

Numerical problems in coupling photon momentum (radiation pressure) to gas

Philip F. Hopkins   and Michael Y. Grudić 

TAPIR, Mailcode 350-17, California Institute of Technology, Pasadena, CA 91125, USA

Accepted 2018 October 24. Received 2018 October 23; in original form 2018 March 19

ABSTRACT

Radiation pressure (RP; or photon momentum absorbed by gas) is important in a tremendous range of astrophysical systems. But we show the usual method for assigning absorbed photon momentum to gas in numerical radiation-hydrodynamics simulations (integrating over cell volumes or evaluating at cell centres) can severely underestimate the RP force in the immediate vicinity around unresolved (point/discrete) sources (and subsequently underestimate its effects on bulk gas properties), unless photon mean free paths are highly resolved in the *fluid* grid. The existence of this error is *independent* of the numerical radiation transfer (RT) method (even in exact ray-tracing/Monte Carlo methods), because it depends on how the RT solution is interpolated back onto fluid elements. Brute-force convergence (resolving mean free paths) is impossible in many cases (especially where UV/ionizing photons are involved). Instead, we show a ‘face-integrated’ method – integrating and applying the momentum fluxes at interfaces between fluid elements – better approximates the correct solution at all resolution levels. The ‘fix’ is simple and we provide example implementations for ray-tracing, Monte Carlo, and moments RT methods in both grid and mesh-free fluid schemes. We consider an example of star formation in a molecular cloud with UV/ionizing RP. At state-of-the-art resolution, cell-integrated methods underestimate the net effects of RP by an order of magnitude, leading (incorrectly) to the conclusion that RP is unimportant, while face-integrated methods predict strong self-regulation of star formation and cloud destruction via RP.

Key words: hydrodynamics – methods: numerical – stars: formation – galaxies: active – galaxies: formation.

1 INTRODUCTION

Radiation-hydrodynamics (RHD) is fundamental to the behaviour of a huge range of astrophysical systems, including galaxies; star-forming giant molecular clouds (GMCs) and star clusters; the interstellar (ISM) circumgalactic (CGM) and intergalactic (IGM) medium; protoplanetary discs (PPDs) and sites of planet formation; stellar structure and evolution; accretion physics, compact objects, and active galactic nuclei (AGNs); supernovae; and more. In most of these systems the ‘radiation pressure forces’ – meaning both optically thin transfer of momentum from the radiation field to the gas via absorbed photon momentum, and optically thick (multiple scattering) local momentum transfer generating a true ‘pressure’ – cannot be neglected. However, because these are complicated, non-linear, turbulent, multiphysics systems, numerical simulations are often necessary, and RHD is a computationally challenging problem for many reasons. As such, a great deal of literature has focused on numerical methods for solving (or approximately solving) the

radiative transfer equation (RTE) on the fly in simulations (for some recent examples, see e.g. Hopkins, Quataert & Murray 2011; Bate 2012; Davis, Stone & Jiang 2012; Kuiper et al. 2012; Wise et al. 2012; Kolb et al. 2013; Rosdahl et al. 2013; Davis et al. 2014; González et al. 2015; Roth & Kasen 2015; Tominaga, Shibata & Blinnikov 2015; Bunttemeyer et al. 2016; Rosen et al. 2017; Kim et al. 2017; Zhang & Davis 2017; Foucart 2018).

But comparably little attention has been paid to how, *given* a solution to the RTE, radiation pressure forces are coupled *onto the gas or fluid* (although see e.g. Lowrie, Morel & Hittinger 1999). In this paper, we show that the most common methods used in the literature can artificially suppress the radiation pressure forces in gas elements nearby unresolved (discrete or point or ‘subgrid’) sources by orders of magnitude, unless the photon mean free paths (MFPs) are very well resolved. This is especially problematic for physical systems with many discrete sources (e.g. star and galaxy formation), or where the source emission region is optically thick and poorly resolved (e.g. AGN and black hole emission), and in those regimes can be particularly important for the large-scale, bulk effects of radiation on the system. But in many of these contexts,

* E-mail: phopkins@caltech.edu

the MFPs of especially ionizing and UV photons in the dense gas around sources (where radiation pressure is most important) can be many orders of magnitude smaller than state-of-the-art numerical resolution.

In Section 2, we briefly review the equations solved for the RTE and radiation pressure forces, and common classes of numerical implementations. Section 3 considers the coupling of photon momentum in the single-scattering limit, and shows how the typical ‘cell-integrated’ approach fails (Section 3.1), and why the resolution requirements to ‘brute force’ convergence with this method are impractical (Section 3.2). We propose instead, in Section 3.3, an alternative, face-integrated approach, which resolves these errors and properly treats the absorbed photon momentum around sources even in the limit where photon MFPs are unresolved. We show how to implement this in various (ray-tracing, Monte Carlo, moments based) RTE methods, and various (fixed grid, mesh-free) hydrodynamic methods (additional details in Appendix A). Section 4 generalizes to multiple scattering, and Section 5 demonstrates how these errors can have a dramatic impact on real astrophysical problems, using a simulation of star formation in a molecular cloud as an example. We conclude in Section 6.

2 REVIEW OF RHD METHODS

Recall, the radiative transfer equation (RTE), to the order of interest (neglecting higher order relativistic and Hubble-flow effects unimportant here), is

$$\left(\frac{1}{c} \frac{\partial}{\partial t} + \hat{\Omega} \cdot \nabla\right) I_\nu = j_\nu - \alpha_\nu I_\nu - \int d\Omega' [I_\nu \sigma_\nu(\Omega, \Omega') - I_\nu(\Omega') \sigma_\nu(\Omega', \Omega)] \quad (1)$$

(Mihalas & Mihalas 1984), where c is the speed of light, $\hat{\Omega}$ is the unit vector in the direction of some differential solid angle $d\Omega$, α_ν and σ_ν are the absorption and scattering coefficients, j_ν is the emissivity, and I_ν is the intensity. Ignoring scattering and sources, this is

$$\left(\frac{1}{c} \frac{\partial}{\partial t} + \hat{\Omega} \cdot \nabla\right) I_\nu = -\rho \kappa_\nu I_\nu = -\frac{I_\nu}{\lambda_{\text{MFP}}}, \quad (2)$$

where ρ is the gas density, κ_ν is the opacity, and $\lambda_{\text{MFP}} \equiv 1/(\rho \kappa_\nu)$ the photon MFP. Within an infinitesimally small differential volume $d\text{Vol} = d^3\mathbf{x}$, the rate of momentum transfer from radiation to gas is

$$\frac{d\mathbf{p}}{d^3\mathbf{x}} = \frac{\partial(\rho \mathbf{u})}{\partial t} = \int d\nu d\Omega \frac{\rho \kappa_\nu I_\nu \hat{\Omega} d\Omega}{c} = \int d\nu \frac{\rho \kappa_\nu \mathbf{F}_\nu}{c}, \quad (3)$$

where $\mathbf{F}_\nu \equiv \int d\Omega I_\nu \hat{\Omega}$ is the flux. We can also define the energy density $e_\nu \equiv c^{-1} \int d\Omega I_\nu$, and source function $S_\nu \equiv \int d\Omega j_\nu$.

A range of methods in the literature attempts to solve equation (1) ‘on the fly’ in simulations, subject to various approximations. The problems highlighted here apply to *all* methods, in principle, but because details of the necessary correction differ, we briefly review broad classes of popular methods for solving the RTE.

2.1 Collisionless (ray-tracing and Monte Carlo) methods

Ray tracing with ‘long characteristics’ (RT-LC) and/or Monte Carlo (MC) methods, which explicitly track I_ν along discrete ‘rays’ or

photon packets from all sources, can (at least in principle) exactly solve the RTE. This is often prohibitively expensive so it is common to make a variety of approximations, e.g. (1) using a limited (finite) number of rays/photon packets; (2) assuming either an infinite-speed-of-light/equilibrium solution (dropping the $c^{-1} \partial I_\nu / \partial t$ term) or ‘reduced speed of light’ ($c \rightarrow \tilde{c} \ll c$) in explicit non-equilibrium solutions, both of which allow for larger time-steps; (3) neglecting scattering (or assuming it is isotropic); (4) neglecting velocity-dependence in κ_ν (for continuum transfer) and/or higher order relativistic terms; (5) limiting to a small number of discrete frequency ‘bins’ by integrating over some finite $\Delta\nu$; (6) replacing direct rays from all sources in the volume to all points with ‘adaptive ray tracing’ (ART; Abel & Wandelt 2002) where rays are ‘split’ or ‘merged’ to sample the domain more/less accurately in some regions, or ‘short-characteristics’ (RT-SC; Olson & Kunasz 1987) where only a fixed set of angles within each cell (which must be interpolated between cells) is used; (7) following direct rays but collapsing non-local shadowing/extinction to a local isotropic region around each source (LEBRON; Hopkins et al. 2014, 2018b); (8) treating absorption and re-emission as elastic scattering (in e.g. ‘implicit Monte Carlo’ methods; Fleck & Cummings 1971).

