optical depth builds up and the gas reaches a turbulent state with the Eddington ratio hovering around unity.

This test has allowed us to gauge the advantages and limitations of our scheme. The general results agree well with the FLD, VET, and M1–R15 schemes used in literature. Our scheme performs better than the other schemes at early times by trapping the radiation more efficiently, increasing the optical depth and powering a higher velocity outflow. The maximum average velocity is about 50 per cent higher than the ones achieved by either of the VET or




Figure 24. Levitation of optically thick gas: plot showing the volume-weighted Eddington ratio, optical depth, and the ratio between the flux-weighted mean optical depth to the volume-weighted optical depth as a function of the characteristic time in our simulation (AREPO-RT; red curves) compared to M1–R15 (green curves), VET (blue curves), and FLD (cyan curves) schemes. All three quantities rise sharply in the beginning and these high values are maintained for about $15t_*$, after which all three quantities decline sharply as the RTIs build up. It is only after most of the gas falls back to the bottom ($\sim 50t_*$) that the optical depth starts to build up again and pushes the Eddington ratio closer to unity.

M1–R15 schemes. However, the late time behaviour of the gas is closer to the FLD scheme than the VET scheme. Most of the gas in our scheme falls back to the bottom of the domain (like FLD and M1–R15 schemes), while the VET scheme still continues to evacuate gas at a significant rate. The reason for this difference is the approximation made in the M1 closure relation. Although the M1 closure locally stores the bulk direction of the radiation field, it is unable to accurately capture the propagation direction in the presence of multiple sources creating artificial diffusion. This causes the radiation to escape out of the chimneys created by the RTI. Unfortunately, this is physical limitation of the of the M1 closure approximation and not a limitation of the implemented numerical scheme. It might turn out that the only way to accurately capture the exact coupling between radiation and gas in this regime might be to run quite expensive ray tracing short/long characteristic RT methods. However, we note that our scheme is better at trapping photons and driving outflows compared to the M1–R15 scheme which has the same physical limitations as our method.

5 SUMMARY AND CONCLUSIONS

In this paper, we have presented AREPO-RT, a novel implementation of an accurate and computationally efficient RHD scheme on

unstructured, moving Voronoi meshes. The scheme is based on a fluid description of the radiation field obtained by taking the zeroth- and first-order moments of the continuity equation of specific luminosity. These moment equations are a pair of hyperbolic conservation laws for photon energy density and photon flux. The system is then closed using the M1 closure relation that equates the pressure tensor to the energy density using a specific form of the Eddington tensor, which locally stores the bulk direction of the radiation field. The ability of the M1 closure to obtain an estimate of the Eddington tensor from just the local properties of the cell renders it very useful for computationally challenging problems.

We employ an operator split approach based on dividing the moment equations into equations for pure radiation transport and equations for the source and sink terms. We achieve high-order accuracy by replacing the PC approximation of Godunov’s scheme with a slope-limited piecewise linear spatial extrapolation and a first-order time prediction to obtain the states of the primitive variables on both sides of the cell interface. The spatial extrapolation is carried out using an LSF gradient estimator that has been shown to work well in meshes where the centre of mass of the cell can be offset from the mesh-generating point. Two different flux functions have been implemented to solve the Riemann problem at the interface: a second-order HLL flux function that uses the exact eigenvalues that represent the wave speeds of the RT transport equation and a GLF flux functions that sets the eigenvalues to the light speed irrespective of the geometry of the problem.

A conservative time-integration scheme is implemented using Heun’s method, which is a variant of the second-order Runge–Kutta scheme. The fluxes are computed as an average of fluxes at the beginning and end of the time-step. An RSLA is implemented in order to overcome the problem of small time-steps required due to the large speed of light. Additionally, a conservative subcycling scheme is implemented that is fully compatible with the individual time-stepping scheme of AREPO.

The radiation field couples to the gas and dust via photoionization, photoheating, and momentum injection. A multifrequency approach is used to model UV, optical, and IR radiation fields. We implement atomic H and He thermochemistry using a semi-implicit approach that is quite stable and allows for reasonably large time-steps. The local nature of the M1 closure relation allows the scheme to account for radiation emitted from collisional recombinations and discard OTSA. IR radiation is accounted for by coupling it to the semi-empirical dust model of McKinnon et al. (2016) which treats dust as a passive scalar, whose motion is closely coupled to the gas motions. This allows us to assume that the system is close to local thermodynamic equilibrium, which is a good approximation for cold high-density regions of the ISM. The main advantage of this coupling is that the dust opacities are self-consistently calculated from the properties of the dust in the cell, thereby eliminating the need for ad hoc scaling relations used in previous works (Bieri et al. 2017; Costa et al. 2017).

We test our implementation on a variety of problems. The implementation works well overall and reproduces analytic results in all the tests performed in this work. We first start with a test designed to gauge the accuracy of our radiation transport scheme in vacuum by simulating the radial advection of a thin Gaussian pulse. We find that using different flux functions (GLF or HLL) to solve the Riemann problem makes little difference to the obtained solution. It is much more important to use gradient extrapolated values at the face as inputs to the Riemann solver instead of the PC approximation used in R13 and R15. Such higher order fiducial schemes have very low numerical diffusion and the

convergence order is ~ 2.0 compared to ~ 0.5 for the PC schemes. In fact, the L^1 error in the simulation with 2×64^2 resolution elements run with the fiducial scheme is much lower than the error for a simulation with 2×256^2 resolution elements using a PC approximation.

The veracity of the multifrequency scheme coupled to the H–He photochemistry is verified by simulating Strömgren spheres around an ionizing source in a constant density medium. The ionization structure and the time evolution of the Strömgren radius match very well with the analytic expectations. We also simulated the case of an I-front trapping inside a dense clump of gas. The I-front is trapped inside the clump forming a shadow behind the clump. The sharpness of the shadow correlates with the diffusivity of the scheme. We show that the PC approximation is unable to produce accurate shadows if the radiation transport is even slightly non-parallel to the mesh geometry. On the other hand, our fiducial schemes which use gradient extrapolations are able to produce sharp shadows irrespective of the flux function or the mesh geometry used.

The coupling between RT and hydrodynamics is tested by simulating the expansion of an HII region in both a constant density medium and in a medium with a steep power-law density slope which mimics more realistic situations. The evolution of the ionization structure, density, temperature, pressure, and the Mach number of the gas matches very well with the results obtained by previous works. Additionally, we also address the accuracy of the momentum injection into the gas due to photon absorption by simulating radiation pressure-driven outflows. We show that after an initial period of supersonic expansion, a linear relation between the gas velocity and time is reproduced which is directly proportional to the luminosity of the source. These tests together validate the accuracy of the coupling between hydrodynamics and the radiation field.

One of the main shortcomings of the M1 closure relation is its inability to accurately determine the direction of the underlying photon field in a multiple source geometry. We quantify this deficiency by simulating the topology of the radiation field emanating from a thin disc surrounded by an optically thick torus. Although we achieve qualitative agreement with the analytic expectation, the simulation slightly overshoots the field geometry perpendicular to the disc and undershoots it in the transverse direction. This is because the rays from one side of the disc intersects the rays emanating from the opposite side causing spurious perpendicular flux. We stress that this is a fundamental limitation of the M1 closure approximation. In fact, R15 find similar results with their M1 closure scheme.

Next, we test the implementation of the multiscattering IR–dust gas coupling. Especially in optically thick regimes multiple scatterings lead to an isotropization of the radiation flux and the radiation tends to diffuse rather than advect through the medium. It is difficult to capture this transition because if the numerical diffusion of the scheme becomes larger than the true radiation diffusion then the operator split approach to solve the RT equations is not valid anymore. Therefore, the ability of a numerical scheme to model the radiation transport in an optically thick regime is extremely sensitive to the inherent numerical diffusivity of the scheme. Our simulations of dust coupling in a optically thick media show that our fiducial scheme manages to reproduce accurate results for $\tau_c \lesssim 6$, while that limit without using a sub-grid diffusion model for R15 is $\tau_c \sim 0.6$. We conclude that our scheme is able to attain the right solution even when the optical depth is moderately under-resolved.

As a final test of our scheme, we explore the ability of a trapped IR radiation field to accelerate a layer of gas in the presence of an external gravitational field that points in the opposite direction to the radiation pressure. Simulations using an FLD scheme argue that as the gas lifts it becomes Rayleigh–Taylor unstable creating chimneys through which the radiation escapes rather than coherently lifting the gas. Simulations performed with a more accurate VET closure scheme, however, coherently lift the gas even in the presence of RTI. The difference stems from the fact that VET manages to accurately estimate the direction of the underlying photon field, while the FLD approximation only transports photons in the direction of energy gradient. Our scheme performs better than the other schemes at early times by trapping the radiation more efficiently, increasing the optical depth and powering a higher velocity outflow. However, the late time behaviour of the gas is closer to the FLD scheme than the VET scheme. Most of the gas in our scheme falls back to the bottom of the domain, while the VET scheme still continues to evacuate gas at a significant rate. This is because, although the M1 closure locally stores the bulk direction of the radiation field, it is unable to accurately capture the propagation direction in the presence of multiple sources, creating artificial diffusion. Unfortunately, this is a physical limitation of the M1 closure approximation. It might turn out that the only way to accurately capture the exact coupling between radiation and gas in this regime might be to run quite expensive ray-tracing short/long characteristic RT methods. However, we note that our scheme is much better at trapping photons and driving outflows compared to the R15 scheme which have the same physical limitations as our scheme.

We conclude that we have implemented an efficient, robust and accurate RHD solver in the moving-mesh code AREPO. In forthcoming work, we plan to use this implementation to study timely problems in astrophysics related to RT, such as the role of radiative stellar feedback in driving galactic scale outflows, radiation pressure from quasars, and its role in quenching high-redshift galaxies and modelling the reionization history of the Universe.

ACKNOWLEDGEMENTS

We would like to thank the referee for giving constructive comments which substantially improved the quality of the paper. We thank Shane Davis and Joakim Rosdahl for kindly sharing their data. RK acknowledges support from NASA through Einstein Postdoctoral Fellowship grant number PF7-180163 awarded by the Chandra X-ray Center, which is operated by the Smithsonian Astrophysical Observatory for NASA under contract NAS8-03060. MV acknowledges support through an MIT RSC award, the Alfred P. Sloan Foundation, NASA ATP grant NNX17AG29G, and a Kavli Research Investment Fund. RM acknowledges support from the DOE CSGF under grant number DE-FG0297ER25308. VS acknowledges support through subproject EXAMAG of the Priority Program 1648 SPPEXA of the German Science Foundation. VS and RP are also supported by the European Research Council through ERC-StG grant EXAGAL-308037. The simulations were performed on the joint MIT-Harvard computing cluster supported by MKI and FAS. The figures in this work were produced by using the MATPLOTLIB graphics environment (Hunter 2007) and the interactive visualization tool SPLASH (Price 2007).