For our purposes, these approximations are not important – these methods are all similar in the key respects that they explicitly follow I_ν along different angles $\hat{\Omega}$, and (at least in principle) I_ν can have structure below the fluid grid scale.

2.2 ‘Moments’ or ‘fluid’ methods

In moments methods, one integrates the RTE over $\int \hat{\Omega}^n d\Omega$, with $n = 0, 1, \dots$, to take the n th moments. The hierarchy of moment equations never closes (each depends on the next-higher moment), so one truncates the series by adopting an ad hoc closure *ansatz*.

Flux-limited diffusion (FLD; Levermore & Pomraning 1981) closes the hierarchy at $n = 0$, giving $\partial e_\nu / \partial t + \nabla \cdot \mathbf{F}_\nu = S_\nu - \rho \kappa_\nu \tilde{c} e_\nu$, with the *ansatz* $\mathbf{F}_\nu \rightarrow (\lambda_L \tilde{c} / \kappa_\nu \rho) \nabla \cdot (e_\nu \mathbb{D})$ with Eddington tensor $\mathbb{D} = \mathbb{I}/3$ (\mathbb{I} is the identity matrix) and the ‘flux-limiter’ $\lambda_L = 3(2 + \mathcal{F})/(6 + 3\mathcal{F} + \mathcal{F}^2)$ (with $\mathcal{F} \equiv |\nabla e_\nu| / (e_\nu \kappa_\nu \rho)$) chosen to interpolate between free-streaming and isotropic-diffusion-like behaviour based on an estimate of the local optical depth. The ‘optically thin variable Eddington tensor’ (OTVET; Gnedin & Abel 2001) approach is identical to FLD but with $\mathbb{D} = \hat{\mathbf{F}}_\nu^{\text{thin}} \otimes \hat{\mathbf{F}}_\nu^{\text{thin}}$, where $\hat{\mathbf{F}}_\nu^{\text{thin}}$ is the flux that would be present with no obscuration anywhere. The ‘moment-one’ (M1; Levermore 1984) method closes at $n = 1$, keeping the energy equation above and adding the flux equation $\tilde{c}^{-1} \partial \mathbf{F}_\nu / \partial t + \tilde{c} \nabla \cdot \mathbb{P}_\nu = -\rho \kappa_\nu \mathbf{F}_\nu$ (so \mathbf{F}_ν is explicitly evolved), with the *ansatz* $\mathbb{P}_\nu \rightarrow e_\nu \mathbb{D}_\nu$, $\mathbb{D}_\nu = (1 - \chi_\nu) \mathbb{I}/2 + (3\chi_\nu - 1) (\hat{\mathbf{F}}_\nu \otimes \hat{\mathbf{F}}_\nu)/2$, $\chi_\nu = (3 + 4f_\nu^2)/(5 + 2\sqrt{4 - 3f_\nu^2})$, $f_\nu \equiv |\mathbf{F}_\nu| / (e_\nu \tilde{c})$ which again interpolates between free-streaming and isotropic diffusion. Other assumptions for the closures are of course possible but change nothing for our study here.

Independent of the closure, the salient feature of these methods is that each moment is *single-valued* at a given \mathbf{x} – one has made the fluid approximation for the radiation. This gives the well-known result that moments methods do not, in fact, converge to the correct RTE solutions (e.g. antiparallel photon streams will ‘collide’) except in the limit when the optical depths are everywhere infinite. These are popular methods because of their simplicity, but introduce unique complications discussed below.

3 ERRORS AND CORRECTIONS: SINGLE SCATTERING

First consider ‘single scattering’: photons free-stream until they are absorbed (no re-emission or scattering). To demonstrate the numerical issues, consider an idealized problem shown in Fig. 1. A single monochromatic isotropic source with luminosity $L = L_\nu$ sits at position \mathbf{x}_0 , centred on a computational domain which we take to be a regular Cartesian grid with cells of side-length Δx (our conclusions are independent of the mesh geometry but this is convenient). Furthermore, assume that (regardless of RT method used) the RTE is solved *perfectly* at all points \mathbf{x} in the domain.

3.1 The problem: cell-integrated coupling

In most RHD implementations, *regardless* of the method used to solve the RTE, the radiation force on the gas is computed by integrating the momentum of absorbed photons over the volume of a given domain (cell) ‘ a ’:¹

$$(\dot{\mathbf{p}}_v)_a = \int_{\Delta V_a} d^3\mathbf{x} \frac{\rho \kappa_\nu \mathbf{F}_v}{c} \quad (4)$$

so the total momentum $\Delta \mathbf{p}_a = \int_{\Delta v} \int_{\Delta t} d\nu dt (\dot{\mathbf{p}}_v)_a$ is simply added to the gas in a each time-step Δt . For discrete (e.g. Monte Carlo) methods, this integral is given by summing over all absorptions in a given cell: $\Delta \mathbf{p}_a = c^{-1} \sum_j E_\nu^{j, \text{abs}} \hat{\Omega}_j$, where $E_\nu^{j, \text{abs}}$ is the total photon energy absorbed in photon packet/ray j , within domain a , over time-step Δt .

Now consider the case where the photon MFP is unresolved by the *fluid* grid, i.e. $\Delta x \gg \lambda_{\text{MFP}}$. Essentially all the photons should therefore be absorbed inside the single cell a surrounding the source. Equation (4) then trivially evaluates to $(\dot{\mathbf{p}}_v)_a = \Delta \mathbf{p}_a = \mathbf{0}$ because the source emits isotropically.

In the absence of other physics, the system would remain (incorrectly) perfectly static as $t \rightarrow \infty$. The correct solution (with negligible gravity and pressure forces) is that a shell of material moves away from the source, sweeping up gas (leaving an empty cavity) with radial momentum flux $=L/c$ (so the total *radial* momentum $\int d^3\mathbf{x} \rho \mathbf{v} \cdot \hat{\mathbf{r}} = Lt/c$). If L and ρ are constant the shell should expand indefinitely with radius $r(t) = (3Lt^2/2\pi c\rho)^{1/4}$.

3.2 Solution 1: increasing resolution (is not practical)

One solution to the problem above is to increase the resolution of the *hydrodynamic* grid until the photon MFPs are all well resolved,

¹Even simpler than the ‘cell-integrated’ equation (4), some RHD implementations (in e.g. moments methods or some SPH/finite-point methods) adopt a ‘cell-centred’ approach where the acceleration $\mathbf{a}_v \equiv c^{-1} \rho \kappa_\nu \mathbf{F}_v$ is evaluated at the cell centre/particle location \mathbf{x}_a and then assigned to the whole cell/particle. As this is equivalent to evaluating equation (4) using only the cell-centred $\mathbf{F}_v \rightarrow \mathbf{F}_v(\mathbf{x}_a)$ (likewise for ρ, κ), it produces the same errors as cell-integrated approaches. In addition, it is noisier than cell or face-integrated methods, and can violate linear momentum conservation, so we will not discuss this particular case in more detail.

i.e. $\Delta x \ll \lambda_{\text{MFP}}$.² This would eventually converge to the correct behaviour (assuming a perfect RTE solution).

But this is not possible for many real simulations. For example, in simulations of star formation, galaxies, or AGNs (outside the accretion disc), the single-scattering absorption comes primarily from neutral hydrogen absorption of ionizing photons, with $\kappa_\nu \sim 4 \times 10^6 \text{ cm}^2 \text{ g}^{-1}$ (for neutral gas), and/or near-UV with $\kappa_\nu \gtrsim 3000 \text{ cm}^2 \text{ g}^{-1}$, in relatively dense gas around the sources (e.g. $n \sim \rho/m_p \gtrsim 10^2\text{--}10^4 \text{ cm}^{-3}$, in H II regions, or $\sim 10^6\text{--}10^{12} \text{ cm}^{-3}$ in the obscuring ‘torii’ around AGNs). In a uniform grid, this would require

$$\Delta x \ll \lambda_{\text{MFP}} \sim 5 \times 10^{-6} n_4^{-1} \text{ pc} \quad (5)$$

(where $n_4 = n/10^4 \text{ cm}^{-3}$) for ionizing photons in neutral gas. Even for just near-UV photons this gives $\Delta x \ll 0.006 n_4^{-1} \text{ pc}$, so a grid of $\sim 3100^3 \sim 10^{10.5}$ elements is needed for a $\sim 10 \text{ pc}$ -radius GMC (even if the maximum density were capped at $\sim 10^4 \text{ cm}^{-3}$). In Lagrangian or AMR methods where the mass resolution Δm is approximately fixed and spatial resolution is adaptive ($\Delta x \approx (\Delta m/\rho)^{1/3}$), the required mass resolution would be

$$\Delta m \ll \lambda_{\text{MFP}}^2/\kappa_\nu \sim 3 \times 10^{-14} n_4^{-2} M_\odot \quad (6)$$

(for ionizing photons in neutral gas) or $\Delta m \ll 6 \times 10^{-5} n_4^{-2} M_\odot$ (for near-UV). This is wildly beyond the resolution of state-of-the-art simulations.

If we consider absorption of hard radiation in the vicinity of accreting supermassive black holes (where the largest spatial scales of interest are similar, but the densities and opacities are even higher), the resolution-discrepancy only becomes more severe.

3.3 Solution 2: face-integrated coupling

Consider instead a *face-integrated* momentum coupling. Instead of considering only the volume element a , note that a is surrounded by a set of faces³ ab (between domain a and each neighbouring element b) each with an oriented vector area \mathbf{A}_{ab} . Now simply integrate the absorbed flux moving ‘towards’ each face \mathbf{A}_{ab} :

$$(\dot{\mathbf{p}}_v)_{ab} = \int_{\Delta V_a} d^3\mathbf{x} \frac{\rho \kappa_\nu \mathbf{F}_v}{c} \Theta_{ab}(\mathbf{x}, \hat{\mathbf{F}}_v) = \int_{\Delta V_a, \Omega_{ab}} d^3\mathbf{x} \frac{\rho \kappa_\nu \mathbf{F}_v}{c} \quad (7)$$

where $\Theta_{ab}(\mathbf{x}, \hat{\mathbf{F}}_v) = 1$ if the face \mathbf{A}_{ab} is the ‘intercepted face’ (i.e. if \mathbf{A}_{ab} is the first face crossed, along the ray originating at \mathbf{x} in direction $\hat{\mathbf{F}}_v$), and $=0$ otherwise. Equivalently, $\Theta = 1$ if $\hat{\mathbf{F}}_v$ ‘points to’ face ab

²Another possibility would be to artificially decrease the opacity in different bands, so that MFPs are always resolved. This amounts to capping the opacity at $\kappa_\nu^{\text{cap}} \ll 1/(\rho \Delta x)$ (in some cases accomplished by treating cells around sources as $\kappa_\nu = 0$ ‘ghost zones’). This immediately creates a number of problems: (1) Most important, this can lead to photons being absorbed at the wrong physical locations, far further from their sources than they should be. (2) It is very difficult to implement such a prescription in anything but an idealized simulation, without risking allowing photons to escape entirely from dense regions where they should have been absorbed (e.g. ‘undershooting’ the opacity). (3) The central cell around the source is still not correctly being ‘swept up’ in a shell in the problem described above, because the photons are free-streaming out to neighbouring cells. So a shell will form, but only external to a central dense region from which the photons escape – thus, the solution at finite resolution will still not correctly represent the converged solution.

³In ‘mesh-free’ methods (e.g. SPH, FPM) one can always define ‘effective’ faces by reference to equation of motion and point locations: we provide a generic method for this in Appendix A.

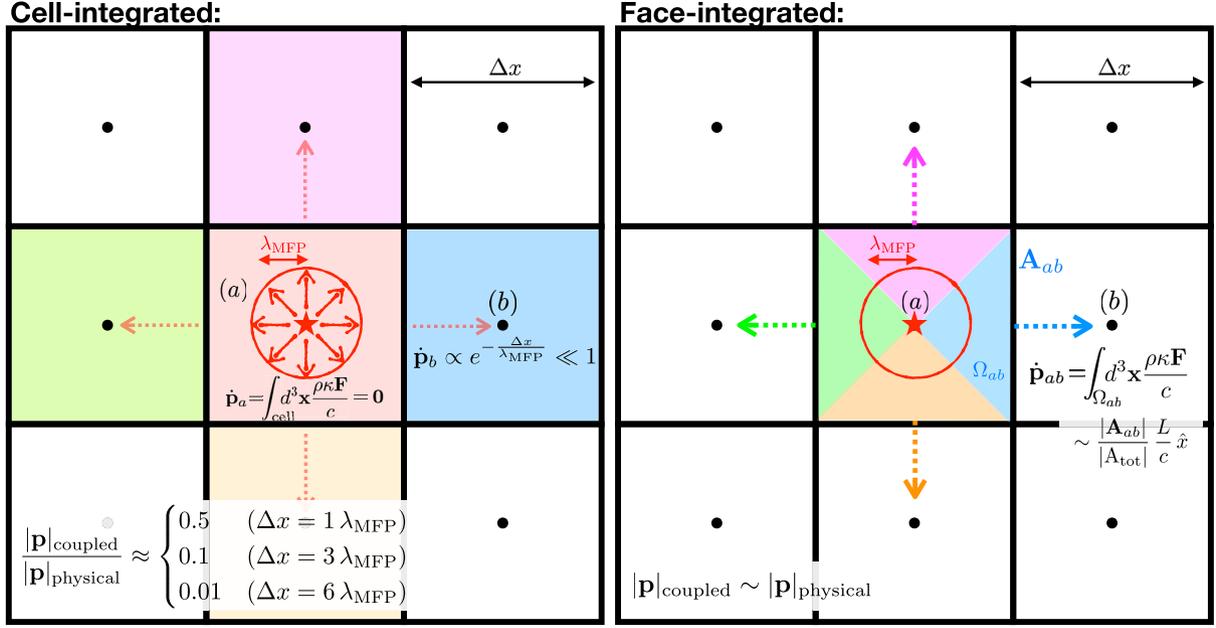


Figure 1. Cartoon illustrating the key errors discussed here (and their solution). A single monochromatic radiation source (red star) sits at the centre (cell a) of the fluid grid (cells with side-length Δx), emitting an isotropic flux $\mathbf{F} = \hat{r} L / (4\pi r^2)$ of photons which are singly scattered in a medium with constant density ρ and opacity κ (photon mean free path $\lambda_{\text{MFP}} = 1/\rho\kappa$). We assume that the RTE (equation 1) is solved perfectly (\mathbf{F} is known exactly everywhere). *Left:* cell-integrated coupling. A momentum flux $\dot{\mathbf{p}}_{a,b}$ is assigned to cell a or b by integrating the absorbed photon momentum density $= \rho\kappa\mathbf{F}/c$ over all points in the volume/domain of a or b (equation 4). If $\Delta x \gg \lambda_{\text{MFP}}$, almost all photons are absorbed in cell a , but since these are emitted (by definition) isotropically, the net $\dot{\mathbf{p}}_a = \mathbf{0}$. Some coherent momentum is imparted to cell b but it is lower by a factor $\sim \exp(-\Delta x/\lambda_{\text{MFP}}) \ll 1$. As a result, the radiation pressure will do nothing to the gas, even in the limit $L \rightarrow \infty$. *Right:* face-integrated coupling. Instead of averaging over volume, we integrate the absorbed photon momentum density in the domain Ω_{ab} flowing towards/through each face A_{ab} between neighbouring fluid elements a and b (equation 7), and assign this as a flux or force between elements (at the face). Even when $\Delta x \gg \lambda_{\text{MFP}}$, the correct momentum ‘outward’ towards each face from the source is now recovered.

from \mathbf{x} within a . Thus, this is the integral over the domain $\Delta V_a[\Omega_{ab}]$ where Ω_{ab} is the range of solid angle subtended by face A_{ab} from \mathbf{x} (i.e. angles where $\Theta(\mathbf{x}, \hat{\Omega}) = 1$). In (exact or approximate) ray-tracing methods where absorption is calculated along rays with fixed $\hat{\Omega}$, it is straightforward to evaluate equation (7) exactly.