REFERENCES

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

- Abel T., Norman M. L., Madau P., 1999, *ApJ*, 523, 66
 Agertz O., Kravtsov A. V., Leitner S. N., Gnedin N. Y., 2013, *ApJ*, 770, 25
 Alpher R. A., Herman R. C., 1948, *Phys. Rev.*, 74, 1737
 Battisti A. J., Calzetti D., Chary R.-R., 2016, *ApJ*, 818, 13
 Bauer A., Springel V., Vogelsberger M., Genel S., Torrey P., Sijacki D., Nelson D., Hernquist L., 2015, *MNRAS*, 453, 3593
 Behrooz P. S., Wechsler R. H., Conroy C., 2013, *ApJ*, 770, 57
 Berthon C., Charrier P., Dubroca B., 2007, *J. Sci. Comput.*, 31, 347
 Bieri R., Dubois Y., Rosdahl J., Wagner A., Silk J., Mamon G. A., 2017, *MNRAS*, 464, 1854
 Booth C. M., Schaye J., 2009, *MNRAS*, 398, 53
 Bouchut F., 2004, *Nonlinear Stability of Finite Volume Methods for Hyperbolic Conservation Laws: And Well balanced Schemes for Sources*. Springer-Verlag, Berlin
 Bouwens R. J. et al., 2011, *ApJ*, 737, 90
 Bouwens R. J. et al., 2015, *ApJ*, 803, 34
 Cantalupo S., 2010, *MNRAS*, 403, L16
 Choi E., Ostriker J. P., Naab T., Johansson P. H., 2012, *ApJ*, 754, 125
 Ciardi B., Ferrara A., Marri S., Raimondo G., 2001, *MNRAS*, 324, 381
 Cicone C. et al., 2015, *A&A*, 574, A14
 Cielo S., Bieri R., Volonteri M., Wagner A., Dubois Y., 2017, *MNRAS*, 477, 1336
 Commerçon B., Debout V., Teyssier R., 2014, *A&A*, 563, A11
 Costa T., Sijacki D., Haehnelt M. G., 2014, *MNRAS*, 444, 2355
 Costa T., Rosdahl J., Sijacki D., Haehnelt M. G., 2017, *MNRAS*, 479, 2079
 Costa T., Rosdahl J., Sijacki D., Haehnelt M. G., 2018, *MNRAS*, 473, 4197
 Dalla Vecchia C., Schaye J., 2008, *MNRAS*, 387, 1431
 Davis S. W., Jiang Y.-F., Stone J. M., Murray N., 2014, *ApJ*, 796, 107
 Dijkstra M., Gronke M., Venkatesan A., 2016, *ApJ*, 828, 71
 Draine B. T., Hao L., 2002, *ApJ*, 569, 780
 Draine B. T., Lee H. M., 1984, *ApJ*, 285, 89
 Dubroca B., Feugeas J., 1999, *Acad. Sci. Paris C. R. Ser. Sci. Math.*, 329, 915
 Dullemond C. P., Juhasz A., Pohl A., Sereshti F., Shetty R., Peters T., Commerçon B., Flock M., 2012, *Astrophysics Source Code Library*, record ascl:1202.015
 Fabian A. C., 1999, *MNRAS*, 308, L39
 Fabian A. C., Iwasawa K., 1999, *MNRAS*, 303, L34
 Faucher-Giguère C.-A., Quataert E., 2012, *MNRAS*, 425, 605
 Finkelstein S. L. et al., 2015, *ApJ*, 810, 71
 Finlator K., Özel F., Davé R., 2009, *MNRAS*, 393, 1090
 Flock M., Fromang S., Turner N. J., Benisty M., 2016, *ApJ*, 827, 144
 Gardiner T. A., Stone J. M., 2005, *J. Comput. Phys.*, 205, 509
 Gardner J. P. et al., 2006, *Space Sci. Rev.*, 123, 485
 Gilmore G., Wyse R. F. G., Kuijken K., 1989, *ARA&A*, 27, 555
 Gnedin N. Y., 1998, *MNRAS*, 299, 392
 Gnedin N. Y., 2000, *ApJ*, 542, 535
 Gnedin N. Y., Abel T., 2001, *New Astron.*, 6, 437
 Gnedin N. Y., Hollon N., 2012, *ApJS*, 202, 13
 Gnedin N. Y., Hui L., 1998, *MNRAS*, 296, 44
 Gnedin N. Y., Ostriker J. P., 1997, *ApJ*, 486, 581
 Godunov S. K., 1959, *Math. Sb.*, 47, 271
 Goldsmith P. F., 2001, *ApJ*, 557, 736
 González M., Audit E., Huynh P., 2007, *A&A*, 464, 429
 González M., Vaytet N., Commerçon B., Masson J., 2015, *A&A*, 578, A12
 Greene J. E., Ho L. C., 2005, *ApJ*, 627, 721
 Greif T. H., 2014, *MNRAS*, 444, 1566
 Guedes J., Callegari S., Madau P., Mayer L., 2011, *ApJ*, 742, 76
 Haardt F., Madau P., 1996, *ApJ*, 461, 20
 Harten A., Lax P. D., van Leer B., 1983, *SIAM Rev.*, 25, 35
 Heckman T. M., Miley G. K., van Breugel W. J. M., Butcher H. R., 1981, *ApJ*, 247, 403
 Heng K., Frierson D. M. W., Philipps P. J., 2011, *MNRAS*, 418, 2669
 Hindmarsh A. C., Brown P. N., Grant K. E., Lee S. L., Serban R., Shumaker D. E., Woodward C. S., 2005, *ACM Trans. Math. Softw.*, 31, 363
 Hopkins P. F. et al., 2017, *MNRAS*, 480, 800
 Hopkins P. F., Quataert E., Murray N., 2011, *MNRAS*, 417, 950