In the optically thick limit of particular interest, equation (7) is approximately $\approx c^{-1} \int_{\Delta V_a, \Omega_{ab}} d^3\mathbf{x} d\Omega \rho \kappa_v I_v \hat{\Omega}$. In discrete (Monte Carlo) methods, this becomes $\Delta \mathbf{p}_{ab} = c^{-1} \sum_j \Theta_{ab}(\mathbf{x}_{\text{abs}}, \hat{\Omega}_j) E_j^{i, \text{abs}} \hat{\Omega}_j$.

Since the momentum flux is defined now at *faces*, rather than cell centres, it can be incorporated either as a simple operator-split flux (i.e. cell ‘ b ’ receives a momentum $\Delta \mathbf{p}_{ab}$ from cell ‘ a ’), or as a force in the Riemann problem or momentum equation solved between cells. Note that in general $|(\dot{\mathbf{p}}_v)_{ab}| \neq |(\dot{\mathbf{p}}_v)_{ba}|$, so the *net* momentum flux between cells, $(\dot{\mathbf{p}}_v)_{ab} + (\dot{\mathbf{p}}_v)_{ba}$, is what will ultimately matter. But for a single source in our example problem, $(\dot{\mathbf{p}}_v)_{ba} = \mathbf{0}$, so the force is purely ‘outwards’ from the source (as it should be), while for a perfectly symmetric or homogeneous source distribution, $(\dot{\mathbf{p}}_v)_{ba} = -(\dot{\mathbf{p}}_v)_{ab}$ and the net forces correctly vanish.

Return to our test problem: it is easy to verify that this produces an outward force from the central cell. The momentum flux into each neighbour b follows the exact solution (assuming a spherical shell propagates out from the centre), averaged over the face. While initially the solution cannot, of course, be exactly spherical (owing to the grid geometry), it will rapidly converge to such as the shell expands.

3.3.1 Implementation in moments based methods: additional complications and errors

In moments based methods, two additional, independent and very important errors arise.

(i) Because the ‘mesh’ on which radiative properties are computed is identical to the fluid mesh, the flux $\mathbf{F}_v(\mathbf{x}) = \langle \mathbf{F}_v \rangle_a$ is single valued within a given cell. So we must insert an *explicit* model for $\mathbf{F}_v(\mathbf{x})$ around each source. Since the error here, where important ($\Delta x \gg \lambda_{\text{MFP}}$), is dominated by the cell a surrounding the source, this can be accomplished during the ‘photon deposition’ step. Recall, in moments methods photons must be ‘deposited’ onto the grid around a source, normally via a scalar kernel/weight function: e.g. $(\Delta E_v)_{ab} = L_a \Delta t_a \omega_{ab}$, where $\sum_b \omega_{ab} = 1$. So *during this step*, one can *explicitly* calculate the integral in equation (7) with \mathbf{F}_v from each source, integrated over domain ΔV_a , Ω_{ab} within cell a towards face ab . A detailed example of how to do this numerically is given in Appendix A. Outside the cell hosting a source, it is a much smaller error to simply adopt the cell-centred $\langle \mathbf{F}_v \rangle$ (if $\Delta x \ll \lambda_{\text{MFP}}$, the solution is converged, if $\Delta x \gg \lambda_{\text{MFP}}$, the flux outside the origin cell a is negligible).

(ii) In moments methods, the flux \mathbf{F} does not directly follow from I_v , but is sourced by the *numerical gradient* of e_v . This means that even outside the ‘origin’ cell, it requires several resolution elements to resolve any gradient in e_v (consider, at the origin, the numerical gradient must vanish if the source is isotropic). So even when Δx

$\lesssim \lambda_{\text{MFP}}$, essentially no photon momentum can be transferred to the gas within the central ~ 10 cells in any direction because the numerical gradient is smoothed by a finite discrete ‘kernel’. This is true even for higher order gradient estimators in regular Cartesian-mesh codes (see fig. B1 in Rosdahl et al. 2015). This means a mass $\sim (4\pi/3)(10\Delta x)^3\rho \sim 4200\Delta m$ has its momentum underestimated, which makes the resolution criterion in equations (5) and (6) much more challenging (the mass resolution must be at least >4000 times better, or spatial resolution >20 times better, to converge). One fix to this is to extend the explicit integration method for \mathbf{F}_v in (i) above, to a ~ 10 -cell radius around each source – essentially, performing a ‘mini ray-trace’ in each $r \sim 10\Delta x$ sphere around each source. This is the most accurate method, but is computationally much more complex. A much simpler, albeit significantly less accurate, fix is proposed and adopted by Rosdahl et al. (2015): simply replace $\mathbf{F}_v \rightarrow e_v c \hat{\mathbf{F}}_v$ whenever calculating the photon momentum/radiation pressure terms. This is exact in the single-source, optically thin limit; however, if there are multiple sources with intersecting rays this can overestimate the true momentum transfer (although the error is usually small unless the sources should exactly cancel).

3.4 Numerical example

Fig. 2 demonstrates the errors and their fixes here with a variant of the simple test problem in Fig. 1. We place a single, isotropic, monochromatic point source of constant luminosity L randomly on an effectively infinite grid of uniform Cartesian cells of size Δx with constant density ρ and opacity κ , and calculate the equilibrium flux \mathbf{F} at all points (assuming the light-crossing time is much shorter than all other time-scales in the problem) according to the labelled RHD method. Again we stress the gas properties are ‘frozen’ – we are just solving the RTE. We then calculate the momentum flux $\dot{\mathbf{p}}$ either according to the cell-integrated or face-integrated approach. For the moments methods our ‘face-integrated’ method also implements the additional fixes from Section 3.3.1. We calculate the total ‘outward’/radial momentum flux from the source, $\dot{p}_r^{\text{tot}} = \sum_{ab} (\dot{\mathbf{p}})_a \cdot \hat{\mathbf{r}}_a$ where $\hat{\mathbf{r}}_a$ is the unit vector pointing from source to the cell centre in the cell-integrated case; for the face-integrated case we sum over $(\dot{\mathbf{p}})_{ab} \cdot \hat{\mathbf{r}}_{ab}$ where $\hat{\mathbf{r}}_{ab}$ points to the face-centre. This has the exact solution $\dot{p}_r^{\text{tot}} = L/c$. We compare \dot{p}_r^{tot} for each method, as a function of $\Delta x/\lambda_{\text{MFP}}$ – in these units, the absolute units of the problem scale out completely.

As expected, when the momentum is deposited in a ‘cell-integrated’ fashion, even *exact* ray-tracing or MC methods severely suppress the momentum unless $\Delta x \ll 0.5\lambda_{\text{MFP}}$. The ‘cell-integrated’, ‘default’ moments (M1/FLD) methods fare even worse, requiring $\Delta x \ll 0.05\lambda_{\text{MFP}}$. With ‘face-integrated’ methods, however, the errors are radically reduced, and depend only very weakly on resolution (owing to geometric effects discussed above).

At low resolution ($\Delta x/\lambda_{\text{MFP}} \rightarrow \infty$), even with an exact RT solution, the face-integrated methods converge to a coupled radial momentum $\dot{p}_r^{\text{tot}} \rightarrow 0.84 L/c$, not L/c (i.e. the total radial momentum is ~ 16 per cent lower than the converged solution). This is a real geometric effect.⁴ Imagine a perfectly resolved spherical shell

emerges from a with total radial momentum $=L/c$ – but upon entering cell b , we must *integrate* the total momentum entering b , which means averaging the momentum of the parts of the shell crossing A_{ab} (producing some cancellation of the momentum components transverse to the face). So the solution is still ‘exact’ in that it reflects the converged solution, ‘averaged down’ to the grid scale after propagating through each cell. More important, this difference is small and constant (while cell-integrated methods have \dot{p}_r^{tot} decreasing exponentially as $\Delta x/\lambda_{\text{MFP}} \rightarrow \infty$).

For this simple test problem, approximate-ray (LEBRON/ART) methods give the same result as exact ray-tracing/MC methods. Moments methods always produce larger errors for the reasons in Section 3.3.1 (because we allow the radiation field to reach equilibrium, FLD and M1 are identical here). Even with the face-integrated method and fixes in Section 3.3.1, moments methods are not perfect: at $\Delta x \ll \lambda_{\text{MFP}}$ they converge (much faster than with cell-integrated coupling), and at $\Delta x \gg \lambda_{\text{MFP}}$ the solution occurs entirely ‘within one cell’ so they are identical to the exact methods, but when $\Delta x \sim \lambda_{\text{MFP}}$, the gradient errors noted in Section 3.3.1 somewhat suppress the coupled momentum.