- Hopkins P. F., Kereš D., Oñorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, *MNRAS*, 445, 581
- Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
- Iliev I. T. et al., 2006, *MNRAS*, 371, 1057
- Iliev I. T. et al., 2009, *MNRAS*, 400, 1283
- Ishibashi W., Fabian A. C., 2015, *MNRAS*, 451, 93
- Ishibashi W., Fabian A. C., 2017, *MNRAS*, 472, 2768
- Jaura O., Glover S. C. O., Klessen R. S., Paardekooper J.-P., 2017, *MNRAS*, 475, 2822
- Kannan R. et al., 2014a, *MNRAS*, 437, 2882
- Kannan R., Stinson G. S., Macciò A. V., Brook C., Weinmann S. M., Wadsley J., Couchman H. M. P., 2014b, *MNRAS*, 437, 3529
- Kannan R., Springel V., Pakmor R., Marinacci F., Vogelsberger M., 2016a, *MNRAS*, 458, 410
- Kannan R., Vogelsberger M., Stinson G. S., Hennawi J. F., Marinacci F., Springel V., Macciò A. V., 2016b, *MNRAS*, 458, 2516
- Kannan R., Vogelsberger M., Pfrommer C., Weinberger R., Springel V., Hernquist L., Puchwein E., Pakmor R., 2017, *ApJ*, 837, L18
- Karouzos M., Woo J.-H., Bae H.-J., 2016, *ApJ*, 819, 148
- Katz N., Weinberg D. H., Hernquist L., 1996, *ApJS*, 105, 19
- Kim J.-G., Kim W.-T., Ostriker E. C., Skinner M. A., 2017, *ApJ*, 851, 93
- Krumholz M. R., Thompson T. A., 2012, *ApJ*, 760, 155
- Krumholz M. R., Thompson T. A., 2013, *MNRAS*, 434, 2329
- Krumholz M. R., Klein R. I., McKee C. F., Bolstad J., 2007, *ApJ*, 667, 626
- Kuiper R., Klahr H., Beuther H., Henning T., 2012, *A&A*, 537, A122
- Laor A., Draine B. T., 1993, *ApJ*, 402, 441
- Levermore C. D., 1984, *J. Quant. Spec. Radiat. Transf.*, 31, 149
- Liebrandt M., Whitehouse S. C., Fischer T., 2009, *ApJ*, 698, 1174
- Liu T.-P., 1987, *Commun. Math. Phys.*, 108, 153
- Lucy L. B., 1977, *AJ*, 82, 1013
- Lupi A., Bovino S., Capelo P. R., Volonteri M., Silk J., 2017, *MNRAS*, 474, 2884
- Madau P., Haardt F., Rees M. J., 1999, *ApJ*, 514, 648
- Mas-Ribas L., Hennawi J. F., Dijkstra M., Davies F. B., Stern J., Rix H.-W., 2017, *ApJ*, 846, 11
- Mathis J. S., Rimpl W., Nordsieck K. H., 1977, *ApJ*, 217, 425
- McKinnon R., Torrey P., Vogelsberger M., 2016, *MNRAS*, 457, 3775
- McKinnon R., Torrey P., Vogelsberger M., Hayward C. C., Marinacci F., 2017, *MNRAS*, 468, 1505
- Mihalas D., Mihalas B. W., 1984, *Foundations of Radiation Hydrodynamics*. Oxford Univ. Press, New York, p. 731
- Moster B. P., Somerville R. S., Maulbetsch C., van den Bosch F. C., Macciò A. V., Naab T., Oser L., 2010, *ApJ*, 710, 903
- Moster B. P., Naab T., White S. D. M., 2013, *MNRAS*, 428, 3121
- Murray N., Quataert E., Thompson T. A., 2010, *ApJ*, 709, 191
- Navarro J. F., Frenk C. S., White S. D. M., 1996, *ApJ*, 462, 563
- Nelson C. H., Whittle M., 1995, *ApJS*, 99, 67
- Oppenheimer B. D., Davé R., 2006, *MNRAS*, 373, 1265
- Osterbrock D. E., Ferland G. J., 2006, *Astrophysics of Gaseous Nebulae and Active Galactic Nuclei*. University Science Books, Sausalito, CA
- Oxley S., Woolfson M. M., 2003, *MNRAS*, 343, 900
- Pakmor R., Springel V., Bauer A., Mocz P., Munoz D. J., Ohlmann S. T., Schaaf K., Zhu C., 2016, *MNRAS*, 455, 1134
- Parrish I. J., Stone J. M., 2005, *ApJ*, 633, 334
- Pawlik A. H., Schaye J., 2008, *MNRAS*, 389, 651
- Penzias A. A., Wilson R. W., 1965, *ApJ*, 142, 419
- Peters T. et al., 2017, *MNRAS*, 466, 3293
- Petkova M., Springel V., 2009, *MNRAS*, 396, 1383
- Petkova M., Springel V., 2011, *MNRAS*, 415, 3731
- Pillepich A. et al., 2018, *MNRAS*, 473, 4077
- Price D. J., 2007, *PASA*, 24, 159
- Reddy N. A., Steidel C. C., Pettini M., Bogosavljević M., Shapley A. E., 2016, *ApJ*, 828, 108
- Ripoll J.-F., Dubroca B., Duffa G., 2001, *Combust. Theor. Model.*, 5, 261
- Rosdahl J., Teyssier R., 2015, *MNRAS*, 449, 4380
- Rosdahl J., Blaizot J., Aubert D., Stranex T., Teyssier R., 2013, *MNRAS*, 436, 2188
- Rosdahl J., Schaye J., Teyssier R., Agertz O., 2015, *MNRAS*, 451, 34
- Rosen A. L., Krumholz M. R., McKee C. F., Klein R. I., 2016, *MNRAS*, 463, 2553
- Roth N., Kasen D., Hopkins P. F., Quataert E., 2012, *ApJ*, 759, 36
- Roškar R., Teyssier R., Agertz O., Wetzstein M., Moore B., 2014, *MNRAS*, 444, 2837
- Rusanov V. V., 1961, *J. Comput. Math. Phys. USSR*, 1, 267
- Sales L. V., Marinacci F., Springel V., Petkova M., 2014, *MNRAS*, 439, 2990
- Schaye J. et al., 2015, *MNRAS*, 446, 521
- Schaye J., Theuns T., Rauch M., Efstathiou G., Sargent W. L. W., 2000, *MNRAS*, 318, 817
- Semelin B., Combes F., Baek S., 2007, *A&A*, 474, 365
- Semenov D., Henning T., Helling C., Ilgner M., Sedlmayr E., 2003, *A&A*, 410, 611
- Shapiro P. R., Giroux M. L., 1987, *ApJ*, 321, L107
- Sijacki D., Springel V., 2006, *MNRAS*, 366, 397
- Silk J., Mamon G. A., 2012, *Res. Astron. Astrophys.*, 12, 917
- Skinner M. A., Ostriker E. C., 2013, *ApJS*, 206, 21
- Smoot G. F. et al., 1992, *ApJ*, 396, L1
- Spitzer L., 1978, *Physical Processes in the Interstellar Medium*. Wiley-Interscience, New York
- Springel V., 2010, *MNRAS*, 401, 791
- Springel V., Hernquist L., 2003, *MNRAS*, 339, 289
- Springel V., Di Matteo T., Hernquist L., 2005, *ApJ*, 620, L79
- Stinson G., Seth A., Katz N., Wadsley J., Governato F., Quinn T., 2006, *MNRAS*, 373, 1074
- Stinson G. S., Brook C., Macciò A. V., Wadsley J., Quinn T. R., Couchman H. M. P., 2013, *MNRAS*, 428, 129
- Strang G., 1968, *SIAM J. Numer. Anal.*, 5, 506
- Strömgren B., 1939, *ApJ*, 89, 526
- Thacker R. J., Couchman H. M. P., 2001, *ApJ*, 555, L17
- Thompson T. A., Quataert E., Zhang D., Weinberg D. H., 2016, *MNRAS*, 455, 1830
- Tielens A. G. G. M., 2005, *The Physics and Chemistry of the Interstellar Medium*. Cambridge Univ. Press, Cambridge, UK
- Tombesi F., Meléndez M., Veilleux S., Reeves J. N., González-Alfonso E., Reynolds C. S., 2015, *Nature*, 519, 436
- Trac H., Cen R., 2007, *ApJ*, 671, 1
- van Leer B., 1979, *J. Comput. Phys.*, 32, 101
- Vogelsberger M. et al., 2014a, *MNRAS*, 444, 1518
- Vogelsberger M. et al., 2014b, *Nature*, 509, 177
- Vogelsberger M., Sijacki D., Kereš D., Springel V., Hernquist L., 2012, *MNRAS*, 425, 3024
- Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, *MNRAS*, 436, 3031
- Walch S. K., Whitworth A. P., Bisbas T., Wünsch R., Hubber D., 2012, *MNRAS*, 427, 625
- Waxman E., Draine B. T., 2000, *ApJ*, 537, 796
- Weinberger R. et al., 2017, *MNRAS*, 465, 3291
- Whalen D., Norman M. L., 2006, *ApJS*, 162, 281
- White S. D. M., Rees M. J., 1978, *MNRAS*, 183, 341
- Wise J. H., Abel T., Turk M. J., Norman M. L., Smith B. D., 2012, *MNRAS*, 427, 311
- Zackrisson E., Inoue A. K., Jensen H., 2013, *ApJ*, 777, 39
- Zakamska N. L. et al., 2016, *MNRAS*, 459, 3144
- Zhang D., Davis S. W., 2017, *ApJ*, 839, 54
- Zubko V., Dwek E., Arendt R. G., 2004, *ApJS*, 152, 211