4 EXTENSION TO MULTIPLE SCATTERING

Briefly, consider multiple scattering: photons are absorbed and re-emitted (or scattered) many times before escaping/being destroyed. Where this is of interest, opacities are usually low $\kappa \sim \kappa_0 \text{ cm}^2 \text{ g}^{-1}$: e.g. in IR dust re-processing ($\kappa_0 \sim 1 - 10$), free-electron scattering ($\kappa_0 \sim 0.4$), X-ray metal-line absorption ($\kappa_0 \sim 2$ at the Fe edge). So resolution requirements are less extreme: $\Delta x \ll \lambda_{\text{MFP}} \sim 20\kappa_0^{-1} n_4^{-1} \text{ pc}$ or $\Delta m \ll \lambda_{\text{MFP}}^2/\kappa_v \sim 2 \times 10^6 \kappa_0^{-3} n_4^{-2} M_\odot$. But even these are not always met; moreover multiple-scattering can also be important in resonance lines where κ is much larger.

It is trivial to see that in this limit once again, a ‘cell-integrated’ coupling of photon momentum to gas (equation 4) gives essentially vanishing radiation pressure ($(\dot{\mathbf{p}}_v)_{ab} \rightarrow \mathbf{0}$ in the cell around a source when $\Delta x \gg \lambda_{\text{MFP}}$ (precisely the opposite of the correct behaviour). But it is also easy to verify that the ‘face-integrated’ coupling (equation 7) gives the correct solution in this limit.

We can repeat the numerical experiment in Fig. 2, for example, but instead of single scattering assume multiple scattering with a grey (constant) opacity κ , take $\lambda_{\text{MFP}} \rightarrow 0$ compared to resolved scales, and assume perfect re-emission (all absorbed energy is re-radiated). In that case the exact solution is now $\dot{p}_r^{\text{tot}}(<r) = \tau(<r) L/c$, where $\dot{p}_r^{\text{tot}}(<r)$ is the total radial momentum flux integrated over all cells within a radius r , and $\tau(<r) = \rho \kappa r$ is the optical depth out to that radius. If we then compare $\dot{p}_r^{\text{tot}}(<r)$ to the correct solution $\tau(<r) L/c$, as a function now of $\Delta x/r$ (since $\lambda_{\text{MFP}} \rightarrow 0$ by assumption), we obtain essentially identical results to those shown in Fig. 2 for single scattering.

5 EFFECTS IN AN EXAMPLE PROBLEM: STAR CLUSTER FORMATION

To explore this in an astrophysical context, we consider as an example a simulation of star cluster formation in an individual GMC following Grudić et al. (2018). The simulations use the code GIZMO

faces) means the mean face-integrated momentum as $\Delta x/\lambda_{\text{MFP}} \rightarrow \infty$ is $\dot{p}_r^{\text{tot}} \rightarrow 0.94 L/c$.

⁴In fact the rectilinear Cartesian mesh we adopt gives nearly the worst-case scenario for this geometric effect. If we instead assume a glass configuration of mesh-generating points with a Voronoi tessellation or the mesh-free MFM/MFV methods from Hopkins (2015) used to calculate the faces, the larger number of faces (cells are closer to regular polyhedra with ~ 16

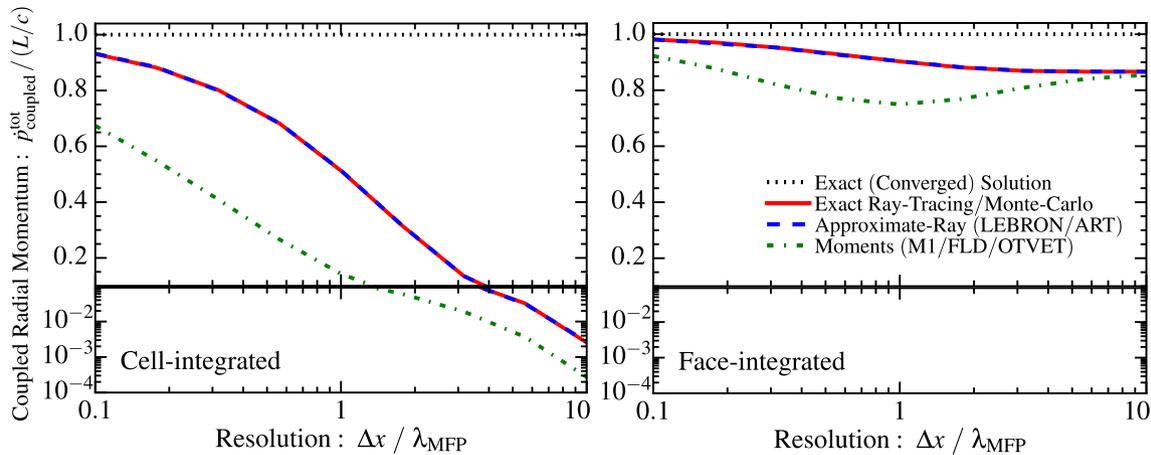


Figure 2. Quantitative demonstration of the error from the simple test in Fig. 1. We take that set-up, place the single source at a random position on the regular Cartesian grid (averaging over ~ 1000 positions), and solve the RTE (waiting for it to come into its equilibrium solution, with the gas properties fixed) using one of three methods: (1) an ‘exact’ (e.g. infinite-resolution ray-tracing or Monte Carlo) method, where the flux is exact at all points; (2) an approximate ray-tracing method (using the ray tracing in LEBRON or ART gives nearly identical results here, both to each other and to exact RT methods); (3) a ‘moments’ method (M1 here, but FLD and OTVET give identical results). We then calculate the total radial momentum flux, $\sum_a \dot{p}_a$ summed over all cells (or faces) according to that RTE solution and either the cell-integrated (equation 4; *left*) or face-integrated (equation 7; *right*) coupling methods. With cell-integrated methods, the true RP force is suppressed by orders of magnitude, even given a perfect RT solution, if $\Delta x \gtrsim \lambda_{\text{MFP}}$ (in moments schemes, additional errors require $\sim 10 \times$ higher resolution; see Section 3.3.1). Face-integrated methods give robust answers at all resolution levels (all solutions asymptote to ≈ 85 per cent of the full radial momentum as $\Delta x \rightarrow \infty$, but this is the correct geometric effect from the finite number of faces in a grid cell).

(Hopkins 2015),⁵ with the meshless-finite-mass (MFM) Godunov MHD solver, and fully adaptive/Lagrangian force-softening. We initialize a solar-metallicity cloud of mass $10^6 M_{\odot}$ and size 50 pc, with intentionally low resolution (fixed $\Delta m = 8 M_{\odot}$), and an initially turbulent velocity and magnetic field spectrum with virial parameter of unity and mean plasma $\beta = 100$, and evolve it including self-gravity, ideal MHD, radiative cooling (with the cooling curves from Hopkins et al. 2018b from 10 to 10^{10} K, assuming a universal Milky Way like radiation field for heating/cooling), and star formation into sink particles in locally self-gravitating, Jeans-unstable, converging flows (all details in Grudić et al. 2018). While the simulations in Grudić et al. (2018) included stellar mass-loss, SNe (Ia and II), and multiband RHD with single and multiscattering RP and photoheating, here we disable these terms and consider *only* monochromatic single-scattering RP (to clearly isolate the effects of interest). To further simplify, each star particle (once formed) is assigned a constant light-to-mass ratio $L/M = 1100 L_{\odot}/M_{\odot}$ and we adopt a constant $\kappa = 2000 \text{ cm}^2 \text{ g}^{-1}$ (approximately appropriate for near-UV luminosities and dust opacities around young populations).