APPENDIX A: SUBCYCLING OF THE RT STEP

In this section, we describe a scheme designed to subcycle multiple RT steps for each hydro step that is compatible with the local time-stepping scheme of AREPO. Implementing such a scheme relies on how well the artificial domain boundaries created by inactive cells are handled. It was shown that the imposition of Dirichlet boundary conditions (Commerçon et al. 2014) can lead to a violation of energy conservation which can get quite acute in the presence of large




Figure A1. Subcycling of the RT step: the radius (top panel) and velocity (bottom panel) of the ionization front as a function of time for the expansion of an H II region in a r^{-2} density profile, in a simulation with $N_{\text{sub}} = 1$ (black curves), $N_{\text{sub}} = 2$ (red curves), $N_{\text{sub}} = 4$ (blue curves), and $N_{\text{sub}} = 8$ (green curves). The simulation results are converged with respect to the number of subcycles used.

energy gradients. We instead chose to follow the method outlined in Pakmor et al. (2016). The number of RT subcycles to every hydro step (N_{sub}) is an input variable. The hydro time-step of every cell in the simulation is then set as

$$\Delta t_{\text{hydro}} = \min(N_{\text{sub}} \Delta t_{\text{RT}}, \Delta t_{\text{hydro}}). \quad (\text{A1})$$

In order to maximize the impact of subcycling, the number of subcycles should be set to $N_{\text{sub}} = 2^n$, since the local time-stepping scheme of AREPO uses a discretization of the allowed time-step sizes into a power-of-two hierarchy (Springel 2010).

For every hydro step, the following RT loop is executed N_{sub} times:

- (i) The time-step of each subcycle step is set to $\Delta t = \Delta t_{\text{hydro}} / N_{\text{sub}}$.
- (ii) The thermochemistry and momentum injection steps are executed for all the active cells.
- (iii) A list of all active interfaces, i.e. interfaces with at least one adjacent active cell is made.
- (iv) All cells that share at least one corner with an active interface are collected. This includes a layer of inactive cells around the active cells.
- (v) The fluxes are exchanged over the active interfaces (as described in Section 3.1), with the time-step set to the minimum of the time-steps of the two adjacent cells divided by N_{sub} .
- (vi) The primitive variables of the active cells are then updated.

This form of subcycling is fully conservative and only requires to solve the RT equations on the active cells plus a one-cell boundary layer.

As a test of this scheme, we resimulate the expansion of an H II region in a varying density field as described in Section 4.5, with a varying number of RT subcycles. Briefly, a constant luminosity source with a blackbody spectrum ($T_{\text{eff}} = 10^5$ K) emitting at a




Figure A2. Subcycling of the RT step: the total run time (solid black line) of the simulation (in CPU hours) as a function of the number of RT subcycles used (N_{sub}) compared to the ideal $\propto 1/N_{\text{sub}}$ scaling (dashed line).

rate of 10^{50} photons s^{-1} is placed at the centre of a spherically symmetric, steeply decreasing power-law (-2) density profile with a small flat central core of gas. A domain of side length $2L_{\text{box}} = 1.6$ kpc is resolved by 2×80^3 resolution elements placed on a regular grid. The central core has a density of $n_0 = 3.2 \text{ cm}^{-3}$ and a radius of $r_0 = 91.5$ pc. The simulation is run for 25 Myr. For a fiducial run with $N_{\text{sub}} = 1$, the position and velocity of the ionization front are plotted in Fig. 17 and the corresponding quantities are shown in Fig. A1 for $N_{\text{sub}} = 1$ (black curves), $N_{\text{sub}} = 2$ (red curves), $N_{\text{sub}} = 4$ (blue curves), and $N_{\text{sub}} = 8$ (green curves). All the runs reproduce the same result irrespective of the number of subcycles used, proving the validity and accuracy of our scheme. We also note that during the simulation the time bin hierarchy reaches up to 4 bins deep, proving that the subcycling scheme is compatible with AREPO's local time-stepping scheme.

Fig. A2 shows the total amount of time taken to run the simulation (in CPU hours) as a function of N_{sub} (solid black curve). Satisfyingly, the run time of the simulation reduces by the expected factor of $\sim 1/N_{\text{sub}}$. This drop in the run time is attributed to the fact that the time-consuming routines are called less frequently due to the larger hydro step (equation A1). This can be seen more clearly in Fig. A3, which compares the amount of time taken by individual sub-routines, such as, RT (red shaded region; including cooling and chemistry), tree-based time-step calculation (cyan shaded region), domain decomposition (yellow shaded region), hydro (green shaded region), and Voronoi mesh construction (blue shaded region) for a simulation with $N_{\text{sub}} = 1$ (left-hand panel) and $N_{\text{sub}} = 8$ (right-hand panel). The time taken by the RT routine stays about constant because it still has to follow the time-step imposed by the speed of light. The increased hydro step reduces the frequency with which the other routines are called thereby massively reducing the time spent on them.

It is important to note that although Voronoi mesh construction is a computationally expensive process, it does not usually take up as much time as depicted in our simulation (~ 70 per cent). This only happens because we start out with a regular mesh and this introduces




Figure A3. Subcycling of the RT step: the total amount of computing time taken by RT (red region), tree-based time-step estimate (cyan region), domain decomposition (yellow region), hydrodynamic flux calculation (green region), and Voronoi mesh construction (blue region) routines in simulations with $N_{\text{sub}} = 1$ (left-hand panel) and $N_{\text{sub}} = 8$ (right-hand panel). Increasing the number of subcycles reduces the frequency with which the computationally expensive routines (such as the Voronoi mesh construction) are called, thereby reducing the total computing time.

a lot of degeneracies in mesh geometry which require a large number of geometric predicates that need to be carried out with exact floating point arithmetic to lift the degeneracies, thereby, artificially boosting the computing time. Starting from more realistic, and therefore less regular, particle geometries will drastically bring down the time taken by this routine. However, this high computing time for Voronoi mesh construction acts as a proxy for other routines not included in the simulation such as gravity and galaxy formation physics that will definitely be present in realistic simulations.

Finally, the speed of AREPO-RT, compared to the parent code AREPO is problem dependent. Specifically, it depends on the value of the reduced speed of light used in the simulations. If the time-step due to the speed of light is comparable to the time-step obtained from hydrodynamic considerations, then the overhead due to the RT flux calculations is pretty small, of order ~ 10 per cent. The non-equilibrium chemistry and cooling can cause significant (~ 30 per cent) overhead in simulations with high-density gas due to extremely small cooling times. The weak scaling of AREPO-RT is as good as AREPO (which has shown excellent weak scaling upto tens of thousands of cores) because it uses the same domain decomposition, parallelization, and communication algorithms.

APPENDIX B: H-HE THERMOCHEMISTRY

In this section, we describe in detail the semi-implicit scheme used to solve the thermochemical network of H and He (equations 47–54). First let us define

$$\begin{aligned} A &= \Delta t \sigma_{\text{eH I}} (n_{\text{e}})^n, \\ B &= \Delta t \tilde{c} \sum_i \tilde{\sigma}_{\text{iH I}} (N_{\gamma}^i)^n \text{ and} \\ C &= \Delta t \alpha_{\text{H II}} (n_{\text{e}})^n, \end{aligned} \quad (\text{B1})$$

where Δt is the time interval over which we are integrating the equation and the superscript ‘ n ’ denotes the values of the quantity

at the present time. Then, equation (49) then can be written in a semi-implicit form

$$\begin{aligned} (\tilde{n}_{\text{H II}})^{n+1} &= (\tilde{n}_{\text{H II}})^n + A(1 - (\tilde{n}_{\text{H II}})^{n+1}) + B(1 - (\tilde{n}_{\text{H II}})^{n+1}) \\ &\quad - C (\tilde{n}_{\text{H II}})^{n+1}, \end{aligned} \quad (\text{B2})$$

which gives

$$(\tilde{n}_{\text{H II}})^{n+1} = \frac{(\tilde{n}_{\text{H II}})^n + A + B}{1 + A + B + C} \quad (\text{B3})$$

where $\{\tilde{n}_{\text{H I}}, \tilde{n}_{\text{H II}}\} = \{n_{\text{H I}}, n_{\text{H II}}\}/n_{\text{H}}$, therefore $\tilde{n}_{\text{H I}} + \tilde{n}_{\text{H II}} = 1$ which in turn sets the value of $(\tilde{n}_{\text{H I}})^{n+1}$.

For He chemistry, we represent $\tilde{n}_j = n_j/n_{\text{He}}$ and define

$$\begin{aligned} D &= \Delta t \sigma_{\text{eHe II}} (n_{\text{e}})^n, \\ E &= \Delta t \alpha_{\text{He III}} (n_{\text{e}})^n, \\ F &= \Delta t \sigma_{\text{eHe I}} (n_{\text{e}})^n, \\ G &= \Delta t \alpha_{\text{He II}} (n_{\text{e}})^n, \\ H &= \Delta t \tilde{c} \sum_i \tilde{\sigma}_{\text{iHe I}} (N_{\gamma}^i)^n, \text{ and} \\ I &= \Delta t \tilde{c} \sum_i \tilde{\sigma}_{\text{iHe II}} (N_{\gamma}^i)^n. \end{aligned} \quad (\text{B4})$$

Therefore, the change in $n_{\text{He III}}$ (equation 51) can be written as

$$\begin{aligned} (\tilde{n}_{\text{He III}})^{n+1} &= (\tilde{n}_{\text{He III}})^n + D (\tilde{n}_{\text{He II}})^{n+1} - E (\tilde{n}_{\text{He III}})^{n+1} \\ &\quad + I (\tilde{n}_{\text{He II}})^{n+1}, \end{aligned} \quad (\text{B5})$$

which gives

$$(\tilde{n}_{\text{He III}})^{n+1} = \frac{(\tilde{n}_{\text{He III}})^n + (D + I) (\tilde{n}_{\text{He II}})^{n+1}}{1 + E}. \quad (\text{B6})$$

We note that $(\tilde{n}_{\text{He II}})^{n+1}$ is still an unknown, which is given by

$$\begin{aligned} (\tilde{n}_{\text{He II}})^{n+1} &= (\tilde{n}_{\text{He II}})^n - D (\tilde{n}_{\text{He II}})^{n+1} + E (\tilde{n}_{\text{He III}})^{n+1} \\ &\quad - G (\tilde{n}_{\text{He II}})^{n+1} - I (\tilde{n}_{\text{He II}})^{n+1} \\ &\quad + F \left(1 - (\tilde{n}_{\text{He II}})^{n+1} - (\tilde{n}_{\text{He III}})^{n+1} \right) \\ &\quad + H \left(1 - (\tilde{n}_{\text{He II}})^{n+1} - (\tilde{n}_{\text{He III}})^{n+1} \right). \end{aligned} \quad (\text{B7})$$

Substituting the value of $(\tilde{n}_{\text{He III}})^{n+1}$ from equation (B6) into equation (B7), we get

$$(\tilde{n}_{\text{He II}})^{n+1} = \frac{(\tilde{n}_{\text{He II}})^n + F + H - \frac{H + F - E}{1 + E} (\tilde{n}_{\text{He III}})^n}{1 + D + F + G + H + I + \frac{(H + F - E)(D + I)}{1 + E}}. \quad (\text{B8})$$