We solve RHD using one of two approximate methods: (1) the ray-based LEBRON (for implementation details, see Hopkins et al. 2018b), or (2) moments based M1 (following Rosdahl et al. 2013). For each, we consider both ‘cell-integrated’ and ‘face-integrated’ implementations as Figs 1 and 2 (with the ‘face-integrated’ M1 also including the Rosdahl et al. 2015 fix to the gradient error discussed in Section 3.3.1). We also run a test with no radiation for comparison. We emphasize that both of these methods will have large errors here: LEBRON is exact in the optically thin limit, M1 in the optically thick multiple-scattering limit, so neither solves the RTE exactly in this problem where there are both thin and thick single-scattering regimes. So since their errors and regimes of applicability are quite

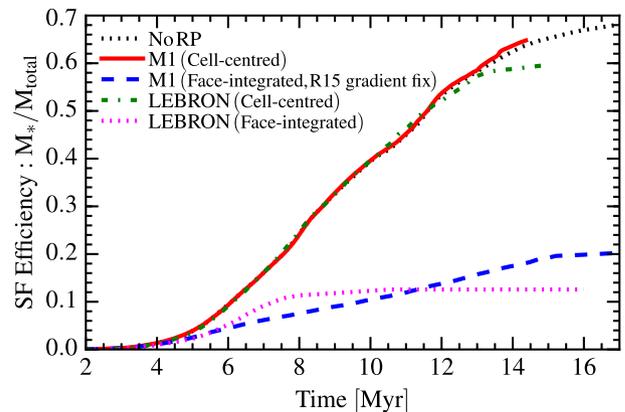


Figure 3. Example problem (see Section 5), illustrating effects of the errors discussed here. We simulate a single GMC collapsing and forming a star cluster, including self-gravity, star formation, and single-scattering RP (each star/sink, once formed, is assigned a constant mono-chromatic luminosity used to solve the RTE assuming single scattering). We compare two approximate methods to solve the RTE, chosen to have opposite regimes of validity: (1) the ray-based LEBRON (exact in the optically thin single-scattering limit), and (2) the moments based M1 (exact in the optically thick *multiple*-scattering limit). For each we compare cell-centred and face-integrated formulations. We plot the stellar mass (as a fraction of total cloud mass) versus time. Without RP, most of the mass turns into stars (no other feedback, e.g. photoionization heating or SNe, is included here). With both RT methods, the cell-centred formulations suppress the RP strongly (it has almost no effect), while the more accurate face-integrated couplings predict a large effect of RP suppressing SF.

different, it is useful to consider how both methods are influenced by these errors.

Fig. 3 plots the ‘star formation efficiency’ (SFE; fraction of the mass turned into stars): over a few dynamical times, gas collapses and rapidly turns into stars until it is exhausted or expelled by RP. In both cell-centred runs, the RP is strongly suppressed at this resolu-

⁵<http://www.tapir.caltech.edu/~phopkins/Site/GIZMO.html>

tion and has no effect on the SFE. In both face-centred runs, the SFE is suppressed by a significant factor – i.e. RP has a dramatic effect on the cluster evolution.⁶ This may explain several apparently discrepant results in the literature. For example, the default Grudić et al. (2018) simulations adopted the ‘face-integrated’ LEBRON method, while a similar study by Raskutti, Ostriker & Skinner (2016) used the ‘cell-integrated’ (and otherwise ‘uncorrected’) M1 method. As expected from our test here, Raskutti et al. (2016) find an order-of-magnitude higher SFE (much weaker effects of radiation pressure), for otherwise similar clouds.

It is worth noting that the errors here can be ‘hidden’ in numerical studies which include e.g. *both* RP and other feedback (e.g. photoionization or stellar winds or SNe). Consider: if some other (non-RP) physics first creates a low-opacity ‘bubble’ around the source (e.g. stellar winds sweeping gas into a shell surrounding a very low density cavity, or SNe destroying dust, or photoionization generating a Stromgren sphere that eliminates the neutral hydrogen opacity inside the sphere and/or pushes gas again into a shell), then one can have an MFP *immediately around the source* which becomes resolved ($\Delta x \ll \lambda_{\text{MFP}}$), even though when the photons eventually encounter a distant high-opacity ‘shell’ (or neutral gas or dust) they are absorbed in a thin layer ($\Delta x \gg \lambda_{\text{MFP}}$ ‘in the shell’). In that regime, as long as the shell radius is well resolved, the errors here are not large (the momenta do not cancel inside a single cell) – so simulations like the GMC study in Fig. 3 which attempt to be more ‘full physics’ may have the inadvertent benefit of enabling better resolution and therefore coupling of the RP forces (making the differences less stark). However, even in these cases, the *early* RP effects (before the shell/Stromgren sphere reaches large radii $r \gg \Delta x$) will still be severely suppressed (if $\Delta x \gg \lambda_{\text{MFP}}$). Moreover, if one takes one of these simulations and ‘turns off’ the other sources of feedback, then RP alone would be dramatically suppressed, leading to the incorrect conclusion that it is unimportant. And more generally, one cannot always rely on these secondary mechanisms to ensure the RP forces ‘become resolved’. In sufficiently dense clouds, photoionization will not generate an expanding bubble or large/well-resolved H II regions, and in galaxy-scale or cosmological-scale simulations even quite large individual H II regions are usually unresolved. And in different physical circumstances, e.g. outflows from AGNs, it is not obvious whether any other process necessarily ‘pre-generates’ a cavity for the RP to act within.

6 CONCLUSIONS

We have shown that ‘cell-integrated’ coupling of absorbed photon momentum (radiation pressure) in radiation-hydrodynamics treatments – the most common approach used in the literature – severely underestimates the true momentum flux around sources, unless the photon MFPs are well resolved in the *hydrodynamic* grid (fluid spatial resolution $\Delta x \ll \lambda_{\text{MFP}}$ or mass resolution $\Delta m \ll \lambda_{\text{MFP}}^2/\kappa$). But the required resolution is often impossible – for example, in simulations of star or galaxy formation, proper treatment of the UV and ionizing photons (most of the single-scattering radiation

pressure) formally requires mass resolution $\Delta m \ll 10^{-13} M_{\odot}$ (or $\ll 10^{-17} M_{\odot}$ in FLD/M1 ‘moments’ methods).

Fortunately, we show that adopting a ‘face-integrated’ coupling – in which the momentum flux is integrated towards each hydrodynamic face, and treated as part of the usual fluxes, instead of being integrated over the entire cell volume and added to the cell-centred momentum – resolves these errors. We show that this produces good convergence (provided the radiative transfer equation is properly solved) even when the gas grid fails to resolve the photon MFPs ($\Delta x \gg \lambda_{\text{MFP}}$).

We stress that the errors identified here, and their fixes, are *independent* of the radiation methods, even if the RTE is solved *perfectly* (with e.g. exact Monte Carlo or ray-tracing methods). The error arises not from the solution of the RTE, but from *how the absorbed photon momentum is assigned to the gas*. As such the errors and their fixes are extremely general. We do also identify some additional (related but distinct) errors in common implementations unique to moments based methods (FLD/OTVET/M1) for RHD and discuss potential fixes. We also demonstrate that the errors and fixes are qualitatively independent of the hydrodynamic method – we outline how to implement these for both grid-based codes and mesh-free (e.g. SPH) methods.

We show that the erroneous formulation can fail entirely to capture the effects of radiation pressure (underestimating the true photon momentum by orders of magnitude) when $\Delta x \ll \lambda_{\text{MFP}}$. We illustrate the effects of this in a fully non-linear example of state-of-the-art simulations of star cluster formation, where we show an improper cell-integrated treatment of the single-scattering photon-momentum leads to incorrect conclusions. Specifically (in the test chosen), with a correct treatment of the radiation pressure (in either M1 or ray-based methods), radiation pressure from massive stars rapidly disrupts the star-forming GMC and greatly suppresses the star formation efficiency. With an incorrect (cell-integrated) treatment (again in both M1 and ray-based methods) the radiation pressure (erroneously) does very little to the cloud. Obviously similar extensions apply to any simulations of unresolved point-like sources, for example most simulations of the effects of AGN radiation on regions larger than the emitting disc (e.g. the broad-line or torus or narrow-line regions).

We stress that simple ‘resolution tests’ (running the same ICs with increasing resolution) applied to this problem will not reveal the problem with cell-integrated approaches, because the behaviour will change sharply (and converge to the correct solution) only when the mass resolution Δm becomes smaller than the ludicrously small value above (which would require $\gtrsim 10^{16}$ particles in our example star-cluster simulation). At achievable resolution, one will instead see ‘false convergence’ until this threshold is reached.

ACKNOWLEDGEMENTS

Support for PFH and MYG was provided by an Alfred P. Sloan Research Fellowship, NSF Collaborative Research Grant #1715847 and CAREER grant #1455342. Numerical calculations were run on the Caltech compute cluster ‘Wheeler’, allocations from XSEDE TG-AST130039 and PRAC NSF.1713353 supported by the NSF and NASA HEC SMD-16-7592.