Then, the value of $(\tilde{n}_{\text{He II}})^{n+1}$ is used in equations (B6) to obtain $(\tilde{n}_{\text{He I}})^{n+1}$. Finally, $(\tilde{n}_{\text{He I}})^{n+1}$ is given as

$$(\tilde{n}_{\text{He I}})^{n+1} = 1 - (\tilde{n}_{\text{He II}})^{n+1} - (\tilde{n}_{\text{He III}})^{n+1}, \quad (\text{B9})$$

and the electron density is set to

$$(n_e)^{n+1} = (n_{\text{H II}})^{n+1} + (n_{\text{He II}})^{n+1} + 2 (n_{\text{He III}})^{n+1}. \quad (\text{B10})$$

Finally, the photon number density and flux are updated as

$$(N_\gamma^i)^{n+1} = \frac{(N_\gamma^i)^n + \Delta t \sum_j s_{ij}}{1 + \Delta t \tilde{c} \left(\sum_j (n_j)^n \tilde{\sigma}_{ij} + \kappa_i \rho \right)}, \quad (\text{B11})$$

and

$$(\mathbf{F}_\gamma^i)^{n+1} = \frac{(\mathbf{F}_\gamma^i)^n}{1 + \Delta t \tilde{c} \left(\sum_j (n_j)^n \tilde{\sigma}_{ij} + \kappa_i \rho \right)}. \quad (\text{B12})$$

APPENDIX C: DUST OPACITIES

The coupling between dust and radiation is a function of both dust grain size and wavelength (or frequency) of incident radiation. A grain of size a and geometric cross-section πa^2 subject to incident radiation at wavelength λ has radiation pressure cross-section $Q(a, \lambda)\pi a^2$, where Q is the dimensionless radiation pressure coefficient $Q(a, \lambda) = Q_{\text{abs}}(a, \lambda) + (1 - g(a, \lambda)) \times Q_{\text{sca}}(a, \lambda)$. Here, Q_{abs} is the contribution from absorption, Q_{sca} is the contribution from

scattering, and $g = \langle \cos \theta \rangle$ is the average cosine of the angle of scattered light.

We use tabulated values of $Q(a, \lambda)$ from Draine & Lee (1984) and Laor & Draine (1993), who present absorption and scattering data for $10^{-3} \leq \lambda/\mu\text{m} \leq 10^3$ and for graphite and silicate grain compositions. Since we do not follow detailed grain chemistry, we average the radiation pressure coefficients for graphite and silicate grains to calculate effective $Q(a, \lambda)$ values. We note that $Q(a, \lambda)$ is related to grain opacity by $\kappa(a, \lambda) = 4Q(a, \lambda)/3a\rho_{\text{gr}}$, where $\rho_{\text{gr}} \approx 2.4 \text{ g cm}^{-3}$ is the internal density of a solid dust grain.

Astrophysical dust grains come in a range of different sizes, and the size distribution is typically defined in terms of a function dn/da , where $dn/da \times da$ gives the number of grains with size in the interval $[a, a + da]$. We assume that dust grains follow a power-law size distribution $dn/da \propto a^{-3.5}$ (Mathis, Rumpl & Nordsieck 1977), with minimum size $a_{\text{min}} = 0.001 \mu\text{m}$ and maximum size $a_{\text{max}} = 1 \mu\text{m}$.

Using these radiation pressure coefficients, the Planck mean opacity κ_{P} at gas temperature T for the frequency interval $[\nu_{\text{min}}, \nu_{\text{max}}]$ averaged over the grain size distribution is given by

$$\kappa_{\text{P}} = \frac{\int_{a_{\text{min}}}^{a_{\text{max}}} \int_{\nu_{\text{min}}}^{\nu_{\text{max}}} B_\nu(T) \kappa(a, c/\nu) \frac{dn}{da} d\nu da}{\int_{a_{\text{min}}}^{a_{\text{max}}} \frac{dn}{da} da \int_{\nu_{\text{min}}}^{\nu_{\text{max}}} B_\nu(T) d\nu}, \quad (\text{C1})$$

and the Rosseland mean opacity κ_{R} is calculated from

$$\frac{1}{\kappa_{\text{R}}} = \frac{\int_{a_{\text{min}}}^{a_{\text{max}}} \int_{\nu_{\text{min}}}^{\nu_{\text{max}}} \frac{\partial B_\nu(T)}{\partial T} \frac{1}{\kappa(a, c/\nu)} \frac{dn}{da} d\nu da}{\int_{a_{\text{min}}}^{a_{\text{max}}} \frac{dn}{da} da \int_{\nu_{\text{min}}}^{\nu_{\text{max}}} \frac{\partial B_\nu(T)}{\partial T} d\nu}, \quad (\text{C2})$$

where $B_\nu(T)$ is Planck blackbody function. By calculating mean opacities over specific frequency ranges, these values can be used in the multifrequency approach detailed in Section 3.2.1.

In practice, we create a lookup table by pre-computing these dust opacities over a range of gas temperatures and during simulations calculate the opacities of individual gas cells by interpolating to the cells' temperatures. Additionally, since the opacities from equations (C1) and (C2) are given in units of area per unit mass of dust, we multiply by a gas cell's dust-to-gas ratio so that opacities are per unit mass of gas.

This paper has been typeset from a \LaTeX file prepared by the author.