REFERENCES

- Abel T., Wandelt B. D., 2002, *MNRAS*, 330, L53
Bate M. R., 2012, *MNRAS*, 419, 3115

⁶Fig. 3 runs with ‘fixed’ schemes give roughly the expected result from simple analytic arguments: comparing the strength of the total RP force $\sim L/c \sim (L/M) M_*/c$ to gravity $\sim G M_{\text{cloud}}^2/R^2$ shows that RP should overwhelm gravity on the cloud scale when $M_* \gtrsim 0.1 M_{\text{cloud}}$.

- Buntemeyer L., Banerjee R., Peters T., Klassen M., Pudritz R. E., 2016, *New Astron.*, 43, 49
- Davis S. W., Stone J. M., Jiang Y.-F., 2012, *ApJS*, 199, 9
- Davis S. W., Jiang Y.-F., Stone J. M., Murray N., 2014, *ApJ*, 796, 107
- Fleck J. A., Jr., Cummings J. D., 1971, *J. Comput. Phys.*, 8, 313
- Foucart F., 2018, *MNRAS*, 475, 4186
- Gnedin N. Y., Abel T., 2001, *New Astron.*, 6, 437
- González M., Vaytet N., Commerçon B., Masson J., 2015, *A&A*, 578, A12
- Grudić M. Y., Hopkins P. F., Faucher-Giguère C.-A., Quataert E., Murray N., Kereš D., 2018, *MNRAS*, 475, 3511
- Hopkins P. F., 2015, *MNRAS*, 450, 53
- Hopkins P. F., Quataert E., Murray N., 2011, *MNRAS*, 417, 950
- Hopkins P. F., Keres D., Onorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, *MNRAS*, 445, 581
- Hopkins P. F. et al., 2018a, *MNRAS*, 477, 1578
- Hopkins P. F. et al., 2018b, *MNRAS*, 480, 800
- Kim J.-G., Kim W.-T., Ostriker E. C., Skinner M. A., 2017, *ApJ*, 851, 93
- Kolb S. M., Stute M., Kley W., Mignone A., 2013, *A&A*, 559, A80
- Kuiper R., Klahr H., Beuther H., Henning T., 2012, *A&A*, 537, A122
- Levermore C. D., 1984, *J. Quant. Spectrosc. Radiat. Transfer*, 31, 149
- Levermore C. D., Pomraning G. C., 1981, *ApJ*, 248, 321
- Lowrie R. B., Morel J. E., Hittinger J. A., 1999, *ApJ*, 521, 432
- Mihalas D., Mihalas B. W., 1984, *Foundations of Radiation Hydrodynamics*, Oxford Univ. Press, Oxford
- Olson G. L., Kunasz P. B., 1987, *J. Quant. Spectrosc. Radiat. Transfer*, 38, 325
- Raskutti S., Ostriker E. C., Skinner M. A., 2016, *ApJ*, 829, 130
- 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., Oishi J. S., Lee A. T., Klein R. I., 2017, *J. Comput. Phys.*, 330, 924
- Roth N., Kasen D., 2015, *ApJS*, 217, 9
- Tominaga N., Shibata S., Blinnikov S. I., 2015, *ApJS*, 219, 38
- Wise J. H., Abel T., Turk M. J., Norman M. L., Smith B. D., 2012, *MNRAS*, 427, 311
- Zhang D., Davis S. W., 2017, *ApJ*, 839, 54

APPENDIX A: EFFECTIVE FACE CONSTRUCTION AND COUPLING: EXAMPLE METHODS

As discussed in the text, our face-integrated methods require defining effective faces \mathbf{A}_{ab} around any source and integrating the absorbed momentum in the direction of each face (equation 7). Regardless of the hydrodynamic grid, in moments based methods (FLD/OTVET/M1) we must explicitly construct this ‘effective grid’ on a subcell scale in order to obtain correct solutions (see Section 3.3.1). And in hydrodynamic methods where the grid is anything but regular (e.g. AMR or moving-mesh or MFM/MFV methods), constructing faces and solving the relevant integrals for an arbitrary geometric distribution of mesh-generating points and source location is highly non-trivial. Moreover, in SPH or finite-point methods (FPM) no explicit grid exists. We therefore present here a general method that can be used for *any* set of mesh-generating points (or ‘particles’ in SPH/FPM), to construct a set of effective faces and solve equation (7) in the immediate vicinity of a radiation source.

(i) Every time-step Δt_a , for each source a (at position \mathbf{x}_a), determine L_v^a and $\Delta E_v^a = L_v^a \Delta t_a$.

(ii) Identify ‘neighbouring’ gas surrounding the source. In regular-mesh methods this is straightforward; in mesh-free methods one identifies all gas elements within a ‘search radius’ H_a centred

on a ($|\mathbf{x}_{ba}| \equiv |\mathbf{x}_b - \mathbf{x}_a| < H_a$), and all elements for which a is within their search radius ($|\mathbf{x}_{ba}| < H_b$).⁷

(iii) Construct ‘effective faces’ \mathbf{A}_{ab} around the source. There are many possible choices for this, corresponding to different hydrodynamic methods. For example, a Voronoi tessellation (using a and all b as mesh-generating points). For the kernel-volume decomposition used in MFM/MFV methods (Hopkins 2015): $\mathbf{A}_{ba} \equiv \bar{n}_a^{-1} \bar{\mathbf{q}}_b(\mathbf{x}_a) + \bar{n}_b^{-1} \bar{\mathbf{q}}_a(\mathbf{x}_b)$, where $\bar{\mathbf{q}}_b(\mathbf{x}_a) \equiv \mathbf{E}_a^{-1} \cdot \mathbf{x}_{ba} W(\mathbf{x}_{ba}, H_a)$ and $\mathbf{E}_a \equiv \sum_c (\mathbf{x}_{ca} \otimes \mathbf{x}_{ca}) W(\mathbf{x}_{ca}, H_a)$. In SPH, $\mathbf{A}_{ba} = [\bar{n}_a^{-2} \partial W(|\mathbf{x}|_{ba}, H_a) / \partial |\mathbf{x}|_{ba} + \bar{n}_b^{-2} \partial W(|\mathbf{x}|_{ba}, H_b) / \partial |\mathbf{x}|_{ba}] \hat{\mathbf{x}}_{ba}$.

(iv) Solve the integral in equation (7), towards each face \mathbf{A}_{ab} . Assume ρ and κ_v are constant within each (subcell scale) subdomain, and take \mathbf{x}_a to be the coordinate origin. Then for single-scattering, $\mathbf{F}_v^a(r) = [L_v / (4\pi r^2)] \exp(-r/\lambda_{\text{MFP}})$ along each ray from the source (and we can add all sources independently). The integral can then be expressed as a set of vector weights $\bar{\mathbf{w}}_{ba}$:

$$(\dot{\mathbf{p}}_v)_{ab} \approx f_{\text{abs}} \frac{L_v^a}{c} \bar{\mathbf{w}}_{ba} \quad (\text{A1})$$

$$\bar{\mathbf{w}}_{ba} \equiv \frac{\mathbf{w}_{ba}}{\sum_c |\mathbf{w}_{ca}|} \quad (\text{A2})$$

$$\mathbf{w}_{ba} \equiv \omega_{ba} \sum_{+, -} \sum_{\alpha} (\hat{\mathbf{x}}_{ba}^{\pm})^{\alpha} (f_{\pm}^{\alpha})_a \quad (\text{A3})$$

$$(f_{\pm}^{\alpha})_a \equiv \left\{ \frac{1}{2} \left[1 + \left(\frac{\sum_c \omega_{ca} |\hat{\mathbf{x}}_{ca}^{\mp}|^{\alpha}}{\sum_c \omega_{ca} |\hat{\mathbf{x}}_{ca}^{\pm}|^{\alpha}} \right)^2 \right] \right\}^{1/2} \quad (\text{A4})$$

$$\omega_{ba} \equiv \frac{1}{2} \left(1 - \frac{1}{\sqrt{1 + (\mathbf{A}_{ba} \cdot \hat{\mathbf{x}}_{ba}) / (\pi |\mathbf{x}_{ba}|^2)}} \right) \approx \frac{\Delta \Omega_{ba}}{4\pi}, \quad (\text{A5})$$

where the $\hat{\mathbf{x}}_{ca}^{\pm}$ are the positive or negative (singly signed) projection vectors:

$$\hat{\mathbf{x}}_{ba} \equiv \frac{\mathbf{x}_{ba}}{|\mathbf{x}_{ba}|} = \sum_{+, -} \hat{\mathbf{x}}_{ba}^{\pm} \quad (\text{A6})$$

$$(\hat{\mathbf{x}}_{ba}^+)_{\alpha} \equiv |\mathbf{x}_{ba}|^{-1} \text{MAX}(\mathbf{x}_{ba}^{\alpha}, 0) \Big|_{\alpha=x, y, z} \quad (\text{A7})$$

$$(\hat{\mathbf{x}}_{ba}^-)_{\alpha} \equiv |\mathbf{x}_{ba}|^{-1} \text{MIN}(\mathbf{x}_{ba}^{\alpha}, 0) \Big|_{\alpha=x, y, z} \quad (\text{A8})$$

and, for single-scattering, $f_{\text{abs}} \approx 1 - \exp(-|\mathbf{x}_{ba}|/\lambda_{\text{MFP}})$. For multiple-scattering from a single source, the expression above is identical up to f_{abs} – if we assume a grey opacity and perfect re-radiation (within the single cell subdomain) then $f_{\text{abs}} \approx |\mathbf{x}_{ba}|/\lambda_{\text{MFP}}$. These expressions are derived in Hopkins et al. (2018a, up to the trivial addition here of calculating f_{abs}); while non-trivial they have three key properties. (1) They maintain manifest conservation of linear momentum. (2) They give fluxes that are statistically isotropic in the rest frame of the source (i.e. the coupled momenta are not numerically biased in any particular direction, regardless of the position of the source within the cell or the position/distribution of mesh-generating points/particles). (3) They approximate, as closely as possible without a computationally expensive exact numerical quadrature, the exact integral of equation (7).

⁷The most common approach to define H_a is via a target ‘neighbour number’ $N_{\text{ngb}} = (4\pi/3) H_a^3 \bar{n}_a(H_a)$, where $\bar{n}_a = \sum W(\mathbf{x}_{ba}, H_a)$ and W is an appropriate kernel function which integrates (over volume) to unity.

(v) Verify conservation: one should ensure that total momentum and energy are conserved, regardless of how equation (7) is solved. For an isotropic source in a uniform density field, this means ensuring the *total linear* momentum coupled vanishes (the momentum-coupling is symmetric). If one solves equation (7) *exactly* at all points, *and* has an exact solution for \mathbf{F}_v at all points, *and* has a set of faces which forms an exactly closed convex hull, then this is guaranteed mathematically, but numerically these conditions are not usually all satisfied. Here, this step is included above: the f_{\pm}^{α} terms in equation (A4) are numerically evaluated in a first-loop over the neighbours, to calculate w_{ba} , and represent the ‘correction’ (vector re-normalization) needed to guarantee conservation (one can verify that if the conditions above were met, then one would always have $f_{\pm}^{\alpha} = 1$ exactly).

(vi) Couple the momentum to the gas fully conservatively, either as direct fluxes/forces into each neighbour b , or in a Riemann problem (just adding the flux to the ‘other’ fluxes evaluated at the face in that step).

The procedure above resolves the primary issues in the text, on arbitrary grid, for arbitrary source positions. Note that we have calculated the face-integrated coupling in the vicinity of the source. Far away from the source, one can extend a method like this if desired. However, far from a source, the direction of the flux is not strongly varying within a single cell – so simply using the cell-centred average flux (equivalently, assuming the net flux $\mathbf{F}_v(\mathbf{x})$ at all positions \mathbf{x} inside a cell b , far from any sources, points in approximately the same direction) is not a significant source of error in this regime.

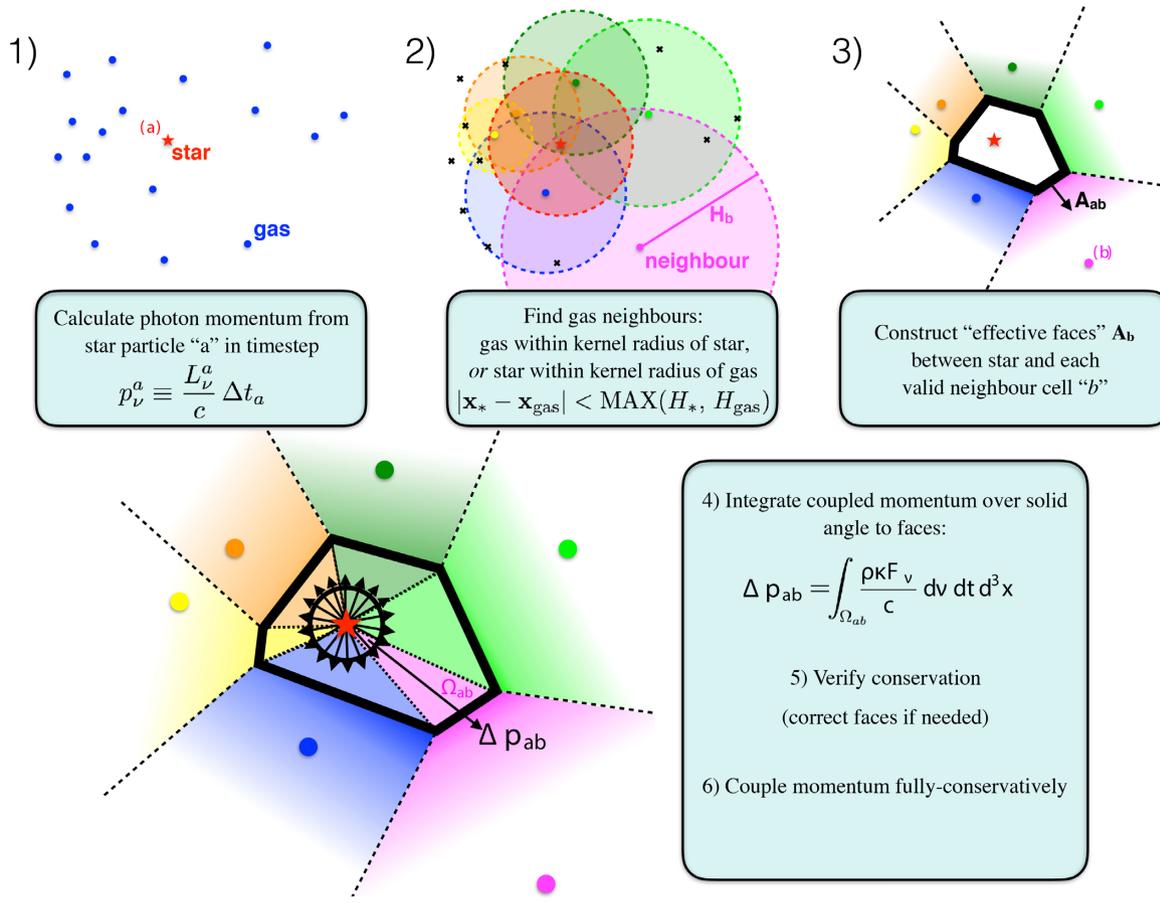


Figure A1. Cartoon illustrating a ‘face-integrated’ numerical algorithm for coupling radiative feedback in arbitrary mesh or particle-based codes (Section A). (1) Determine whether the source a (e.g. a star or black hole) is active in a given time-step, and calculate its luminosity in each waveband. (2) Identify valid interacting neighbours b for the source (fluid elements within a search radius H_a from the source, or for which the source is within their search radius H_b), taking care to cover all directions from the source. (3) Construct the ‘effective faces’ A_{ab} of the interacting fluid elements, as seen by the source, via e.g. a Voronoi tessellation or kernel-volume decomposition. (4) Integrate the absorbed photon momentum over the solid angle subtended by each face A_{ab} (i.e. in the domain Ω_{ab}), according to equation (7), to determine the momentum flux $(\dot{\mathbf{p}}_\nu)_{ab}$ associated with each face. (5) Verify the fluxes maintain machine-accurate momentum conservation and statistical isotropy; if not, correct the faces to ensure this is true (implicit in the algorithm here). (6) Couple the momentum flux at faces in a fully conservative manner, either as an operator-split force or added to the force on the face from the Riemann problem: e.g. $\mathbf{p}_b = \mathbf{p}_b(t) + \sum_{\text{sources}} \sum_{\text{faces}} \Delta \mathbf{p}_{ab}$ with $\Delta \mathbf{p}_{ab} = \int_{\Delta\nu} \int_{\Delta t} (\dot{\mathbf{p}}_\nu)_{ab} d\nu dt$.

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